A research monograph

The Geometry and Arithmetic of Iterated Theta Lifts

A reconstruction and study of the shortened Alpöge–Claude proof of Dirichlet GRH

Expanded reference edition

Working mathematical manuscript – November 2026

Scope. This fourth exposition aims at a largely self-contained research-level account of the shortened argument. It integrates the preceding monograph with substantial additional mathematical proofs, including classical analytic foundations, the global Weil data, CM reciprocity and unfolding, the quantitative filling argument, the integral period and congruence comparisons, and crystalline lifting estimates. It assumes standard foundational theorems in the pertinent disciplines. The audit at the end distinguishes proved arguments, applications of classical results, and the remaining specialized statements imported from the Lean source.

Attribution. The original mathematical argument is attributed by the supplied artifact to Levent Alpöge with Claude. The shortened Lean development is the primary source of this exposition. No claim of independent kernel verification, new authorship of the GRH proof, or completion of every remaining paper-proof obligation is made. The shortened Lean code on which this exposition is based is credited to Luke Sernau; see his repository, `https://github.com/lsernau/grh-lean`.

AI disclaimer. The text of this document was written by AI, ChatGPT-6.2 Sol, and has not been checked by a human.

Preface: the purpose of this monograph

The compressed formal proof of the generalized Riemann hypothesis for Dirichlet L-functions presents a very unusual task for an expository work. There is a clearly identifiable central argument, but the mathematics needed to implement its quantitative estimates comes from several distinct areas. A reader who knows the analytic theory of L-functions will not necessarily be familiar with the integral model of a unitary theta lift. A specialist in p-adic Hodge theory may have little reason to know the exact Bessel normalization in a trace calculation. And a specialist in automorphic forms may be familiar with the norm formulas of the theta correspondence, but not with the torsion or crystalline comparison estimates that are needed uniformly over all coefficient primes. The difficulty is not merely that there are many auxiliary lemmas; the lemmas belong to different mathematical languages and are linked in numerical ways that must survive translation.

The 49-page first exposition supplied a guide to the mathematical architecture. The 162-page specialist version then expanded many of the central estimates, organized the proofs of the analytical endgame, and gave the interfaces to the remaining arithmetic statements. The 333-page third exposition supplied systematic technical companion chapters. The present volume is the fourth and most extensive exposition. It preserves that logical spine but extends it to include proofs of more of the hypotheses of the earlier companion chapters. In particular we reconstruct the elementary analytic inputs, rather than invoking uniform prime distribution, and develop the normalized global data, the exact finite and local computations, and the specialized geometric and arithmetic reductions of the main argument. The order is chosen to make the relationship between the conceptual proof and its technical realization visible.

For whom the book is written. We assume a graduate-level acquaintance with algebra, number fields, linear algebra and complex analysis. A reader with experience in automorphic forms, representation theory, arithmetic geometry or Galois representations should be able to follow the chapters of their own specialty and to understand the role of the others. In a largely self-contained mathematical monograph it is reasonable to take major classical theorems, such as global reciprocity, classical Hermitian mass formulas, and standard foundations of p-adic Hodge theory, from the literature rather than re-establish the entire theory. The distinction is that specialized strengthenings used here, especially exact integral or quantitative variants, must be separately justified. Where this has not been possible, the exact missing assertion is identified rather than silently treated as a textbook theorem.

What constitutes a proof here. The label *formal-source input* marks an assertion taken from a specified theorem in the Lean repository whose conventional proof, including its precise hypotheses and constants, is not completely reconstructed in this manuscript. Such an input is not a new conjecture or an additional axiom of the formal development. It is an outstanding

obligation for a proof that a mathematician could verify without using that development as a certificate. In the passages presented as ordinary proofs, the argument is given from the stated hypotheses and inputs. The geometric constructions and some of the most technical integral and uniform estimates are still mathematical work requiring further checking. The final chapter lists them by theme. The aim is to make genuine progress in conventional exposition without disguising a gap by calling a result "standard".

Reading strategy. Read the introductory chapter, the analytic reduction and the final arithmetic assembly first. These explain why three specific estimates suffice and where the critical exponent arises. Read the parameter and common arithmetic-model chapters next. The upper-bound, spanning and lower-bound parts can then largely be studied separately. Each includes both the main chain of deductions and complementary chapters which derive supporting lemmas and specialize classical theorems. Earlier statements are sometimes named using their formal declaration to preserve concordance even when individual components are subsequently proved. Consult the final dependency audit and the internal cross-references to distinguish the complete parts of a proof from those whose exact specialized hypotheses still need to be established.

Part	Main mathematical question
I. Analytic reduction and models	How can a one-sided smoothed prime sum exclude a zero, and what are the fields, weights, levels and integral spaces that will produce it?
II. Upper bounds	Why do derivatives of CM theta functions contain the valuations needed to count prime powers, and how can lifts of integral forms be made short without a prohibitively large index?
III. Spanning	How can the short forms be made to span almost all of a large automorphic space, with an explicit bound on the number of translates?
IV. Lower bounds	Why must an almost-basis have large total volume, and how do exchange identities, special values, and congruences control the integral costs?
V. Assembly and audit	How do these estimates cancel, yield a square-root saving, and imply GRH? Which mathematical dependencies remain to be independently justified?

A conceptual point. Products of positive integers that encode primes can give large archimedean values; factorials are the simplest example. The proof at issue is much more elaborate, but one can recognize a familiar product-formula pattern. A value whose finite valuations are large cannot simultaneously be too small at every infinite place, unless it vanishes. Here the relevant object is not a single nonzero value but a family of automorphic vectors, its determinant, and an auxiliary family of twists measuring different finite places. The main term of the upper estimate measures the cost of derivatives after subtracting guaranteed divisibilities. The lower estimate recovers exactly that cost with the opposite sign via volume ratios. The canonical cancellation between them, not an absolute estimate of a high-degree automorphic L-function, gives the saving. This is the mathematical thread to keep in mind across the many local constructions.

Conventions. The prime d defines $K = \mathbb{Q}(\sqrt{-d})$, and $\mathcal{D}$ denotes a chosen admissible datum. The integer m is the number of theta lifts, r the initial row rank, $k = r + 1 - m$ the initial column rank, and $t = 5r$ the rank increment on the row chain. The symbol $n = \dim W$ is the dimension of the row form space, not a Hermitian rank. We reserve d_{ind} for a logarithmic index

cost. The scale b is real and large, $B = fa(\log b)^3$ is the basic error allowance, $L = B \log b$, and $H = mk\mathcal{U}$ is the common main term. The normalized prime error R_M differs from the arithmetic remainder $\mathcal{R}$. A glossary at the end of the volume collects these and other important symbols, with references to the chapters where they enter. We use the compact edition's numerical constants unless otherwise specified. For source references, a module path beginning with `New/` or `Lib/` is relative to `grh-compact/src/` of the supplied repository.

Contents

Preface: the purpose of this monograph **i**

I The analytic reduction and arithmetic models **1**

1 Introduction and mathematical status **2**
1.1 The architecture of the proof 2
1.2 Primary material, conventions, and the remaining challenge 4

2 The analytic problem: a one-sided bound for weighted prime sums **5**
2.1 The particular smoothing used in the proof 5
2.2 Why the prime sum contains information about zeros 7
2.3 An elementary form of Landau's one-sided principle 8
2.4 Quadratic twists avoiding a prescribed half-strip 9
2.5 The one-sided zero-free criterion 12
2.6 A good second scale in a long window 13
2.7 From the target bound to GRH 14

3 Classical analytic foundations without a prime number theorem **16**
3.1 Finite characters, Gauss sums, and the Fourier bound 16
3.2 The quadratic symbol and the choice of the family 18
3.3 Holomorphy and nonvanishing on the line at one 19
3.4 Reciprocal primes in a fixed residue class 21
3.5 The theta transform and the full functional equations 22
3.6 Analytic details of the bad-twist estimate 24
3.7 A uniform disk at one for powers of a fixed modulus 25
3.8 Prime existence in long logarithmic windows 27
3.9 Final conditional theorem and quantifier audit 29

4 The admissible data and coefficient spaces **32**
4.1 The CM extension and the different 32
4.2 The scales, and the reason for so many parameters 33
4.3 Admissibility and construction of the data 34
4.4 Derived weights, level and integral structures 36
4.5 Two kinds of automorphic forms 36

4.6 Precise weights and levels of the source and target . . . 37
4.7 Integral gauges and the precise uniform index . . . 39
4.8 The theta chains and their relation . . . 40
4.9 Norms, squared covolume and index . . . 40
4.10 The common leading term and the prime remainder . . . 41

5 Adelic coefficient models and integral linear algebra 43
5.1 Hermitian spaces, places and adelic lattices . . . 43
5.2 Why the extra level removes stabilizer denominators . . . 45
5.3 Normalization of class sums and change of level . . . 46
5.4 A useful elementary form of weak approximation . . . 47
5.5 Algebraic modular forms and the finite-dimensional model . . . 48
5.6 The integral Weyl model . . . 48
5.7 Fischer pairing and the highest-weight hook constant . . . 50
5.8 Examples of the coefficient normalizations . . . 51
5.9 The class-sum metric and the real matrix-coefficient model . . . 51
5.10 Transported lattices and gauges . . . 52
5.11 Finite-place Hecke operators on explicit lattices . . . 54
5.12 Class correspondences and adjoints . . . 55
5.13 Where this elementary model stops . . . 56

6 The splitting character, integral unitary models and the definite mass 57
6.1 Notation and the local norm theorem . . . 57
6.2 A character extension with exactly the required sign . . . 58
6.3 The self-dual lattice and half of the different . . . 61
6.4 Reduction of finite ideals and their adelic generators . . . 62
6.5 Completions of the coefficient primes and the choice of a side . . . 63
6.6 Unimodular lattices over split and unramified places . . . 64
6.7 Existence and equivalence of the framed global spaces . . . 65
6.8 Counting the local finite groups and their congruence quotients . . . 66
6.9 An exact norm-one class-sum identity . . . 67
6.10 The classical mass formula in the present normalization . . . 68
6.11 Useful consequences and the separation of normalizations . . . 69

7 Heisenberg operators and the finite vacuum multiplier 71
7.1 The saturating lattice at a finite place . . . 71
7.2 The multiplier and its transformation laws . . . 72
7.3 The finite vacuum and the Heisenberg cocycle . . . 73
7.4 The lattice model and the absolute vacuum coefficient . . . 74
7.5 The Gauss-sum phase and the finite splitting . . . 75
7.6 The real harmonic model and its basic normalization . . . 77
7.7 The rational theta functional in a polarized model . . . 78

II Upper bounds: CM jets, theta chains and Hecke repair **80**

8 The CM differential, theta normalization and integral frame **81**
8.1 An elliptic curve normalized at a point of order two . . . 81
8.2 Theta sections and their integral size . . . 83
8.3 Isogenies, their multipliers, and Tate's inequality . . . 84
8.4 The integral formal parameter and the Gauss principle . . . 85
8.5 Logarithms, norm-pair limits, and the formal height . . . 87
8.6 A unit-adjusted Chowla–Selberg formula . . . 89
8.7 The finite idele and the period scale . . . 91
8.8 Constructing arithmetic tight frames of prescribed length . . . 92
8.9 Frobenius on the frame matrix . . . 95
8.10 Theta-group orthogonality and the finite realization . . . 96
8.11 From an arithmetic frame to a theta section . . . 98
8.12 Integral vacuum and the remaining reciprocity law . . . 99
8.13 What the arithmetic normalization proves . . . 101

9 Local CM theta analysis and Schur jets **102**
9.1 The arithmetic local scales . . . 102
9.2 The CM period and the differential coordinate . . . 103
9.3 Integral and logarithmic theta identities . . . 104
9.4 The regularized CM parameter at nonsplit primes . . . 105
9.5 The exceptional exponential factor . . . 106
9.6 Tight frames and theta vectors . . . 107
9.7 Frobenius pairings in tight-frame products . . . 108
9.8 Schur matrices and determinant losses . . . 110
9.9 Application to the prescribed weights . . . 113
9.10 From theta products to the standard integral kernel data . . . 114

10 The definite chain, collisions and Hecke repair **117**
10.1 Coefficient lattices, class spaces and index . . . 117
10.2 The groups and the chain of theta kernels . . . 119
10.3 The geometric seesaw and its denominator control . . . 121
10.4 Ramification and short generators on the definite group . . . 124
10.5 The collision loss and the marked set . . . 125
10.6 Moving the marked Hecke operators through the chain . . . 127
10.7 The Frobenius-twisted elementary-divisor theorem . . . 130
10.8 Application to the actual column lattice . . . 136
10.9 The prime N: moving denominators to a Fourier operator . . . 138
10.10 The simultaneous local divisibility statement . . . 140
10.11 Archimedean growth and the level normalization . . . 141
10.12 Projection, product formula scaling and Theorem A . . . 142

11 The local Hecke estimate in full detail, and the integral seesaw 144
11.1 The cyclic coefficient representation and the two filtrations 144
11.2 The root estimate, including its coefficient extraction 145
11.3 The dimension bound from flags and convex volume 147
11.4 Deep coefficients: a weighted Nullstellensatz and its application 151
11.5 Integral diagonal pairings and class means 153
11.6 Frobenius products of matrices and the approximate projector 155
11.7 Primitive half-depth matrices 158
11.8 The Smith pairing lemma and the local Hecke theorem 160
11.9 The finite Fourier operator at the level prime 162
11.10Geometry and integrality in passing to a seesaw 165

12 Normalized CM reciprocity and the geometry of the iterated seesaw 168
12.1 A marked algebraic curve and its sections 168
12.2 Simultaneous torsion translations and finite-adelic reciprocity 171
12.3 Removing one string: the Hermitian linear algebra 173
12.4 The surviving rational orbit 174
12.5 The one-step period transfer 177
12.6 The glued lattice and integral Witt extension 180
12.7 Adapted bases and equality of lattice lines 182
12.8 Local supports of zero terms 183
12.9 A one-step bound for the lattice lines 184
12.10Finite cuts with prime-supported multiplicity 187
12.11Two edges at once: removal of the extra plane 189
12.12The split neighbour lattice in the paired evaluation 191
12.13Weighted integrality of evaluations 192
12.14The admissible combinations and the iterated identity 194
12.15The remaining representation-theoretic comparison 196

III Spanning: Bessel and Poincare series 197

13 Formal degrees, elliptic trace comparison and the row dimension 198
13.1 Parameters and the compact-dual Haar measure 198
13.2 Evaluation of the constant in the formal degree 199
13.3 The exact elliptic orbital character identity 202
13.4 Absolute convergence and a direct ellipticity argument 203
13.5 Changing the signatures of the rational elliptic classes 205
13.6 The trace of the elliptic finite test functions 206
13.7 Primitive labels and the finite cubic correction 207
13.8 Equality of all good Hecke traces and determinants 209
13.9 The precise row-dimension inequality 210

14 Quantitative spanning by projected Bessel–Poincaré series 213
14.1 Statement and finite-dimensional framework . . . 213
14.2 Finite level, central projection, and the transfer bound . . . 214
14.3 Hermitian geometry of the Bessel subgroup . . . 216
14.4 The archimedean Bessel system . . . 217
14.5 Construction and exact unfolding of the Poincaré map . . . 220
14.6 Energy, pairs of congruent terms, and removal of lower cells . . . 223
14.7 The sharp coherent main term . . . 226
14.8 A finite-field positivity argument for singular minors . . . 230
14.9 First moment and bad finite-type loss . . . 232
14.10From the moment bounds to the covering theorem . . . 233

15 The analytic and arithmetic mechanism of the Bessel sieve 236
15.1 Norm balls in the local groups . . . 236
15.2 A rank-one convolution estimate and its differentiated form . . . 238
15.3 Threshold projectors and a lowest-type coefficient . . . 240
15.4 Two-factor Laplace estimate . . . 242
15.5 Interpolation, isotropic frames, and the Bessel peak . . . 244
15.6 A peak-to-cut inequality with all mass factors . . . 247
15.7 From an L^2 reproducing estimate to the bandlimit tower . . . 249
15.8 The cell L^2 estimate and its limited differentiation cost . . . 251
15.9 A finite-subset version of adelic Poisson summation . . . 252
15.10An explicit induction for the sieve system . . . 253
15.11Frequency summation and the product formula . . . 255
15.12Finite packets, their normalization, and the global comparison . . . 256
15.13Degenerate terms and the accumulation of energy errors . . . 258

16 Quantitative Bessel construction and the arithmetic of the filling datum 260
16.1 The compact integral and the change of splitting . . . 260
16.2 Logarithmic principal-minor estimates by resolvents . . . 262
16.3 Gaussian shape of the coefficient . . . 265
16.4 Determinant decay and a small-ball bound . . . 266
16.5 Removal of the common compact subgroup . . . 267
16.6 A phase-cancelling realization of the Bessel group . . . 269
16.7 Marked integral lattice and the exact finite level indices . . . 272
16.8 Haar coordinates and the arithmetic unipotent covolume . . . 274
16.9 Cutoffs, approximation and the quotient character . . . 277
16.10Orbital cutoffs and Bessel transport between classes . . . 280
16.11Agreeing characters and propagation of a shared flag vector . . . 282
16.12The link as a Hermitian Lagrangian and its congruence defect . . . 284
16.13Saturated lattice heights and their zeta sum . . . 286
16.14The global frequency estimate and explicit sieve budget . . . 288
16.15The finite zero form and central projections . . . 290

16.16The projection covariance and the complement-coset error 292
16.17Assembly of the sharp energy bound 294
16.18What remains of the filling argument 295

IV Lower bounds: periods, congruences and crystalline defects 296

17 The determinant lower bound 297
17.1 Conventions for norms of conjugate coefficient vectors 298
17.2 Heights with varying finite-place entries 299
17.3 The determinant lemma with index 300
17.4 Special-value data and fit twists 301
17.5 The correspondence of complementary rank and the finite square 303
17.6 The two-grid exchange and the crude rows 306
17.7 The triangle of weighted lifts 309
17.8 Comparison forms for one lift 312
17.9 Compressed Gram determinants 315
17.10The height saving in a rank-one special value 317
17.11Combining the special values and closing the lower bound 320

18 Period identities, finite matching, and the geometry of the lower bound 322
18.1 Complementary frames and matched data 322
18.2 Fourier transformations of the matching identity 325
18.3 The graph term in doubling 328
18.4 Sharp support, level factors, and the common scalar 330
18.5 A rank-one Euler factor test 332
18.6 Hyperbolic-plane removal and the signed product 333
18.7 What makes an exchange an identity on all partners? 334
18.8 Level orders in the paths 336
18.9 Why a false lift with one common root contributes no period 337
18.10Two finite ingredients of integral comparison 339
18.11Limits of the reconstruction 341

19 The exchange construction, the two grids, and comparison lattices 342
19.1 Local integrals of complementary frames 342
19.2 Generating all partner data by stable spaces 345
19.3 Construction of the real data on an enlarged Hermitian space 347
19.4 A plane removal on determinant vectors 350
19.5 How the equal-rank normalization is obtained 351
19.6 From a convergent exchange to an operator square 354
19.7 Terminal interlacing and the signatures of the two grids 355
19.8 The finite square move with no accumulated denominator 357
19.9 Moving the descent along the path 358
19.10The divided lattice that glues the two comparison groups 360

19.11 The loss of the Cartan product . . . 362
19.12 Descent of a glued point form . . . 363
19.13 Rank-one reflections and the small-root descent argument . . . 364
19.14 The comparison map and the precise scalar it sees . . . 367
19.15 The Archimedean budget of the comparison operator . . . 369
19.16 Two further finite calculations for the rank-one values . . . 370
19.17 Summary of the strengthened reconstruction in C . . . 373

20 Congruences, auxiliary primes, and the fit family 375
20.1 The integral projection problem . . . 375
20.2 Labels and Galois representations of lifted eigensystems . . . 376
20.3 Selmer modules over finite coefficient rings . . . 379
20.4 New level characters and tied cocycle conditions . . . 380
20.5 Shrinking several tied Selmer modules . . . 382
20.6 Finite presentations and interpolation of lengths . . . 384
20.7 Degree-free selection of roots of unity . . . 385
20.8 The global twisting theorem and value families . . . 387
20.9 The Tate characters and triangular Taylor–Wiles places . . . 389
20.10 The first-order Ribet construction . . . 391
20.11 Hecke operators realizing the first-order relations . . . 394
20.12 The denominator bound from the local length budgets . . . 396
20.13 Global numerical summation and proof of the fit theorem . . . 398

21 The foundations of the congruence argument 400
21.1 From weak Hecke compatibility to a trace identity . . . 400
21.2 Tame labels and full matrix span . . . 403
21.3 A duality calculation at the mirror . . . 405
21.4 Anticyclotomic primes from Kummer extensions . . . 407
21.5 Odd laws and the triangular removal of even classes . . . 408
21.6 The algebra of adjoining relations . . . 412
21.7 Passing from scalar rows to matrix blocks . . . 413
21.8 Two applications of Cramer's rule over a family . . . 414
21.9 A transfer calculation at one auxiliary level . . . 415
21.10 Separation of the arithmetic and algebraic budgets . . . 418

22 Arithmetic of the congruence argument: cyclotomic classes, descent and simultaneous twists 419
22.1 The integer in the odd-class budget . . . 419
22.2 The Jacobi family and its distribution relation . . . 421
22.3 The derivative cocycles from the Jacobi family . . . 423
22.4 The odd component, local vanishing and reciprocity . . . 425
22.5 The Kolyvagin–Rubin detection argument . . . 427
22.6 Uniform detection for the small primes . . . 429

22.7 Returning from Dirichlet components to the CM field 430
22.8 Uniform cardinality of Selmer pencils 432
22.9 Dropping ties along a line of approximate characters 434
22.10 The two specializations in the common-twist theorem 435
22.11 Weak attachment: the trace and descent reductions 438
22.12 Hecke-polynomial continuations 442
22.13 Rank-one normalization of the projection bound 444
22.14 Dependency status of the congruence argument 445

23 Local crystalline defects and the integral cup-product estimate 446
23.1 Lattices, crystalline cocycles and imperfectly dual reductions 446
23.2 Why this defect bounds the local discrepancy of congruent cocycles 449
23.3 The Kummer chart, letter decomposition and height of a frame 450
23.4 Exponent lattices and the cost of increasing depth 452
23.5 Cokernel estimates: divisibility by p and by the Eisenstein element 454
23.6 Cup-product vanishing from the three sides 455
23.7 The window estimates: the three ways to obtain safe points 458
23.8 The tame type bound, uniform in the ramification degree 461
23.9 Case analysis for the primes and the explicit local loss 462
23.10 Insertion of the global weights and the sum over all blocks 463

24 Integral period frames and the algebra of the local loss 466
24.1 The rank calculation and duality over a finite coefficient ring 466
24.2 The coordinates in a Kummer chart 468
24.3 A coefficient retraction and why it is used 470
24.4 From crystalline periods to an integral frame 471
24.5 The radius estimate and its fourth power 474
24.6 Exponent lattices, Frobenius letters and the small part 476
24.7 Exact cokernel multipliers via finite differences 478
24.8 A rank-sensitive bound for raising filtered conditions 480

25 Integral crystalline comparison, depth operators and the Frobenius walk 484
25.1 The period rings and an integral contraction 484
25.2 The period-basis argument 486
25.3 Reduction of the determinant and the integral frame 489
25.4 Scalar depth steps at each Frobenius divisor 491
25.5 Depth zero and finite averaging 493
25.6 From scalar depths to crystalline lattice depths 494
25.7 From a single operator to the cost of a window 497
25.8 The monic Taylor coefficient at a long gap 498
25.9 The triple pairing and the safety rule 501
25.10 From safe pairs to an actual cup coboundary 503
25.11 The walk along the Frobenius permutation 505

25.12Plain cuts and long-gap cuts . . . 507
25.13Quantitative killers at sparse cut positions . . . 509
25.14Two choices of cuts . . . 511
25.15A multiplicity-sensitive step in the few-low-jumps case . . . 513
25.16Scope of the integral lifting reconstruction . . . 516

V The arithmetic conclusion and proof audit 517

26 A numerical appendix: asymptotic budgets and the passage to the limit 518
26.1 Uniform scale asymptotics . . . 518
26.2 Rounding of large coefficients and a Lipschitz estimate . . . 519
26.3 The exact level cancellation revisited . . . 519
26.4 An elementary bound on the logarithm of the class number . . . 520
26.5 Summation of the fit-family costs . . . 521
26.6 The final numerical inequalities . . . 522

27 The arithmetic inequality and conclusion 523
27.1 Collecting the arithmetic inputs . . . 523
27.2 From A and B to an almost-basis of short forms . . . 524
27.3 Comparison with the lower bound . . . 524
27.4 Extracting the analytic remainder . . . 525
27.5 Choosing the second cutoff and the power saving . . . 526
27.6 GRH . . . 526

28 Glossary and concordance 528
28.1 Reading the main theorem in the Lean source . . . 529

29 Dependency ledger and the status of the reconstruction 531
29.1 Three levels of mathematical justification . . . 531
29.2 Where to find the added derivations . . . 532
29.3 The exact logical role of the remaining work . . . 533
29.4 What the compact and readable editions contribute . . . 535

Part I

The analytic reduction and arithmetic models

Chapter 1

Introduction and mathematical status

The compact edition of the Alpöge–Claude formal development [1] is a remarkably large argument for the following assertion.

Theorem 1.1 (Dirichlet GRH; asserted in the formal source)**.** *For every integer $q \geq 1$, every Dirichlet character χ modulo q, and every $s \in \mathbb{C}$ with $\Re(s) > 1/2$,*

$$L(s, \chi) \neq 0,$$

where the pole of a principal character at 1 is excluded. In particular every zero in $0 < \Re(s) < 1$ lies on the line $\Re(s) = 1/2$, and the nontrivial zeros of ζ lie on this line.

The qualifier Dirichlet is important: the paper concerns the L-functions of degree one over $\mathbb{Q}$, not the generalized hypothesis for all automorphic L-functions. The proof uses higher rank unitary groups over auxiliary CM fields, but it uses them to construct and compare algebraic objects, not to assert GRH for their standard L-functions.

This monograph is the third stage of the series. It retains the interfaces of the shorter specialist treatment, expands their technical proofs, and supplies additional mathematical background, examples, and context for readers trained in neighboring areas. Definitions are fixed in their operative form, and references connect each mathematical package to the relevant part of the compact Lean development. It is a *working reconstruction of the formal proof*, not a claim that every arithmetic argument has yet been independently translated into a conventional proof. Every assertion called a "formal-source input" is a mathematical statement extracted from, or a stated interface to, the supplied Lean development that still needs an independent written proof or a verified literature replacement. A statement followed by a proof records a mathematical derivation of that statement from the explicitly stated inputs; it does not thereby establish those inputs. In particular the manuscript must not be cited as an independently complete proof of Theorem 1.1 until the remaining formal-source inputs have been discharged. The main outstanding packages, including the exact geometric and integral assertions needed for subsequent deductions, are listed at the end. This convention allows a specialist to test an argument at the level of its precise hypotheses, instead of having to infer whether an omitted detail is formal bookkeeping or a substantial arithmetic claim.

1.1 The architecture of the proof

The principal idea is to use the same automorphic object for two opposed quantities. At the archimedean places it carries a metric; at the finite places it carries an integral coefficient

structure. These are linked by a product formula. To encode primes, consider jets of a theta kernel for a CM elliptic curve in a specified algebraic normalization. Their divisibility at a nonsplit rational prime depends on how the local formal-group parameter transforms under Frobenius. The valuation bounds contain sums of positive parts indexed by powers of the prime. A choice of the infinite-place weights turns their global sum into the weighted prime remainder.

One cannot simply apply the product formula to an arbitrarily chosen jet. The proof instead uses a family of forms obtained through an iterated theta correspondence. Suitable source forms are transferred by a seesaw identity to the lifts of Poincaré series on an indefinite unitary group. The arithmetic choice of fields, levels and Harish-Chandra parameters serves three tasks at once: the finite-place integrality estimate, a nearly sharp second moment for Poincaré series, and a matching lower bound on the determinant of the lift. The latter requires a family of twists, one at each finite embedding, rather than a single twist valid at every place.

Here are the three central estimates, in the notation that will be fixed in Chapter 4. Let W be a row space of dimension n, C the source space of dimension c, and $C^{\mathrm{ram}} \subseteq C$ its ramified part. Let Θ_θ be a composite of m lifts, and $P_s : C^{\mathrm{ram}} \to W$ the projected translates of a Poincaré map, $s \in \mathcal{G}$. The large positive numbers $B = fa(\log b)^3$ and $L = B \log b$ are error budgets, while $H = mk\mathcal{U}$ denotes a signed common main term. The weighted prime remainder is $\mathcal{R}$.

A. There exists a positive rational scaling of the P_s and a spanning set of integral elements of C^{ram} for which all trivial-twist lifts have squared norm at most $\exp(H - m\mathcal{R} + L)$ at every archimedean conjugate. A family from the images of J translates has logarithmic index at most $|J|cB$ simultaneously in the admissible integral structures.

B. Both c and n are positive. There are translates indexed by a finite set J with $|J|c \le 5n \log(1 + b^2)$ whose images span at least $(1 - 4/b^2)n$ dimensions.

C. Every independent family of at least $(1-4/b^2)n$ rows whose trivial-twist lifts have squared covolume at most e^V at all conjugates and whose logarithmic index is at most d_{ind} satisfies $n(H - 3L) \le V + 2d_{\mathrm{ind}}$.

Assertions A, B, C are the interfaces `Data.short_forms`, `Data.enough_forms`, and `Data.integral_forms_long` of the compact source. Their proof decompositions occupy Parts II–IV of the monograph. Notice that C is uniform in the choice of row family, unlike A, and that B is quantitative; neither of these features can be dropped in the endgame.

To see the algebraic mechanism, choose $q \ge (1 - 4/b^2)n$ independent short rows using A and B. Hadamard's inequality gives $V \le q(H - m\mathcal{R} + L)$, while the index is at most $5n \log(1 + b^2)B$. Combining with C, and bounding the missing-dimension term by a crude estimate, gives

$$m\mathcal{R} \le 4L + 4x^{99} + 10 \log(1 + b^2)B \le fa(\log b)^5.$$

The two instances of H cancel exactly. The number m of lifts is of size $\sqrt{x}$; this is the source of the square-root factor in the final saving. The normalized prime remainder $R_{N_0}(x)$ is related to $\mathcal{R}$ by the exact identity

$$\frac{\mathcal{R}}{a} = |\mathcal{S}|(R_{N_0}(x) - R_{N_0}(x_1)) + |\mathcal{T}|R_1(x_1),$$

where $\mathcal{T}, \mathcal{S}$ are two subsets of infinite places. The compact analytic argument chooses x_1 in a long logarithmic window to give a one-sided upper bound for $R_{N_0}(x_1) - |\mathcal{T}|R_1(x_1)/|\mathcal{S}|$. With $|\mathcal{S}| \ge |\mathcal{T}|/\log b$, $f = \varphi(N_0)|\mathcal{T}|$, and $\log b = x^\varepsilon$, one concludes $R_{N_0}(x) \le x^{-1/2+8\varepsilon}$ eventually. A Landau argument then excludes zeros to the right of $\frac{1}{2} + 8\varepsilon$. Quadratic twists are varied to prevent cancellation of the singularities of logarithmic derivatives.

1.2 Primary material, conventions, and the remaining challenge

The supplied snapshot `grh-lean-main` contains the two variants `grh/` (readable) and `grh-compact/` (compact), of approximately 300,000 and 76,000 Lean lines respectively. Their high-level interfaces agree but some proofs differ. We take the compact variant as primary and use the readable one when clarifying its definitions or constituent steps. Unless qualified otherwise, paths in the text are relative to `grh-compact/src/`. The root README, attribution and guide identify the underlying proof as due to Levent Alpöge with Claude and the shortening as a collaboration using Claude. The formal source reports kernel checks of all exported results and dependence only on `propext`, `Classical.choice`, and `Quot.sound`; this manuscript makes no claim of a new independent machine replay. The statement of the result imports only Mathlib's definitions, not the auxiliary arithmetic theory of this argument. The exact goal is in `Statements.lean`, and the assembly is `New.grh` in `New/Main.lean`.

The mathematical difficulty is not measured by Lean line count. In particular the division into `Lib` and `New` is not a proof that all of the former can be replaced by standard citations. Uniformity in growing ranks, large weights and varying auxiliary fields, integral versions of the correspondence, normalization of class sums and precise determinant comparisons must all survive such replacement. We therefore take care to mark specialized inputs instead of silently replacing them by named theorems from the literature. Classical results used in the analytic and parameter-construction sections are identified explicitly and do not presuppose GRH.

Notation. The letter d always denotes the chosen auxiliary rational prime; d_{ind} denotes an index bound. The size b is a scale, not a weight coordinate. We use Φ for the set of embeddings of E above the chosen embedding of K, $\mathcal{T}$ for the tops, and $\mathcal{S}$ for the thick places. We write $A(y,\sigma)$ for the weighted sum in `Data.primeSum`, $A_p(y,\sigma)$ for its normalized p-part, and $J(j)$ for `Data.jetSize`. In some local arguments letters may denote temporary variables, but the main budgets $B, L, H, \mathcal{R}$ retain their meanings. The notation W_j for the auxiliary Hermitian spaces is different from W, the vector space of row forms. The group $\mathcal{G}$ is also denoted $\mathfrak{G}$ in the finite moment estimates; $C_{\mathrm{ram}} = C^{\mathrm{ram}}$.

Chapter 2

The analytic problem: a one-sided bound for weighted prime sums

The next chapter, Chapter 3, supplies the classical analytic prerequisites and the proof of the special prime-window theorem. The final conditional theorem there also verifies the quantifiers in the passage from the three arithmetic inequalities to the assertion for all Dirichlet characters.

Before constructing the automorphic objects, we isolate the precise problem that they solve. The analytic reduction is useful in its own right: it separates the difficult arithmetic estimate from the comparatively short argument about zeros. In particular, the arithmetic construction need not produce a two-sided error estimate for the ordinary prime counting function.

Throughout this section $M \geq 1$ is a modulus, d is a prime with $d \equiv 7 \pmod 8$ and $(d, M) = 1$, and

$$K = \mathbb{Q}(\sqrt{-d}), \qquad \chi_K(\nu) = \Big(\frac{\nu}{d}\Big).$$

The discriminant of K is $-d$; quadratic reciprocity identifies the displayed character with its quadratic Dirichlet character. Write $\varphi(M)$ for Euler's totient. The dependence on d in the notation R_M below is suppressed, but it is essential that the auxiliary prime d can eventually be varied.

2.1 The particular smoothing used in the proof

Let $\Lambda(p^j) = \log p$ for prime powers, and let $\Lambda(\nu) = 0$ otherwise. Put

$$\lambda_{d,M}(\nu) = \mathbf{1}_{\nu \equiv 1 \bmod M} \Lambda(\nu)(1 - \chi_K(\nu)). \tag{2.1}$$

Thus $\lambda_{d,M}$ is non-negative. A prime split in K contributes zero, an inert prime contributes twice its ordinary von Mangoldt weight, and the ramified prime has factor one. The same rule applies to prime powers; for instance an even power of an inert prime has factor zero. Congruence $\nu \equiv 1 \pmod M$ is imposed on the integer ν, not on its underlying prime. In particular, the weighted sum is not literally a sum over primes alone.

For any Dirichlet character ξ, with its modulus understood, denote by

$$\widetilde{L}(s, \xi) = \begin{cases} L(s, \xi), & \xi \text{ non-principal}, \\ (s-1)L(s, \xi), & \xi \text{ principal} \end{cases}$$

the entire function obtained by removing the possible pole at 1 (the second line is interpreted by continuation). The character $\xi\chi_K$ is taken modulo Md; it is non-principal because d is a

new odd prime. Define the meromorphic function

$$G_{d,M}(s) = \frac{1}{\varphi(M)} \left\{ \sum_{\xi \bmod M} \left(\frac{\widetilde{L}'}{\widetilde{L}}(s, \xi\chi_K) - \frac{\widetilde{L}'}{\widetilde{L}}(s, \xi) \right) - 1 \right\}. \tag{2.2}$$

Here and throughout, a sum over characters includes the principal and imprimitive characters. The constant subtracted in the prime error is

$$c_{d,M} = \operatorname{Re} G_{d,M}(1).$$

The classical non-vanishing on $\operatorname{Re} s \geq 1$ of the pole-removed Dirichlet L-functions ensures that this constant exists. In fact $G_{d,M}(1)$ is real, by conjugation of characters.

The error in the compact Lean development, called `primeError`, is exactly

$$R_M(x) = \sum_{1 \leq \nu \leq x} \lambda_{d,M}(\nu) \left(\frac{1}{\nu} - \frac{1}{x} \right) - \frac{\log x}{\varphi(M)} - c_{d,M} \qquad (x > 0). \tag{2.3}$$

It is important to include the constant as well as the logarithm in the normalization. Omitting it would prevent the Laplace transform used below from being regular at $s = 1$. In modulus one, the constant has a more specific expression. If γ is Euler's constant, write

$$\kappa_K = \gamma - \operatorname{Re} \frac{L'}{L}(1, \chi_K).$$

Since $(s-1)\zeta(s) = 1 + \gamma(s-1) + O((s-1)^2)$,

$$c_{d,1} = -1 - \kappa_K. \tag{2.4}$$

The number κ_K, called `chowlaSelbergConst` in the code, recurs in the arithmetic estimates: the appearance of the same constant in the norm of a CM theta period is part of the eventual cancellation.

The choice of smoothing in (2.3) is particularly convenient for the arithmetic proof. For $y > 0$ set $(y)_+ = \max(y, 0)$, and extend this notation to all real numbers. For an integer a and $0 < z \leq a$ one has the exact identity

$$\frac{1}{a} \sum_{\nu=1}^{a} \lambda_{d,M}(\nu) \left(\frac{a}{\nu} - \frac{a}{z} \right)_+ = R_M(z) + \frac{\log z}{\varphi(M)} + c_{d,M}. \tag{2.5}$$

Indeed the positive part vanishes for $\nu \geq z$ and gives $a(1/\nu - 1/z)$ for $\nu < z$. No asymptotic replacement is involved. The importance of this identity is that the arithmetic construction bounds valuations of jets expressed using precisely such positive parts, with integer order parameter a and scale z. The apparently unusual smoothed prime sum is therefore dictated by a divisibility calculation, not just by analytic convenience.

Our ultimate target is, for every $0 < \varepsilon < 1/100$ and every suitable d,

$$R_{N_0}(x) \leq x^{-1/2+8\varepsilon} \qquad (x \text{ sufficiently large}), \tag{2.6}$$

where $18480 \mid N_0$ and $d \equiv N_0/2 - 1 \pmod{N_0}$. Only the upper bound occurs in (2.6). The threshold for x may depend on the fixed parameters N_0, d and ε. We shall explain why such a bound is already strong enough.

Source. The definitions, the constant in modulus one and (2.5) are in `New/Endgame/PrimeError.lean`, declarations `lambdaClass`, `logDerivMean`, `primeError`, `re_logDerivMean_one`, `primeError_eq_sum`. The function $\widetilde{L}$ is `entireL` in `Lib/General/NumberTheory/QuadField.lean`.

2.2 Why the prime sum contains information about zeros

Set

$$E_M(v) = R_M(e^v), \qquad \Phi_M(s) = \int_0^\infty E_M(v) e^{(1-s)v}\, dv.$$

The variable s is the same complex variable as in the L-functions. The initial integral converges for $\operatorname{Re} s > 2$: for $v \geq 0$, the elementary bound $|\lambda_{d,M}(\nu)| \leq 2\nu$ gives $|E_M(v)| \leq Ce^v$, with C depending on M, d.

Proposition 2.1 (Transform and residues)**.** *The integral Φ_M continues meromorphically, with the expression*

$$\Phi_M(s) = \frac{G_{d,M}(s)/s - c_{d,M}}{s-1}. \tag{2.7}$$

It is holomorphic in a neighborhood of the closed half-plane $\operatorname{Re} s \geq 1$. For $0 < \operatorname{Re}\rho < 1$ its residue at ρ is

$$\operatorname{res}_{s=\rho} \Phi_M(s) = \frac{\sum_{\xi \bmod M} \left(\operatorname{ord}_\rho \widetilde{L}(\cdot,\xi) - \operatorname{ord}_\rho \widetilde{L}(\cdot,\xi\chi_K)\right)}{\varphi(M)\rho(1-\rho)}. \tag{2.8}$$

In $\operatorname{Re} s > 0$ any poles are at most simple.

Proof. By orthogonality of characters and the logarithmic derivative of the Euler product, the Dirichlet series of (2.1), absolutely convergent for $\operatorname{Re} s > 1$, is

$$\begin{aligned} D_{d,M}(s) &= \sum_{\nu \geq 1} \frac{\lambda_{d,M}(\nu)}{\nu^s} \\ &= \frac{1}{\varphi(M)} \sum_{\xi \bmod M} \left(\frac{L'}{L}(s,\xi\chi_K) - \frac{L'}{L}(s,\xi)\right) \\ &= G_{d,M}(s) + \frac{s}{\varphi(M)(s-1)}. \end{aligned}$$

In the last equality the unique principal character modulo M contributes $1/(s-1)$ besides the pole-removed logarithmic derivative, while no twisted character is principal. The additional 1 in the definition of $G_{d,M}$ has now combined with that pole term.

Absolute convergence justifies integrating the sum in (2.3) term by term for $\operatorname{Re} s > 2$. For an integer $\nu \geq 1$,

$$\int_{\log\nu}^\infty (\nu^{-1} - e^{-v}) e^{(1-s)v}\, dv = \nu^{-s}\left(\frac{1}{s-1} - \frac{1}{s}\right) = \frac{\nu^{-s}}{s(s-1)}.$$

The transform of v is $1/(s-1)^2$, and that of the constant function one is $1/(s-1)$. Consequently

$$\begin{aligned} \Phi_M(s) &= \frac{D_{d,M}(s)}{s(s-1)} - \frac{1}{\varphi(M)(s-1)^2} - \frac{c_{d,M}}{s-1} \\ &= \frac{G_{d,M}(s)/s - c_{d,M}}{s-1}. \end{aligned}$$

This gives the desired continuation. The apparent singularity at 1 is removable because $c_{d,M} = G_{d,M}(1)$. The classical non-vanishing in $\operatorname{Re} s \geq 1$ rules out all other poles there.

If an entire function F has a zero of order j at ρ, then we can write $F(s) = (s-\rho)^j f(s)$ locally, with $f(\rho) \neq 0$. Hence

$$\frac{F'}{F}(s) = \frac{j}{s-\rho} + \frac{f'}{f}(s).$$

This proves (2.8), including its sign, and shows that the poles are simple. □

The numerator in (2.8) is a *difference* of zero multiplicities. Thus a bound proving that Φ_M has no poles in a certain half-plane would not, by itself, show that every $L(s,\xi)$ is zero-free there: a pole might cancel between the two sums. The remedy is to choose d so that the finitely many twisted L-functions have no zero at a hypothetical offending zero, and no real zero to the right of the relevant line.

In the code the properties of E_M and its transform are packaged by the predicate `IsPolar`. The assertion `isPolar_primeError` in `New/Endgame/PrimeError.lean` establishes this predicate with the residue function (2.8) (zero at 1). This abstraction enables the two subsequent analytic arguments to be proved without referring to primes or automorphic forms.

2.3 An elementary form of Landau's one-sided principle

Here is the precise analytic feature that makes an upper bound sufficient. We give the argument, since it explains the sign required in the later choice of the quadratic twist.

Lemma 2.2 (Landau principle)**.** *Let $E : [0,\infty) \to \mathbb{R}$ be measurable with $|E(v)| \le Ce^v$. Suppose its Laplace transform*

$$\Phi(s) = \int_0^\infty E(v)e^{(1-s)v}\,dv$$

has a meromorphic continuation, holomorphic near $\operatorname{Re} s \ge 1$ and with at most simple poles in $\operatorname{Re} s > 0$. Let $\beta > 0$. Assume that

$$E(v) \le e^{(\beta-1)v} \quad (v \text{ sufficiently large})$$

and that the residue of Φ at each real $\rho > \beta$ is non-negative (or zero when the point is regular). Then Φ is holomorphic in $\operatorname{Re} s > \beta$.

Proof. Choose $A \ge 1$ large enough that

$$h(v) = Ae^{\beta v} - e^v E(v) \ge 0 \qquad (v \ge 0).$$

The eventual upper bound guarantees the inequality at large v, and the growth assumption gives it on the remaining bounded interval after enlarging A. For sufficiently large $\operatorname{Re} s$ we then have

$$\Psi(s) := \int_0^\infty h(v)e^{-sv}\,dv = \frac{A}{s-\beta} - \Phi(s).$$

Let σ_c be the real abscissa of convergence of the integral for Ψ. Because $h \ge 0$, convergence holds in the right half-plane $\operatorname{Re} s > \sigma_c$, where the integral is holomorphic; the displayed meromorphic expression agrees with it there by analytic continuation.

Suppose $\sigma_c > \beta$. If Ψ were to have a simple pole at the real point σ_c, its residue would necessarily be non-negative: for real $s > \sigma_c$, the integral is non-negative, so $(s-\sigma_c)\Psi(s) \ge 0$. But the meromorphic formula gives that the residue equals minus the residue of Φ and is therefore non-positive. This shows that the residue is zero and the singularity removable.

A Laplace transform of a non-negative function cannot be holomorphic at its finite abscissa of convergence. To see this directly, take a real $c > \sigma_c$ sufficiently close to σ_c that the Taylor series of the continued Ψ about c converges at a real $\sigma < \sigma_c$. Differentiation under the integral at c gives

$$(-1)^j\Psi^{(j)}(c) = \int_0^\infty v^j h(v)e^{-cv}\,dv \ge 0.$$

Taylor convergence and monotone convergence therefore give

$$\begin{aligned}\infty &> \sum_{j\ge 0} \frac{(-1)^j \Psi^{(j)}(c)(c-\sigma)^j}{j!} \\ &= \int_0^\infty h(v) e^{-cv} \sum_{j\ge 0} \frac{((c-\sigma)v)^j}{j!}\, dv \\ &= \int_0^\infty h(v) e^{-\sigma v}\, dv,\end{aligned}$$

contrary to the definition of σ_c. Thus $\sigma_c \le \beta$, and the formula for Ψ implies the assertion about Φ on $\operatorname{Re} s > \beta$. □

This argument is based on positivity of an auxiliary integrand, not on a hypothesis that E changes sign or has small absolute value. It is the usual Landau principle for Dirichlet integrals with one sign, in the form formalized by `IsPolar.landau` in `Lib/Analytic/Landau/PolarLandau.lean`.

2.4 Quadratic twists avoiding a prescribed half-strip

We now give a proof of the zero-density ingredient in the form required here. In particular it is not necessary to invoke a zero-density theorem with constants that are uniform in the original modulus M. Fix M throughout this subsection. The following elementary argument follows the compact source `Lib/Analytic/QuadraticTwist/ZeroFreeTwist.lean`. Jutila's classical work on quadratic-character mean values [3] provides the context; the particular assertion needed here is proved below.

Put $g_j(t) = t^j e^{-\pi t^2}$ for $j = 0, 1$, and, for a sequence c of modulus at most 1, put

$$\mathcal{V}_j(c;t) = \sum_{h\ge 1} c(h) g_j(ht), \qquad I_j(c;T) = \int_{1/T}^\infty |\mathcal{V}_j(c;t)| \frac{dt}{\sqrt{t}}, \qquad T > 0.$$

The following properties of the Gaussian are used repeatedly:

$$|g_j(t)| \le 3e^{-2t} \quad (t \ge 0), \tag{2.9}$$

$$|g_j(t)| \le C_\alpha t^{-\alpha} \quad (t > 0,\ \alpha > 0), \tag{2.10}$$

$$|\mathcal{V}_j(c;t)| \le 3e^{-t}/t. \tag{2.11}$$

For the first inequality use $t^j \le 1 + t \le e^t$ and maximize $3t - \pi t^2$; this maximum is $9/(4\pi) < \log 3$. The second follows from boundedness of $t^\alpha g_j(t)$. For the third, summing the first gives $3e^{-2t}/(1 - e^{-2t}) = 3/(e^{2t} - 1) \le 3e^{-t}/t$, since $2\sinh(t) \ge t$.

Lemma 2.3 (Theta transformation and Mellin integral)**.** *Let ψ be a primitive nonprincipal Dirichlet character of conductor q, and choose $j \in \{0, 1\}$ by $\psi(-1) = (-1)^j$. The Mellin transform on $\Re s > 0$ of $\mathcal{V}_j(\psi;t)$ is*

$$M_\psi(s) = \int_0^\infty \mathcal{V}_j(\psi;t) t^{s-1}\, dt = \Gamma_j(s) L(s,\psi), \qquad \Gamma_j(s) = \tfrac{1}{2}\pi^{-(s+j)/2}\Gamma((s+j)/2).$$

Moreover

$$\sqrt{q}\, t\, |\mathcal{V}_j(\psi;t)| = |\mathcal{V}_j(\psi;1/(qt))|, \tag{2.12}$$

and if $q \le T^2$ then

$$\int_0^\infty |\mathcal{V}_j(\psi;t)| \frac{dt}{\sqrt{t}} \le 2 I_j(\psi;T). \tag{2.13}$$

Proof. Poisson summation on each progression modulo q applies to the Schwartz function g_j on $\mathbb{R}$, whose Fourier transform (with kernel $e^{-2\pi ixy}$) is $(-i)^j g_j$. Expanding the periodic character in finite Fourier coefficients, the Gauss-sum formula for a primitive character gives

$$\mathcal{V}_j(\psi;t) = \frac{\varepsilon(\psi)}{\sqrt{q}\,t}\,\mathcal{V}_j(\overline{\psi};1/(qt)), \quad |\varepsilon(\psi)| = 1.$$

Indeed the two bilateral sums are twice their positive parts, by parity and $\psi(0) = 0$; the modulus of the primitive Gauss sum is $\sqrt{q}$. As g_j is real, the absolute values for ψ and $\overline{\psi}$ coincide. This proves (2.12). At zero the transformation and (2.11) show rapid decay; at infinity (2.11) applies directly. Termwise Mellin integration on $\Re s > 1$ yields the displayed formula, by substituting $u = ht$ and then $z = \pi u^2$. Both sides are holomorphic on $\Re s > 0$, so the equality extends there. Finally the substitution $u = 1/(qt)$ in the integral over $(0, 1/T)$, using (2.12), identifies it with the integral over $(T/q, \infty)$ against $du/\sqrt{u}$. Since $T/q \geq 1/T$, the result follows. □

This is the ordinary theta Mellin formula; compare [2, Ch. 10] and the formal declarations `gaussianMellin_eq`, `norm_gaussianTheta_functional_equation`, `thetaTailIntegral_zero_le_two_mul` in `Lib/Analytic/DirichletTheta/MellinTheta.lean`.

Proposition 2.4 (A three-lines zero detector)**.** *For $j = 0, 1$ there is an absolute constant C_j such that, whenever $\eta > 0$, $T \geq 1$, $q \leq T^2$ and ψ is as in Lemma 2.3, the conditions*

$$I_j(\psi;T) + I_j(\mu\psi;T) \leq T^{2\eta}, \qquad L(s,\psi) = 0, \quad \tfrac{1}{2} + 9\eta \leq \Re s \leq 1$$

imply $|\Gamma_j(s)|^2 \leq C_j T^{-2\eta}$.

Proof. Here μ is the Möbius function. Define a truncated Mellin integral

$$G(s) = \int_{1/T}^{\infty} \mathcal{V}_j(\mu\psi;t)t^{s-1}\,dt, \qquad F(s) = \Gamma_j(s)^2 - M_\psi(s)G(s).$$

It is holomorphic and bounded on the closed strip $\frac{1}{2} \leq \Re s \leq 2$ (with a bound depending on q, T): this follows from the integral estimates just proved, the cutoff in G, and the Gaussian Mellin formula. We use the three-lines theorem with boundary estimates uniform in the imaginary coordinate. On $\Re s = 1/2$, write $I = I_j(\psi;T)$, $J = I_j(\mu\psi;T)$. Lemma 2.3 implies $|M_\psi(s)| \leq 2I$ and $|G(s)| \leq J$. Thus

$$|F(s)| \leq \sup_{\Re z = 1/2} |\Gamma_j(z)|^2 + 2IJ \ll_j T^{4\eta}.$$

The bound on the gamma factor is immediate from its defining integral. On $\Re s = 2$, absolute convergence of the Euler products gives $L(s,\psi)L(s,\mu\psi) = 1$, where the second factor denotes $\sum_{n\geq 1} \mu(n)\psi(n)n^{-s}$. Therefore

$$F(s) = M_\psi(s) \int_0^{1/T} \mathcal{V}_j(\mu\psi;t)t^{s-1}\,dt.$$

By (2.11), the first factor is bounded in modulus by $3\int_0^\infty e^{-t}\,dt = 3$, and the second by $3/T$. Hence $|F(s)| \leq 9/T$, which in particular is $\ll_j T^{4\eta-1}$. Interpolation across the strip of width $3/2$ gives

$$|F(s)| \ll_j T^{4\eta-(\Re s-1/2)/(3/2)}.$$

When $L(s,\psi) = 0$, $F(s) = \Gamma_j(s)^2$. For $\Re s \geq 1/2 + 9\eta$, the exponent is at most -2η, as required. □

The formal counterpart is `three_lines_bound_of_LFunction_eq_zero`.

Lemma 2.5 (Smoothed quadratic-character mean value)**.** *For $\eta > 0$, $j = 0, 1$, $X \in \mathbb{N}$, $0 < T$ with $X \le T^2$, $t \ge 1/T$, and $|c(h)| \le 1$, one has*

$$\sum_{1 \le v \le X} \left| \sum_{h \ge 1} c(h) J(-v \mid h) g_j(ht) \right|^2 \ll_\eta T^2 t^{-1-2\eta}.$$

Here $J(v \mid h)$ is the product of the Legendre symbols over odd primes dividing h, each with its multiplicity, and the principal character at 2 when h is even. In particular it is multiplicative in the denominator and bounded in absolute value by 1.

Proof. After expanding the square (absolutely convergent by (2.11)), it is enough to bound

$$\sum_{n,m \ge 1} |g_j(nt) g_j(mt)| \left| \sum_{v \le X} J(-v \mid nm) \right|.$$

If neither nm nor $2nm$ is a square, the Dirichlet character $v \mapsto J(v \mid nm)$ is nonprincipal modulo nm: its odd-prime squarefree kernel is nontrivial. The Pólya–Vinogradov inequality gives the inner bound $\ll \sqrt{nm}(1 + \log nm)$. Use (2.10) with exponent $3/2 + \eta$ for both Gaussian factors and the estimate $1 + \log(nm) \ll_\eta (nm)^{\eta/2}$. The sum of the resulting factors $(nm)^{-1-\eta/2}$ is finite. The contribution of this part is $\ll_\eta t^{-3-2\eta} \le T^2 t^{-1-2\eta}$.

If knm is a square for some $k \in \{1, 2\}$, use the trivial bound $X \le T^2$. Write $kn = au^2$, $m = av^2$ with $a \ge 1$ squarefree; each pair occurs in this parametrization for the indicated k. Use (2.10) with exponent $1/2 + \eta$. Ignoring the squarefree restriction on a and the divisibility requirement on u, the sum is bounded, for each k, by

$$C_\eta T^2 t^{-1-2\eta} k^{1/2+\eta} \sum_{a,u,v \ge 1} (auv)^{-1-2\eta}.$$

This proves the assertion. We have used only the classical Pólya–Vinogradov bound for a (possibly imprimitive) nonprincipal character, not a quadratic large sieve. The finite-character inequality can be found in [2, Ch. 9]. □

Corollary 2.6 (Tail-integral bound)**.** *For $X \le T^2$ and $|c(h)| \le 1$,*

$$\sum_{1 \le v \le X} I_j(c(\cdot) J(-v \mid \cdot); T) \ll_\eta T^{2+\eta}.$$

Proof. By Cauchy–Schwarz and Lemma 2.5, the sum over v of the absolute values inside the integral is $\ll_\eta \sqrt{X} T t^{-1/2-\eta} \le C_\eta T^2 t^{-1/2-\eta}$. Integrating against $dt/\sqrt{t}$ from $1/T$ to infinity gives T^η/η times this constant. □

Proposition 2.7 (Zero density in the family of prime quadratic twists)**.** *Fix a Dirichlet character ξ modulo M, $Y \ge 0$, and $0 < \eta < 1/18$. The number of primes $d \le X$, $d \equiv 7 \pmod 8$, $(d, M) = 1$, for which $\widetilde{L}(s, \xi\chi_K)$ has a zero in $\Re s \ge 1/2 + 9\eta$, $|\Im s| \le Y$, is $O_{M,Y,\eta}(X^{1-\eta/2})$.*

Proof. Let ψ of conductor $M' \mid M$ induce ξ. The twist $\psi\chi_K$ is primitive of conductor $q = M'd$, by coprimality of the conductors. As $d \equiv 7 \pmod 8$, its parity j is independent of d and quadratic reciprocity gives $\chi_K(h) = J(-d \mid h)$, including the factor at 2. The extra Euler

factors for $\xi\chi_K$ are nonzero at $\Re s > 0$, and classical non-vanishing gives $\Re s < 1$ for any indicated zero. Put $T = \sqrt{M'X}$. By the zero detector, if

$$I_j(\psi\chi_K; T) + I_j(\mu\psi\chi_K; T) \le T^{2\eta},$$

then $|\Gamma_j(s)|^2 \ll T^{-2\eta}$ at the putative zero. But Γ_j has a positive lower bound in modulus on the fixed compact rectangle $1/2 \le \Re s \le 1$, $|\Im s| \le Y$; hence for large enough X the putative zero forces the sum of tail integrals to be larger than $T^{2\eta}$. Corollary 2.6, for $c = \psi$ and $c = \mu\psi$, bounds the sum of the tail integrals over $d \le X$ by $O_\eta(T^{2+\eta})$. Dividing by $T^{2\eta}$ proves the announced count. □

The mean-square, tail-integral and density proofs above translate `sum_sq_gaussianTheta_jacobiSym_le`, `sum_thetaTailIntegral_le`, and `zero_density_quadTwist` in `ZeroFreeTwist.lean`.

Proposition 2.8 (Auxiliary zero-free twists)**.** *Fix a modulus M, a unit class $a_0 \pmod M$ all of whose primes are $7 \pmod 8$ and prime to M, a number $\beta > 1/2$ and $Y \ge 0$. There exists a prime $d \equiv a_0 \pmod M$ such that every $\widetilde{L}(s, \xi\chi_K)$, $\xi \bmod M$, is nonzero in $\Re s \ge \beta$, $|\Im s| \le Y$.*

Proof. Choose $0 < \eta < 1/18$ such that $1/2 + 9\eta < \beta$. By Proposition 2.7 the union, over the finitely many characters modulo M, of the bad sets has count $O_{M,Y,\eta}(X^{1-\eta/2})$. Its sum of reciprocals of primes converges by dyadic summation. The sum of reciprocals of all primes in the unit class a_0 diverges, by the elementary reciprocal-prime argument in Section 3.6; hence the union does not contain that class. Choose d outside the union. This is `exists_entireL_quadTwist_ne_zero` in the formal source. □

2.5 The one-sided zero-free criterion

Proposition 2.9 (Criterion used in the proof)**.** *Let $\beta > 1/2$, and let M, a_0 satisfy the conditions of Proposition 2.8. Suppose that for* each *prime $d \equiv a_0 \pmod M$ the corresponding prime error (2.3) satisfies*

$$R_M(x) \le x^{\beta-1} \qquad (x \textit{ sufficiently large}).$$

Then all $\widetilde{L}(s, \xi)$ for $\xi \bmod M$ are zero-free in the half-plane $\operatorname{Re} s > \beta$.

Proof. If $\widetilde{L}(\rho, \xi_0) = 0$ for some ρ to the right of β, apply Proposition 2.8 with $Y = |\operatorname{Im}\rho|$. Choose d so that no twisted $\widetilde{L}(s, \xi\chi_K)$ vanishes anywhere in the indicated half-strip. Then, at ρ, the numerator of (2.8) is a positive integer: the first sum includes the zero of $\widetilde{L}(\cdot, \xi_0)$, while the second sum is zero.

Furthermore at each real point $\sigma \in (\beta, 1)$ the numerator of the same residue formula is non-negative and $\sigma(1-\sigma) > 0$. Thus the real residues of Φ_M to the right of β are non-negative; those at and to the right of 1 are zero by classical non-vanishing and normalization. The assumed bound for the chosen d gives $E_M(v) \le e^{(\beta-1)v}$ eventually. The Landau principle (Lemma 2.2) says that Φ_M is holomorphic on $\operatorname{Re} s > \beta$, contradicting its non-zero residue at ρ. □

The logical order matters: the same arithmetic inequality has to be proved for each admissible auxiliary prime d in the given class, while the proof of the criterion chooses d *after* assuming that an offending zero exists. The threshold in the arithmetic estimate need not be uniform in d, and the choice of d need not work for all zeroes at once. This flexibility is what

avoids a zero-cancellation problem that an argument involving just one fixed quadratic field would face.

The formal declaration is `entireL_ne_zero_of_primeError_le` in `New/Endgame/OneSidedCriterion.lean`.

2.6 A good second scale in a long window

The compact edition employs a second important analytic argument. The arithmetic construction gives a bound involving two scales x, x_1 rather than one, with a combination of the two prime errors $R_{N_0}(x_1)$ and $R_1(x_1)$. We cannot simply discard those terms: their size relative to the desired saving at x would need justification. Instead, we have the freedom to vary $\log x_1$ in a sufficiently long interval. A high-moment argument produces a good choice of the second scale without requiring a pointwise prime number theorem error estimate there.

Lemma 2.10 (Window lemma)**.** *Let $E, E' : [0,\infty) \to \mathbb{R}$ be measurable with $|E(v)| + |E'(v)| \le Ce^v$. Suppose their Laplace transforms, with kernel $e^{(1-s)v}$, continue holomorphically to a neighborhood of the closed half-plane $\operatorname{Re} s \ge 1$. For every sufficiently large T and every $0 \le X \le e^T$ there exists $u \in [T, T^2]$ such that*

$$E(u) - XE'(u) \le e^{-T}.$$

The threshold for T may depend on the two functions but not on X in this range.

Proof. The closed disc of radius 3 about 4 lies in $\operatorname{Re} s \ge 1$. Choose a fixed $3 < R < 4$ such that both continued transforms are holomorphic on a neighborhood of the closed disc of radius R about 4. Such an R exists by compactness. Differentiation under the original integral near $s = 4$ gives

$$\Phi^{(j)}(4) = \int_0^\infty (-v)^j e^{-3v} E(v)\, dv.$$

The Cauchy estimates for Taylor coefficients yield, for all sufficiently large T, simultaneously for every integer $j \ge 0$,

$$\left| \int_0^\infty v^j e^{-3v} E(v)\, dv \right| \le \frac{j!T}{R^j}.$$

The same holds for E'. We can also require $|E(v)|, |E'(v)| \le Te^v$ for all $v \ge 0$.

Write $k = \lceil T^2/9 \rceil$ and suppose, for a contradiction, that $E(v) - XE'(v) > e^{-T}$ throughout $[T, T^2]$. Consider the function

$$v^k \left(e^{7T - Rv} + e^{2T - T^2 - v} + e^{-3v}(E(v) - XE'(v) - e^{-T}) \right).$$

On the window this is non-negative by hypothesis. Outside the window, the elementary growth bounds give

$$e^{-3v}|E(v) - XE'(v) - e^{-T}| \le e^{2T-2v} \qquad (v \ge 0),$$

for sufficiently large T, uniformly in $0 \le X \le e^T$. For $0 \le v \le T$, the first exponential in the parentheses dominates this upper bound, since $R < 4$. For $v \ge T^2$, the second exponential dominates it. Thus the integrand is non-negative everywhere.

Integrate over $v \ge 0$, use the moment bounds above, and recall $\int_0^\infty v^k e^{-qv} dv = k!/q^{k+1}$ for $q > 0$. The result, after division by $k!$, is

$$0 \le \frac{e^{7T}}{R^{k+1}} + e^{2T-T^2} + \frac{T(1+X)}{R^k} - \frac{e^{-T}}{3^{k+1}}.$$

This is impossible for large T. Divide all terms by the magnitude $e^{-T}/3^{k+1}$ of the negative term. The three positive ratios are respectively

$$e^{8T}(3/R)^{k+1}, \quad e^{3T-T^2}3^{k+1}, \quad 3e^{T}T(1+X)(3/R)^{k}.$$

The first and third tend to zero as $T \to \infty$, since k is quadratic in T and $3/R < 1$; the convergence is uniform for the third ratio by $X \le e^T$. The second ratio tends to zero because $(\log 3)/9 < 1$. Their sum is less than one for all sufficiently large T, giving the desired contradiction. □

This proof is an explicit example of what a genuine shortening of the formal development looks like when translated into ordinary mathematics. The readable edition uses a pointwise prime number theorem error bound at its much larger second cutoff. The compact edition needs only regularity of the two transforms in $\operatorname{Re} s \ge 1$ and a good point in a window. It is the ability to choose that point, not an assumed strong prime number theorem, that makes the argument work.

Apply the lemma with $E(v) = R_{N_0}(e^v)$ and $E'(v) = R_1(e^v)$ for the same auxiliary prime d. The hypotheses follow from Proposition 2.1. If $T = 10^6 \log x$ and X is the positive multiplier supplied by the arithmetic data, with $X \le e^T$, the lemma permits the choice $x_1 = e^u$ with

$$\log x_1 \in [T, T^2], \qquad R_{N_0}(x_1) - XR_1(x_1) \le x^{-10^6}. \tag{2.14}$$

This is a one-sided bound for a particular linear combination at a particular point; it should not be confused with a uniform pointwise estimate for either error. In the final argument the change of x_1 has to preserve the admissibility of the automorphic data. The scales are chosen in the arithmetic construction precisely so that this is possible.

Source. The moment and window lemmas are `IsPolar.eventually_moment_le` and `IsPolar.exists_window_sub_mul_le` in `New/Endgame/OneSidedCriterion.lean`; `New/Main.lean` applies them to the remainder in the final bound `eventually_primeError_le_rpow`.

2.7 From the target bound to GRH

To see precisely the quantifiers needed at the end, suppose that (2.6) has been established by the arithmetic construction. Let $q \ge 1$ be any modulus, $\chi \bmod q$ a Dirichlet character, and ρ a hypothetical zero with $\operatorname{Re}\rho > 1/2$, with the pole of a principal character excepted. Take $N_0 = 18480q$ and the unit class $a_0 = N_0/2 - 1 \pmod{N_0}$. Since $16 \mid N_0$, $a_0 \equiv 7 \pmod 8$ and $a_0^2 \equiv 1 \pmod{N_0}$. Choose $\varepsilon > 0$ with $\varepsilon < 1/100$ and $\beta = 1/2 + 8\varepsilon < \operatorname{Re}\rho$. The estimate (2.6) holds for every prime $d \equiv a_0 \pmod{N_0}$ and gives $R_{N_0}(x) \le x^{\beta-1}$ eventually. Proposition 2.9 implies that all pole-removed Dirichlet L-functions modulo N_0 are non-zero at ρ.

Changing the level of χ to N_0 multiplies its Dirichlet L-function by a finite product of omitted Euler factors, each non-zero in $\operatorname{Re} s > 0$. Thus a zero of $\widetilde{L}(\cdot, \chi)$ at ρ persists at the new level. This contradicts the preceding conclusion. We have proved, subject only to the arithmetic target estimate,

$$L(s, \chi) \ne 0 \qquad (\operatorname{Re} s > 1/2),$$

apart from the pole of the principal character at $s = 1$.

For a primitive character the functional equation, whose gamma factors are non-zero and finite in the open critical strip, sends a zero at s to a zero for the dual character at $1 - s$. Thus there can be no zero with $0 < \operatorname{Re} s < 1/2$ either. For an imprimitive character the omitted Euler factors again do not vanish in the open strip. Hence every zero there has real part $1/2$.

This includes the Riemann zeta function (modulus one); its standard functional equation and trivial-zero classification give the usual full statement of RH.

The rest of the article is concerned with the provenance of (2.6): how a problem about the lengths of automorphic forms and the divisibility of their theta lifts yields exactly the exponent needed for this analytic reduction.

Chapter 3

Classical analytic foundations without a prime number theorem

The two analytic uses of prime distribution in Chapter 2 and in the construction of the admissible record deserve particular attention. The first is the selection of the auxiliary quadratic twist; the second is the choice of the large split level primes. The former requires only divergence of a reciprocal-prime sum in a fixed residue class and an upper bound on the number of bad twists. The latter requires existence, not an asymptotic count, of a prime in a long window at a modulus that is a growing power of a fixed integer. Neither application requires the prime number theorem for arithmetic progressions. We prove these facts here, including the limited uniformity in the modulus needed for the prime windows.

This chapter supplements, rather than changes, the analytic inequalities proved earlier. In particular no zero-free region to the left of the line $\Re s = 1$ is used for a varying family of all possible moduli. The elementary disk in Lemma 3.9 is for the restricted family of powers of a fixed modulus. That restriction makes a substantial difference: its real characters are induced from one fixed finite set.

3.1 Finite characters, Gauss sums, and the Fourier bound

Write $e(t) = \exp(2\pi i t)$. For a positive integer q, a character modulo q is the extension by zero to the nonunits of a homomorphism $(\mathbb{Z}/q)^\times \to \mathbb{C}^\times$. The extension is periodic and completely multiplicative. The character of the trivial group for $q = 1$ takes the value 1 on every integer.

The existence and elementary properties of these characters can all be read from the structure theorem for finite abelian groups: for a cyclic group of order n the characters are given by the n-th roots of unity, and the characters of a finite direct product are tensor products. Thus there are $\varphi(q)$ characters and

$$\frac{1}{\varphi(q)} \sum_{\chi \bmod q} \overline{\chi(a)}\chi(n) = \mathbf{1}_{n \equiv a \bmod q} \quad ((a,q) = 1). \tag{3.1}$$

Indeed the sum vanishes on a nonidentity unit by multiplying it by a character which is nontrivial on that unit. On a nonunit both sides vanish. The same argument with the roles of the group and its character group exchanged proves that the sum of a nonprincipal character over a complete residue system is zero.

Lemma 3.1 (Conductor and change of level)**.** *For each character χ modulo q there is a*

unique primitive character ψ modulo some divisor $q' \mid q$ inducing χ. If $q = q'r$, then

$$\chi(n) = \psi(n)\mathbf{1}_{(n,r)=1}, \qquad L(s,\chi) = L(s,\psi)\prod_{\substack{p\mid q\\ p\nmid q'}}(1-\psi(p)p^{-s}). \tag{3.2}$$

If ψ_1, ψ_2 are primitive of coprime conductors q_1, q_2, their product is primitive of conductor q_1q_2.

Proof. Factor the unit group by the Chinese remainder theorem into groups modulo prime powers. For each $p^\alpha \| q$, choose the smallest $\beta \le \alpha$ such that the local component factors through reduction modulo p^β (allow $\beta = 0$ for the trivial factor). The product q' of the p^β is the conductor, and the factorization defines the unique induced primitive character. Its extension by zero detects the primes dividing q'; the indicator in (3.2) detects all remaining primes dividing q. The asserted Euler-factor identity follows in the domain of absolute convergence, and hence throughout continuation. In the coprime case the local components of the product are precisely the original local components, so their minimal exponents are unchanged. □

For a character ψ modulo q put

$$\tau_a(\psi) = \sum_{x \bmod q} \psi(x)e(ax/q), \qquad \tau(\psi) = \tau_1(\psi).$$

Lemma 3.2 (Primitive Gauss sums)**.** *If ψ is primitive, then*

$$\tau_a(\psi) = \overline{\psi(a)}\tau(\psi), \qquad |\tau(\psi)| = \sqrt{q}. \tag{3.3}$$

In particular, for a real primitive character of prime conductor d,

$$\overline{\tau(\psi)} = \psi(-1)\tau(\psi), \qquad \tau(\psi)^2 = \psi(-1)d.$$

Proof. For a a unit change variables by multiplication by a^{-1}. If a is not a unit, the additive character $x \mapsto e(ax/q)$ factors through $m = q/(a,q)$, a proper divisor of q. Since ψ is primitive, there is a unit u congruent to 1 modulo m with $\psi(u) \ne 1$. (The reduction of unit groups is surjective, as is seen prime-power by prime-power.) Changing variables by u shows that $\tau_a(\psi) = \psi(u)\tau_a(\psi)$, so this sum is zero. This proves the first assertion. Orthogonality of the finite additive characters gives

$$\sum_{a \bmod q} |\tau_a(\psi)|^2 = q \sum_{x \bmod q} |\psi(x)|^2 = q\varphi(q).$$

There are $\varphi(q)$ nonzero summands on the left, all of modulus $|\tau(\psi)|^2$. This proves the second assertion. For a real character, conjugation of the sum changes $a = 1$ to $a = -1$, giving the remaining identities. □

The following proof includes the imprimitive case, since the imprimitive character arising from the generalized Jacobi symbol in Chapter 2 need not have conductor equal to its written modulus.

Proposition 3.3 (Pólya–Vinogradov with the source constant)**.** *If χ is nonprincipal modulo q, then for every integer $X \ge 0$,*

$$\left|\sum_{1\le n\le X} \chi(n)\right| \le 2\sqrt{q}(1+\log q).$$

For primitive nonprincipal χ the factor 2 may be omitted.

Proof. First suppose χ primitive; then $q \geq 2$ and $\tau_0(\chi) = 0$. Fourier inversion and Lemma 3.2 give

$$\left|\sum_{n\leq X}\chi(n)\right| \leq q^{-1/2}\sum_{a=1}^{q-1}\left|\sum_{n=1}^{X} e(-an/q)\right|.$$

The inner geometric sum is bounded by $|\sin(\pi a/q)|^{-1}$. On $0 \leq v \leq 1/2$ the concavity of $\sin(\pi v)$ implies $\sin(\pi v) \geq 2v$, so that the inner bound is at most $q/(2\min(a, q-a))$. Sum first over $a \leq q/2$ and then over $a > q/2$. Since $\sum_{1\leq j\leq q/2} j^{-1} \leq 1 + \log q$, the primitive assertion follows. For general χ let $q = q'r$ as in Lemma 3.1; the inducing ψ is nonprincipal. Möbius inversion of the coprimality indicator gives

$$\sum_{n\leq X}\chi(n) = \sum_{d\mid r}\mu(d)\psi(d)\sum_{j\leq X/d}\psi(j).$$

The number of divisors of r is at most $2\sqrt{r}$ by pairing a divisor with its complementary divisor. Thus the already proved bound for ψ gives $2\sqrt{r}\sqrt{q'}(1+\log q')$, as required. □

This is the finite Fourier argument of `norm_gaussSum_of_isPrimitive` and `polya_vinogradov` in `Lib/Analytic/PolyaVinogradov/PolyaVinogradov.lean`. In particular no large-sieve theorem has been used.

3.2 The quadratic symbol and the choice of the family

Here are the elementary details behind the compatibility of the prime family with the Jacobi symbol used in the mean square. For an odd prime p, let (a/p) be 0 if $p \mid a$, 1 if a is a nonzero square modulo p, and -1 otherwise. For $p \nmid a$ the equality $(a/p) \equiv a^{(p-1)/2} \pmod p$ follows from cyclicity of the multiplicative group of the finite field.

For clarity we derive the particular reciprocity rule needed here. Put $h = (p-1)/2$. The signed least residues of $a, 2a, \ldots, ha$ have distinct absolute values: equality of absolute values would give $ai \equiv \pm aj \pmod p$, and hence $i = j$ in the indicated range. Their absolute values are therefore $1, \ldots, h$. Comparing products modulo p gives Gauss's lemma $(a/p) = (-1)^m$, where m counts those least nonnegative residues larger than $p/2$. If a is odd, comparison of sums modulo 2 gives

$$m \equiv \sum_{i=1}^{h}\lfloor ai/p\rfloor \pmod 2 :$$

indeed the sum of the signed residues is $a\sum i - p\sum\lfloor ai/p\rfloor - pm$, while modulo two it equals the sum of their absolute values. If q is another odd prime, count the integral points of the rectangle

$$1 \leq i \leq (p-1)/2, \qquad 1 \leq j \leq (q-1)/2$$

on the two sides of $qi = pj$. There is no point on the separating line. The two counts are respectively $\sum_i\lfloor qi/p\rfloor$ and $\sum_j\lfloor pj/q\rfloor$. Their sum is $(p-1)(q-1)/4$. Consequently

$$(q/p)(p/q) = (-1)^{(p-1)(q-1)/4}.$$

Likewise $(-1/p) = (-1)^{(p-1)/2}$. Applying Gauss's lemma to $a = 2$ gives $m = h - \lfloor p/4\rfloor$, whose parity by the four odd residue classes modulo 8 is $(p^2-1)/8$. This yields $(2/p) = (-1)^{(p^2-1)/8}$.

Let now $d \equiv 7 \pmod 8$ be prime, and $K = \mathbb{Q}(\sqrt{-d})$. The discriminant of K is $-d$. Its quadratic character is $\chi_K(n) = (n/d)$. For any odd prime $p \neq d$, reciprocity gives

$(-d/p) = (p/d)$; at $p = d$ both sides are zero. Moreover $(2/d) = 1$. Hence, with the generalized Jacobi symbol J of Chapter 2,

$$J(-d \mid n) = \chi_K(n) \qquad (n \geq 1). \tag{3.4}$$

Here the contribution to $J(-d \mid n)$ of an even factor in n is 1 because d is odd. The identity extends to all n by multiplicativity. The usual splitting characterization follows directly by reducing the minimal polynomial $X^2 - X + (d+1)/4$ of $(1 + \sqrt{-d})/2$ modulo p, with the 2-adic case giving two distinct roots because $(d+1)/4$ is even. Thus $\chi_K(p) = 1$ exactly at split primes and $\chi_K(p) = -1$ at inert primes, with $\chi_K(d) = 0$. Lemma 3.2 also shows that $\mathbb{Q}(\sqrt{-d})$ is contained in $\mathbb{Q}(\zeta_d)$ by the Gauss sum of its quadratic character. Neither use of quadratic reciprocity involves a choice of the sign of that square root.

For an arbitrary positive denominator h, the function $v \mapsto J(-v \mid h)$ is a constant sign times a (possibly principal) character modulo h. If at least one odd prime occurs in h to an odd exponent, this character is nonprincipal: choose a nonsquare unit at that prime and a square unit at all other primes, using the Chinese remainder theorem. Otherwise its value on every unit is 1 and it is principal, whether or not the exponent of 2 is odd. Hence

$$\left| \sum_{1 \leq v \leq X} J(-v \mid h) \right| \leq \begin{cases} X, & h \text{ or } 2h \text{ is a square,} \\ 2\sqrt{h}(1 + \log h), & \text{otherwise.} \end{cases} \tag{3.5}$$

This is exactly the dichotomy needed in the expansion of the quadratic-character mean square. Note that the modulus h here is the product of the two summation indices, not the discriminant parameter being averaged. This distinction explains why an ordinary Pólya–Vinogradov estimate suffices.

3.3 Holomorphy and nonvanishing on the line at one

We give direct proofs of the classical analytic properties of the Dirichlet L-functions that are used by the polar transform, including the case of a real character at $s = 1$. It is essential not to use the desired zero-free result in this step.

For a nonprincipal character χ modulo q, let $S_\chi(t) = \sum_{1 \leq n \leq t} \chi(n)$, with the real variable $t \geq 1$. Periodicity and the zero sum over a full period imply

$$|S_\chi(t)| \leq \min(q, t). \tag{3.6}$$

Partial summation gives

$$L(s, \chi) = s \int_1^\infty S_\chi(t) t^{-s-1}\, dt \qquad (\Re s > 0).$$

Initially this is equality with the Dirichlet series on $\Re s > 1$. The right hand side is holomorphic on $\Re s > 0$ by uniform convergence on compact subsets, so provides continuation there. In the principal case the corresponding formula for $\zeta(s)$ is

$$\zeta(s) = \frac{s}{s-1} - s \int_1^\infty \{t\} t^{-s-1}\, dt, \tag{3.7}$$

where $\{t\} = t - \lfloor t \rfloor$. This yields a simple pole of residue 1 at $s = 1$ and holomorphy elsewhere on $\Re s > 0$. The principal character modulo q has

$$L(s, \chi_0) = \zeta(s) \prod_{p \mid q} (1 - p^{-s}),$$

and its residue at 1 is strictly positive. Also

$$[(s-1)\zeta(s)]'_{s=1} = \gamma. \tag{3.8}$$

Indeed the constant term of (3.7) is $1 - \int_1^\infty \{t\}t^{-2}\,dt$. Summing on integer intervals gives $\int_1^n \{t\}t^{-2}\,dt = \log n - H_n + 1$; as $n \to \infty$ the assertion follows from the definition $\gamma = \lim_n (H_n - \log n)$.

The Euler products over p on $\Re s > 1$ are absolutely convergent and yield

$$\log L(s,\chi) = \sum_p \sum_{j\geq 1} \frac{\chi(p)^j}{jp^{js}}, \tag{3.9}$$

where the logarithm is specified by this analytic series. Absolute convergence follows by comparing to $\sum_{n\geq 2} n^{-\Re s}$. In particular all the L-functions are nonzero in this open half-plane, including the principal one. Differentiating gives the usual logarithmic derivative series, since the factor $\log p$ introduced by differentiation still leaves a convergent series on each smaller half-plane $\Re s \geq 1 + \delta$.

We recall the precise positivity principle for Dirichlet series that is needed for the boundary case.

Lemma 3.4 (Positive Dirichlet series)**.** *Suppose $a(n) \geq 0$ and $\sum_n a(n)n^{-s}$ has a finite abscissa σ_c of convergence. Its analytic function on $\Re s > \sigma_c$ cannot extend holomorphically to a neighborhood of the real point σ_c.*

Proof. Convergence at a real point implies absolute and locally uniform convergence on the right, by positivity. In the half-plane of convergence every derivative is obtained by introducing the corresponding power of $-\log n$: on a fixed smaller half-plane that power can be absorbed by n^δ for small fixed $\delta > 0$. If the function continued holomorphically near σ_c, choose a real center $c > \sigma_c$ sufficiently close to σ_c that its Taylor series converges at a real point $\sigma < \sigma_c$. Then

$$\sum_{j\geq 0} \frac{(-1)^j F^{(j)}(c)(c-\sigma)^j}{j!} = \sum_n a(n)n^{-c} \sum_{j\geq 0} \frac{((c-\sigma)\log n)^j}{j!}.$$

All terms are nonnegative. The right hand side equals $\sum_n a(n)n^{-\sigma}$ by monotone convergence and the exponential series, but the Taylor series on the left is finite, a contradiction. □

Proposition 3.5 (Classical line nonvanishing)**.** *For every character χ, the pole-removed function $\widetilde{L}(s,\chi)$ of Chapter 2 has no zeros in $\Re s \geq 1$.*

Proof. Nonvanishing for $\Re s > 1$ and the principal case at $s = 1$ have already been shown. First let χ be real and nonprincipal, and suppose $L(1,\chi) = 0$. Then $\zeta(s)L(s,\chi)$ is holomorphic on $\Re s > 0$, since the only possible pole, at 1, would be cancelled. In $\Re s > 1$ its Dirichlet series has coefficients $a(n) = \sum_{d|n} \chi(d)$. This arithmetic function is multiplicative and

$$a(p^j) = 1 + \chi(p) + \cdots + \chi(p)^j.$$

For $\chi(p) \in \{1, 0, -1\}$ its value is always nonnegative and is at least one when j is even. Thus $a(n) \geq 0$ for all n, and $a(m^2) \geq 1$ for all m. Its abscissa of convergence is at least $1/2$ because the subseries on squares diverges there; it is at most 1 by the original absolute convergence. Lemma 3.4 contradicts holomorphy on $\Re s > 0$.

For the remaining case at $s = 1$, consider the finite product $\prod_{\xi \bmod q} L(\sigma, \xi)$ for real $\sigma > 1$. Taking the sum of the Euler logarithms (3.9) and using finite-character orthogonality gives

$$\prod_{\xi \bmod q} L(\sigma, \xi) = \exp\left(\varphi(q) \sum_{p^j \equiv 1 \bmod q} \frac{1}{j p^{j\sigma}}\right) \geq 1.$$

The unique principal factor has a pole of order one at 1. A zero at 1 for a nonreal character would be accompanied by a zero for its distinct conjugate character (conjugate the Dirichlet series at real σ and continue). The product would then tend to zero as $\sigma \downarrow 1$, contrary to the inequality.

Finally for any character χ and any real $t \neq 0$, the Euler products give

$$\zeta(\sigma)^3 |L(\sigma + it, \chi)|^4 |L(\sigma + 2it, \chi^2)| \geq 1 \qquad (\sigma > 1). \tag{3.10}$$

To see this take the real logarithm of the left hand side using (3.9). At a prime not dividing q the contribution of each p^j is the positive weight $p^{-j\sigma}/j$ times $3+4\cos v+\cos(2v) = 2(1+\cos v)^2 \geq 0$, where v is the argument of $\chi(p)^j p^{-ijt}$. At a prime dividing q the contribution from the two L-functions is zero and that from ζ is positive. If $L(1 + it, \chi) = 0$, then its fourth power in modulus is $O((\sigma-1)^4)$. The ζ factor is $O((\sigma-1)^{-3})$ and the χ^2 factor is bounded, since $2t \neq 0$ even if χ^2 is principal. Their product would tend to zero, contradicting (3.10). Removing the principal pole does not affect any of these nonzeros away from 1. □

It follows in particular that the normalization $\widetilde{L}$ in the prime-error transform is nonzero at 1, and that the entire function in the source has the claimed locally nonvanishing restriction on the closed right half-plane. The proof did not use a prime number theorem or a zero-free region of the form $\Re s > 1 - c/\log(q(|\Im s| + 2))$. Equation (3.10) is the Euler-product positivity used by `norm_LFunction_product_ge_one` in the analytic library; its quantitative use is made below.

3.4 Reciprocal primes in a fixed residue class

Proposition 3.6 (The required elementary Dirichlet theorem)**.** *For every positive integer q and every unit class $a \bmod q$ the series*

$$\sum_{p \equiv a \bmod q} \frac{1}{p}$$

diverges. More precisely,

$$\sum_{p \equiv a \bmod q} p^{-\sigma} = \frac{1}{\varphi(q)} \log \frac{1}{\sigma - 1} + O_q(1) \quad (\sigma \downarrow 1,\ \sigma > 1).$$

Proof. For each $\chi \bmod q$ use the Euler logarithm (3.9) at the real argument $\sigma > 1$. Terms with $j \geq 2$ have total absolute value bounded independently of σ, since

$$\sum_p \sum_{j \geq 2} p^{-j}/j \leq \sum_{n \geq 2} \frac{1}{n(n-1)} = 1.$$

Thus by (3.1),

$$\sum_{p \equiv a \bmod q} p^{-\sigma} = \frac{1}{\varphi(q)} \Re \sum_{\chi \bmod q} \overline{\chi(a)} \log L(\sigma, \chi) + O(1).$$

The principal Euler logarithm is real and equals $\log(1/(\sigma-1))+O_q(1)$ by its positive residue at 1. By Proposition 3.5, a nonprincipal L-function is holomorphic and nonzero on some disk about 1 and hence admits a holomorphic logarithm there. To see that this logarithm stays compatible with the Euler logarithm, note that their difference on the real interval to the right of 1 is a continuous function in $2\pi i\mathbb{Z}$, and is therefore constant. Thus each nonprincipal Euler logarithm is bounded as $\sigma \downarrow 1$. This proves the estimate. If the reciprocal-prime series converged, each of the Dirichlet sums on the left would be bounded by that fixed sum, a contradiction. □

In particular, the choice of the auxiliary twist in Proposition 2.8 now follows without the previous appeal to the prime number theorem for a fixed arithmetic progression. The bad set up to X has size $O(X^{1-\eta/2})$. Its reciprocal sum over $2^h \le p < 2^{h+1}$ is $O(2^{-h\eta/2})$, so converges. Proposition 3.6 leaves at least one prime outside the bad set in the prescribed class, and in fact infinitely many. It is unnecessary to prove that a positive proportion of the primes in that class is good by a prime-counting asymptotic.

3.5 The theta transform and the full functional equations

We give details of the theta and Mellin ingredients, including the endpoints of the critical strip. For a Schwartz function f on $\mathbb{R}$, use the Fourier transform

$$\widehat{f}(\xi)=\int_{\mathbb{R}} f(t)e(-t\xi)\,dt.$$

Periodizing f with a positive period h gives a smooth periodic function whose n-th Fourier coefficient is $h^{-1}\widehat{f}(n/h)$, by termwise integration over one period. The series of these coefficients converges absolutely and uniformly (integrate by parts twice in the Fourier transform). Equality with the periodization follows from uniqueness of Fourier coefficients, or from convergence of the Cesàro means of its Fourier series. This is Poisson summation on an arithmetic progression. For $g_0(t)=e^{-\pi t^2}$, the Fourier transform is g_0 (differentiate under the integral and solve the resulting first-order equation, with value 1 at the origin from the Gaussian integral). Since multiplying by t transforms to $-(2\pi i)^{-1}\frac{d}{d\xi}$, the function $g_1(t)=te^{-\pi t^2}$ has transform $-ig_1$.

Take a primitive nonprincipal ψ modulo q, choose $j\in\{0,1\}$ with $\psi(-1)=(-1)^j$, and write $V_j(\psi;t)=\sum_{n\ge1}\psi(n)g_j(nt)$ as in the preceding chapter. Fourier periodization on each residue class and Lemma 3.2 give the exact (not just absolute-value) identity

$$V_j(\psi;t)=\frac{\epsilon_\psi}{\sqrt{q}\,t}V_j(\bar\psi;1/(qt)),\qquad \epsilon_\psi=(-i)^j\tau(\psi)/\sqrt{q},\quad |\epsilon_\psi|=1. \tag{3.11}$$

Indeed the bilateral sum is twice its positive part (the $n=0$ term is zero), and the finite Fourier coefficient for a is $\tau(\psi)\bar\psi(a)$. The factor $(-i)^j$ is determined by the Fourier convention above. The bilateral Poisson formula has exactly the factor $1/(qt)$ before the Gauss sum.

Define $\Gamma_j(s)=\frac12\pi^{-(s+j)/2}\Gamma((s+j)/2)$ and

$$M_\psi(s)=\int_0^\infty V_j(\psi;t)t^{s-1}\,dt.$$

At infinity, the Gaussian bound of Chapter 2 is exponential, and at zero (3.11) gives exponential decay in $1/t$ times a power of $1/t$. Thus this integral is entire: on any compact set of s all derivatives obtained by multiplying the integrand by powers of $\log t$ are absolutely integrable uniformly. Substituting each nt in the sum shows that $M_\psi(s)=\Gamma_j(s)L(s,\psi)$ when $\Re s>1$.

The gamma function is nonzero and holomorphic in $\Re z > 0$ and $1/\Gamma(z)$ is entire with simple zeros at the nonpositive integers. For completeness these statements follow from Euler's limit

$$\Gamma(z) = \lim_{N\to\infty} \frac{N^z N!}{z(z+1)\cdots(z+N)} \quad (\Re z > 0).$$

It is obtained by scaling the beta integral $\int_0^1 u^{z-1}(1-u)^N\,du$ by $u = v/N$ and applying dominated convergence; the product formula for the beta integral follows by repeated integration by parts. On taking reciprocals the limit becomes

$$\frac{1}{\Gamma(z)} = ze^{\gamma z}\prod_{n\ge 1}(1+z/n)e^{-z/n}.$$

The logarithms of the nonzero factors in the tail converge normally on every compact set where $n > |z|+1$, by the $O(|z|^2/n^2)$ bound for their logarithms. Hence the reciprocal is entire and has precisely the displayed zeros, each simple. This also establishes the gamma lower bound on fixed compact rectangles used by the zero detector.

It follows that $L(s,\psi)$ extends to an entire function, given by $M_\psi(s)/\Gamma_j(s)$. In (3.11) substitute $u = 1/(qt)$ inside the Mellin integral to obtain

$$M_\psi(s) = \epsilon_\psi q^{1/2-s} M_{\bar\psi}(1-s).$$

Therefore the completed functions

$$\Lambda(s,\psi) = (q/\pi)^{(s+j)/2}\Gamma((s+j)/2)L(s,\psi)$$

satisfy $\Lambda(s,\psi) = \epsilon_\psi \Lambda(1-s,\bar\psi)$. In the open critical strip the gamma factors have neither poles nor zeros; this is exactly the symmetry required to pass from the right-half-plane conclusion of the arithmetic estimate to the full strip for a primitive nonprincipal character. For an imprimitive character the finite factors (3.2) have no zeros in $\Re s > 0$, so the open-strip conclusion follows as well. This also proves entireness of the nonprincipal $\widetilde{L}$ of the previous chapter. No symmetry beyond this open strip is required for the imprimitive functions because their additional Euler factors may vanish on $\Re s = 0$.

The trivial primitive character (modulo 1) is handled separately. Let $V(t) = \sum_{n\ge1} e^{-\pi n^2 t^2}$. Ordinary Poisson summation gives $1 + 2V(t) = t^{-1}(1 + 2V(1/t))$. Split its Mellin integral at $t = 1$ and substitute $u = 1/t$ in the part from 0 to 1. In $\Re s > 1$ this yields

$$\Gamma_0(s)\zeta(s) = \int_1^\infty V(t)(t^{s-1} + t^{-s})\,dt + \frac{1}{2}\left(\frac{1}{s-1} - \frac{1}{s}\right). \tag{3.12}$$

The integral is entire by exponential decay. The right hand side is meromorphic everywhere, with poles only at 0 and 1, and invariant under $s \mapsto 1-s$. Division by $\Gamma_0(s)$ gives a meromorphic continuation of ζ, with the pole at 0 removable and $\zeta(0) = -1/2$. At $s = -2n$ for each integer $n \ge 1$, the right hand side equals its value at $1+2n$, which is positive and finite; Γ_0 has there a simple pole, so these are exactly the standard simple trivial zeros. At any other point with $\Re s < 0$ the right hand side is nonzero by symmetry and the Euler product, and Γ_0 is finite and nonzero. At $s = it \ne 0$ the same reasoning uses Proposition 3.5 at $1 - it$. Hence all remaining zeros of ζ lie in the open critical strip. The theorem for ζ ultimately deduced from the arithmetic inequality therefore agrees with the full trivial-zero exclusion in `src/Statements.lean`, not just with the assertion about zeros strictly inside $0 < \Re s < 1$.

The theta proof matches the statements `mellin_gaussianTheta_char`, `norm_gaussianTheta_functional_equation` and `gaussianMellin_eq` of `Lib/Analytic/DirichletTheta/MellinTheta.lean`. The elementary line-nonvanishing argument was separated from it to avoid a circular use of the functional equations for zero-freeness.

3.6 Analytic details of the bad-twist estimate

The mean-square argument already given in Chapter 2 can be made quite explicit using (3.5). Here we record the complete absolute-convergence and uniformity checks, which are important when the conductor of a quadratic twist itself varies.

For $j = 0, 1$ and any $\alpha > 0$ let

$$C_{j,\alpha} = \sup_{u>0} u^{\alpha} |g_j(u)| < \infty.$$

A direct derivative calculation (maximize $u^{\alpha+j}e^{-\pi u^2}$) gives $C_{j,\alpha} = ((\alpha + j)/(2\pi e))^{(\alpha+j)/2}$. One also has $|g_j(u)| \le 3e^{-2u}$: use $u^j \le e^u$ and $3u - \pi u^2 \le 9/(4\pi) < \log 3$. Hence, for any c with $|c(n)| \le 1$, the series $V_j(c; t)$ converges absolutely, locally uniformly for $t > 0$, and is bounded by $3e^{-t}/t$. Products and finite averages may consequently be expanded into absolutely convergent double sums. The bound $3e^{-t}/t$ also justifies the ordinary Mellin integrals on $\Re s > 1$ without appealing to cancellation.

For the mean square abbreviate $S_h(X) = |\sum_{v \le X} J(-v \mid h)|$. With $n, m \ge 1$ and $X \le T^2$ one needs to bound

$$A(t) = \sum_{n,m \ge 1} |g_j(nt) g_j(mt)| S_{nm}(X), \qquad t \ge 1/T.$$

If nm and $2nm$ are both nonsquares, (3.5) and the power estimate with $\alpha = 3/2 + \eta$ give

$$A_{\mathrm{nonsq}}(t) \le 2C_{j,3/2+\eta}^2 \, t^{-3-2\eta} \sum_{n,m \ge 1} \frac{1 + \log(nm)}{(nm)^{1+\eta}}.$$

For $\eta > 0$ the last series converges. The identity

$$\sum_{n,m \ge 1} (nm)^{-1-\eta}(1 + \log(nm)) = \zeta(1+\eta)^2 - 2\zeta(1+\eta)\zeta'(1+\eta)$$

follows by termwise differentiation at a point of absolute convergence. Thus this part has a constant depending only on j, η, not on the modulus of any original character. Since $t^{-2} \le T^2$, it is at most a constant times $T^2 t^{-1-2\eta}$.

If knm is a square with $k \in \{1, 2\}$, write $kn = au^2$ and $m = av^2$ with a squarefree; this is possible by uniqueness of the odd-exponent kernel in prime factorization. For $\alpha = 1/2 + \eta$ the product of the power estimates for these two variables is bounded by

$$C_{j,\alpha}^2 t^{-2\alpha} k^{\alpha} a^{-2\alpha} (uv)^{-2\alpha}.$$

Use $S_{nm}(X) \le X \le T^2$, sum over a, u, v without their divisibility restrictions, and also sum over the two k. The resulting upper bound is

$$C_{j,\alpha}^2 (1 + 2^{\alpha}) \zeta(1 + 2\eta)^3 T^2 t^{-1-2\eta}.$$

This proves the asserted uniform mean square and supplies explicit choices of its constants. Cauchy–Schwarz with the finite sum over v, followed by integration against $dt/\sqrt{t}$, produces the tail bound $O_{j,\eta}(T^{2+\eta})$, since $\int_{1/T}^{\infty} t^{-1-\eta}\, dt = T^{\eta}/\eta$. Both sums for $c(n) = \psi(n)$ and $c(n) = \mu(n)\psi(n)$ are covered.

We also recall why the interpolation estimate in the zero detector is uniform in the imaginary coordinate. Suppose a bounded holomorphic function f on the interior of a strip $a < \Re z < b$ is continuous and bounded on its closure, with positive edge bounds M_a, M_b. Divide by

$$\exp\left(\frac{b-z}{b-a}\log M_a + \frac{z-a}{b-a}\log M_b\right).$$

The quotient is bounded on the strip and has boundary modulus at most one. Multiply by $\exp(\delta(z-a)^2)$ for $\delta > 0$. On horizontal segments at imaginary parts $\pm R$ its modulus tends uniformly to zero as $R \to \infty$. On vertical edges its modulus is at most $\exp(\delta(b-a)^2)$. The maximum modulus principle on rectangles, followed by $R \to \infty$ and then $\delta \downarrow 0$, gives the three-lines theorem in this situation.

Apply this to the function F of Proposition 2.4. Boundedness of each Mellin integral on its strip follows by taking absolute values: if $a \le \Re s \le b$ then $t^{\Re s - 1}$ is bounded by $t^{a-1} + t^{b-1}$, and the bounds of the previous paragraph give integrability away from zero. For the full M_ψ near zero use the transformation (3.11), whose exponential decay gives integrability for every real exponent. The truncated integral has a fixed positive lower limit $1/T$ and hence no such issue. Γ_j on the strip is bounded from above by its integral representation. Thus the boundary interpolation gives, for constants independent of the twisting prime,

$$|F(s)| \le C_j T^{4\eta - (2/3)(\Re s - 1/2)} \qquad (1/2 \le \Re s \le 2).$$

At a zero with $\Re s \ge 1/2 + 9\eta$ the left side is $|\Gamma_j(s)|^2$ and the right side is at most $C_j T^{-2\eta}$. By the Euler product for the reciprocal gamma function, its modulus on the closed fixed rectangle $1/2 \le \Re s \le 1$, $|\Im s| \le Y$ has a strictly positive lower bound, independently of the prime twist. This is why we can count bad twists using only the first moment of the two tail integrals and obtain the exponent $1 - \eta/2$. The occurrence of the conductor in the zero detector is only through $q \le T^2$, used to fold the lower half of the Mellin integral. If $\xi \bmod M$ is imprimitive, the inducing ψ has fixed conductor $M' \mid M$, and the twists $\psi\chi_K$ are primitive of conductor $M'd$ by Lemma 3.1. Their extra Euler factors at modulus Md cannot introduce zeros in $\Re s > 0$, by (3.2) (each has modulus bounded away from zero at every point with $\Re s > 0$). Thus the bad-twist bound genuinely holds for all the characters in the mean defining R_M, not just the primitive ones.

This completes the classical dependencies of the density and auxiliary-twist propositions. The relevant source anchors are `sum_sq_gaussianTheta_jacobiSym_le`, `sum_thetaTailIntegral_le`, `zero_density_quadTwist`, `three_lines_bound_of_LFunction_eq_zero` and `exists_entireL_quadTwist_ne_zero` in `Lib/Analytic/QuadraticTwist/ZeroFreeTwist.lean`.

3.7 A uniform disk at one for powers of a fixed modulus

We now prove the precise restricted uniformity that replaces the invocation of Siegel–Walfisz in the construction of the arithmetic data. Fix an integer M divisible by 4 and a real number $A > 0$. In this subsection only, $q = M^k$ with $k \ge 2$.

Lemma 3.7 (Real characters have bounded conductor in the tower)**.** *Every real character modulo M^k is induced by one modulo M^2. The corresponding L-functions, including the pole-removed principal function, belong to a single fixed finite collection as k varies.*

Proof. A real character on the unit group is trivial on every square. We show every unit u reducing to 1 modulo M^2 is a square modulo M^k. By the Chinese remainder theorem it suffices to check modulo $p^{kv_p(M)}$. If p is odd, every $u \equiv 1 \pmod p$ has a square root modulo each p^n: starting with 1 modulo p, if $x^2 \equiv u \pmod{p^n}$ and $p \nmid x$, the square of $x + p^n t$ differs from x^2 modulo p^{n+1} by $2xp^n t$. The coefficient $2x$ is invertible modulo p, so a suitable t completes the induction.

At $p = 2$ the assumption $4 \mid M$ gives $u \equiv 1 \pmod{16}$, in particular $u \equiv 1 \pmod 8$. Every such unit is a square modulo any higher power of two. Starting with an odd solution modulo

8, for $n \geq 3$ the two candidates x and $x+2^{n-1}$ have squares differing by $2^n x$ modulo 2^{n+1} (the additional term 2^{2n-2} is divisible by 2^{n+1}). Exactly one corrects the next binary digit of the square as needed. This proves the assertion by induction. Thus the kernel of the unit-group reduction to M^2 is contained in the squares, and every real character factors through M^2. Since M^2 and M^k have the same prime divisors, the values on all integers agree with those of the lower-level character. Their Dirichlet L-functions therefore coincide, proving the final assertion. □

We next give a simple uniform bound in a fixed small disk, using the bounded sums (3.6) rather than a theta estimate.

Lemma 3.8 (Crude bounds in a disk)**.** *For every $0 < r < 1/4$ there is a constant $C_r > 0$ such that for every nonprincipal character χ modulo any $q \geq 2$,*

$$|L(z,\chi)| \leq C_r q^{2r} \quad (|z-1| \leq 2r), \qquad |L'(z,\chi)| \leq C_r q^{2r} \quad (|z-1| \leq r).$$

Proof. On the first disk, $\sigma = \Re z \geq 1 - 2r > 0$ and $|z| \leq 1 + 2r$. From partial summation and (3.6), split the integral at q to obtain

$$\frac{|L(z,\chi)|}{|z|} \leq \int_1^q t^{-\sigma}\, dt + q \int_q^\infty t^{-\sigma-1}\, dt \leq \frac{q^{2r}}{2r} + \frac{q^{2r}}{1-2r}.$$

This gives the first estimate. For $|z-1| \leq r$ apply Cauchy's estimate on the circle of radius r about z contained in the first disk. Absorb $1/r$ into the constant. □

Proposition 3.9 (Small uniform zero-free disk and logarithmic derivative)**.** *Fix M, A as above. For every sufficiently small real $\delta > 0$, uniformly in $k \geq 2$ with $q = M^k \leq \delta^{-A}$, every pole-removed function $\widetilde{L}(s,\chi)$ modulo q is nonzero on $|s-1| \leq \delta$ and satisfies*

$$\delta \left| \frac{\widetilde{L}'(s,\chi)}{\widetilde{L}(s,\chi)} \right| < 1 \qquad (|s-1| \leq \delta).$$

In fact the supremum of this quantity over all these data tends to zero as $\delta \downarrow 0$.

Proof. For real characters, including the principal character, Proposition 3.5 and Lemma 3.7 show that the functions form a fixed finite nonvanishing collection at 1. In each small disk their logarithmic derivatives are bounded; take a common disk and common bound. Multiplication by δ then gives the assertion.

For a nonreal character χ, both χ and χ^2 are nonprincipal. Let $r > 0$ be small enough that $r < 1/4$ and $10Ar < 1$. Assume $\delta < r$. From (3.10) with $t = 0$ and Lemma 3.8,

$$1 \leq \zeta(1+\delta)^3 |L(1+\delta,\chi)|^4 |L(1+\delta,\chi^2)| \leq (2/\delta)^3 |L(1+\delta,\chi)|^4 C_r q^{2r}.$$

Here $\zeta(1+\delta) \leq 1 + 1/\delta \leq 2/\delta$ by integral comparison. Thus

$$|L(1+\delta,\chi)| \geq c_r \delta^{3/4} q^{-r/2}, \qquad c_r = (8C_r)^{-1/4}. \tag{3.13}$$

Every point s of the disk $|s-1| \leq \delta$ may be joined to $1+\delta$ by a segment of length at most 2δ, still in $|s-1| \leq r$. Consequently

$$|L(s,\chi) - L(1+\delta,\chi)| \leq 2C_r \delta q^{2r}.$$

The ratio of this bound to the lower bound in (3.13) is at most $2C_r c_r^{-1} \delta^{1/4} q^{5r/2} \leq 2C_r c_r^{-1} \delta^{1/4 - 5Ar/2}$. By our strict choice of r the exponent is positive (if $A = 0$ the case of bounded q is simpler). For sufficiently small δ this ratio is at most $1/2$ uniformly, so the function is nonzero everywhere on the disk and its modulus at least $\frac{1}{2} c_r \delta^{3/4} q^{-r/2}$. Divide the Cauchy derivative bound by this lower bound and multiply by δ; the same positive power of δ controls the result. Since $\widetilde{L} = L$ for these nonprincipal characters, the proof is complete. □

The conductor restriction cannot simply be removed in this proof: real characters of arbitrarily increasing conductor need not belong to a fixed family. The preceding disk argument is the conventional proof of the special estimate `eventually_mul_norm_deriv_entireL_lt` in `Lib/Analytic/Primes/AuxiliaryPrimes.lean`. It uses no ineffective lower bound for an increasing real quadratic conductor. It is sufficient because the modulus used for the level-prime searches is itself a power of $N_0 d$, which is fixed once the auxiliary quadratic field is chosen.

3.8 Prime existence in long logarithmic windows

We now prove the source's restricted prime-window theorem. The proof is a second, distinct use of a one-sided positivity argument: if the indicated interval contains no primes in the class 1, analytic continuation at 1 conflicts with the obvious positive residue there. We give the complete Cauchy calculation, with polynomial slack on the boundary instead of the smallest possible absolute constants.

Theorem 3.10 (A prime in the required logarithmic window)**.** *Let M be fixed and divisible by 4, and let $A \geq 1$ be fixed. For every sufficiently large real W, every $k \geq 2$ such that $q = M^k \leq W^A$, and every*

$$W \leq U \leq W^A,$$

there is a prime p with

$$p \equiv 1 \pmod{q}, \qquad U \leq \log p \leq U + W.$$

The threshold for W is independent of k, U in these ranges.

Proof. Put $\delta = W^{-1/2}$ and $a(n) = \Lambda(n)\mathbf{1}_{n \equiv 1 \bmod q}$. By Proposition 3.9, applied with exponent $2A$, the function

$$F(z) = \frac{1 - z\sum_{\chi \bmod q} \widetilde{L}'(1+z,\chi)/\widetilde{L}(1+z,\chi)}{\varphi(q)} \tag{3.14}$$

is holomorphic on a neighborhood of the closed disk $|z| \leq \delta$. Its value at zero is $1/\varphi(q)$, and $|F(z)| \leq 1/\varphi(q) + 1 \leq 2$ on the disk: first on the boundary from Proposition 3.9, and then in the interior by the maximum modulus principle. For $\Re z > 0$, character orthogonality and the Euler logarithmic derivative give the identity

$$\sum_{n \geq 1} \frac{a(n)}{n^{1+z}} = \frac{F(z)}{z}. \tag{3.15}$$

The principal pole contributes exactly the 1 in the numerator of (3.14). The series is absolutely convergent in its indicated domain. Let $D = 1/U$. For large W, $0 < D \leq \delta$. In particular,

$$\sum_n a(n) n^{-1-D} = \frac{F(D)}{D} \leq 2U. \tag{3.16}$$

This is the only bound for the full positive Dirichlet series that we will require.

Suppose, to the contrary, that the interval $[U, U+W]$ contains no logarithm of a prime congruent to 1 modulo q. Define the truncated initial sum and the holomorphic function

$$S_U(z) = \sum_{\log n < U} a(n) n^{-1-z}, \quad G(z) = (F(z) - zS_U(z)) \exp\left(Uz + \frac{\sqrt{W}}{4}(1 + \sqrt{W}z)^2\right).$$

The initial sum is finite. On the circle $|z| = \delta$, write $v = \Re z$ and set

$$t = \frac{W}{2}(1 + \sqrt{W}\, v).$$

Since $-\delta \le v \le \delta$, we have $0 \le t \le W$. A direct calculation using $W|z|^2 = 1$ gives

$$\left|\exp\left(Uz + \frac{\sqrt{W}}{4}(1 + \sqrt{W}z)^2\right)\right| = \exp((U + t)v). \tag{3.17}$$

At the center, in contrast, the exponential factor is $\exp(\sqrt{W}/4)$ and hence

$$|G(0)| = \frac{\exp(\sqrt{W}/4)}{\varphi(q)} \ge W^{-A}\exp(\sqrt{W}/4). \tag{3.18}$$

We give a bound polynomial in W for G on the boundary, separately according to the sign of $v - D$.

If $v \le D$, then $(U + t)v \le (U + W)D \le 2$. The F contribution to $|G|$ is consequently at most $2e^2$. For S_U, positivity and (3.16) give

$$|S_U(z)| \le e^{U(D-v)} \sum_n a(n)n^{-1-D} \le 2Ue^{U(D-v)}.$$

Multiplying by $\delta\exp((U + t)v)$ gives at most $2\delta Ue^{1+tv} \le 2e^2\delta U$. This proves the boundary estimate when $v \le D$, including the points with $v \le 0$ where the untruncated series might not be defined.

For $v > D > 0$, subtract the finite sum from (3.15). This is now legitimate, because both series converge absolutely. The resulting tail

$$F(z) - zS_U(z) = z \sum_{\log n \ge U} a(n)n^{-1-z}$$

has two types of terms. For the portion $\log n > U + W$, the absolute-value sum before multiplication by z is at most $2U\exp(-(U + W)(v - D))$ by (3.16). After multiplication by the factor δ and the exponential in (3.17), its contribution has modulus at most

$$2\delta U\exp((U + W)D - (W - t)v) \le 2e^2\delta U.$$

The remaining indices have $U \le \log n \le U + W$. By the supposition and the support of Λ, their nonzero terms occur only for nontrivial prime powers. The absolute series

$$C = \sum_p \sum_{j\ge 2} (\log p)\, p^{-3j/4}$$

is finite: for each p, the inner geometric sum is at most a fixed constant times $p^{-3/2}$, and $\sum_{n\ge 2}(\log n)n^{-3/2}$ converges. Therefore for $v > D$ the absolute-value sum over this remaining portion is at most $C\exp(-U(1/4 + v))$; indeed $n^{-1-v} = n^{-3/4}\exp(-(1/4 + v)\log n)$. Multiply by $\delta\exp((U + t)v)$ to bound its contribution by

$$C\delta\exp(-U/4 + tv) \le C\delta\exp(-W/4 + \sqrt{W}).$$

This tends to zero uniformly in the asserted range.

Combining the two cases gives, for all boundary points,

$$|G(z)| \le 2e^2 + 2e^2\delta U + C\delta\exp(-W/4 + \sqrt{W}).$$

This is $O_A(1 + W^A)$ uniformly in k, U, and even $O_A(W^{A-1/2} + 1)$ since $\delta = W^{-1/2}$. The function G is holomorphic on a neighborhood of the closed disk and so the maximum modulus principle bounds $|G(0)|$ by the same quantity. But (3.18) eventually exceeds every one of these polynomial bounds. This contradiction proves the theorem. □

This is the specialized theorem `exists_prime_modEq_one_log_mem` of `Lib/Analytic/Primes/AuxiliaryPrimes.lean`. Notice the two key features of the argument. First, the identity (3.15) controls the sum only on the right of the origin, but on the left it is enough to compare a *finite* truncated sum to its positive bound at D. Second, the smoothing exponential has exactly the boundary norm $\exp((U+t)v)$ with $0 \le t \le W$, allowing opposite tail comparisons at U and $U+W$. Its value at the center is superpolynomial. The absence of primes in the interval is incompatible with the pole residue. The prime powers within the interval do not obstruct this conclusion, because their sum is absolutely convergent already at $3/4$.

Corollary 3.11 (The primes for the admissible data, without Siegel–Walfisz)**.** *For fixed N_0, d, ε and the scales of Proposition 4.2, all the requirements (4.5), the selection of g_0, g_1 and the subsequent selection of N follow from Theorem 3.10.*

Proof. Write $Y = x^{\varepsilon}$ as there and take $W = \sqrt{Y}$. The modulus to which the prime classes are referred is

$$q = M_\omega^2 = (N_0 d)^{2(s+6)}.$$

Its base $M = N_0 d$ is a multiple of 4, fixed while x tends to infinity. Its exponent $k = 2(s+6) \ge 2$. Since $s \le 32 \log x$ and $2^s = m^{22} \le (2\sqrt{x})^{22}$, we have $q \le W^A$ and $2^s \le W^A$ for one fixed sufficiently large $A = A(N_0, d, \varepsilon)$. For every $U \in [100Y, 2^s]$ we have $W \le U \le W^A$. The theorem therefore gives a prime congruent to 1 modulo M_ω^2 with logarithm in $[U, U+W]$, exactly as in (4.5).

To select g_0, g_1, use $h_K + 1$ disjoint intervals of width W inside $[10Y, 11Y]$, for instance those starting at $10Y + 2iW$ for $1 \le i \le h_K + 1$. Here h_K depends on the fixed field only, so the intervals are disjoint and lie in the desired range for all sufficiently large x. Their starting points again lie between W and W^A after increasing A if necessary. Thus there are distinct split primes p_i in these intervals. For each choose a prime ideal $\mathfrak{p}_i$ of K of norm p_i. Pigeonhole in the finite ideal class group gives distinct i, j with $[\mathfrak{p}_i] = [\mathfrak{p}_j]$. Consequently $\mathfrak{p}_i\overline{\mathfrak{p}_j}$ is principal: in the class group $[\overline{\mathfrak{p}_j}] = [\mathfrak{p}_j]^{-1}$ since their product is (p_j). Its integral generator π_g has absolute norm $p_i p_j$, hence $\pi_g \overline{\pi}_g = g_0 g_1$ with $g_0 = p_i, g_1 = p_j$ in the notation of the source.

Finally the level exponent λ_N lies in $[100Y, 2^s]$ for sufficiently large x, by the estimate following (4.8). Therefore (4.5) gives N in the required interval regardless of the particular second cutoff eventually chosen in its allowed window. The prime search applies uniformly to the entire permitted range for λ_N, not just to a single choice made before the adjustment of x_1. □

In particular, when reading the proof of Proposition 4.2, Corollary 3.11 can be used in place of its invocation of a classical uniform prime number theorem. There is no circularity: the prime-window theorem concerned only the restricted powers of fixed $N_0 d$ and relied on nonvanishing at the line $\Re s = 1$, proved in Proposition 3.5. The level allocation and class-group pigeonhole are still as in the arithmetic argument, and no assumption about the desired critical line is involved.

3.9 Final conditional theorem and quantifier audit

The conventional analytic deductions are now self-contained relative to elementary real and complex analysis, finite abelian groups, elementary algebraic number theory (including finiteness of the ideal class group in the construction of the record) and the stated arithmetic inputs in the later chapters. To make the remaining mathematical burden precise, we state a theorem recording their exact logical role. The notation is as in Chapters 4 and 27.

Theorem 3.12 (Sufficient arithmetic inequalities)**.** *Suppose that for every fixed N_0 divisible by 18480, every prime $d \equiv N_0/2-1 \bmod N_0$, and every $0<\varepsilon<1/100$, the following holds for all sufficiently large x. For the admissible records of Proposition 4.2, including each permitted replacement of the second cutoff, the arithmetic estimates A, B, C of Chapter 27 hold with the same objects and normalizations. Then every Dirichlet L-function has no zero in $\Re s > 1/2$ (the principal pole convention at $s=1$ excepted), all its zeros in the open critical strip lie on $\Re s = 1/2$, and every zero of $\zeta(s)$ other than its standard trivial zeros lies on that line.*

Proof. Fix the three parameters in the hypothesis. Set $Y = x^{\varepsilon}$, m to be the least power of two at least $\sqrt{x}$, and $T = 10^6 \log x$. By the admissibility construction (with Corollary 3.11) there is such a record for all sufficiently large x. Use the ratio $X = |\mathcal{T}|/|\mathcal{S}| \le Y \le e^T$. Holomorphy of the polar transforms at and to the right of 1 follows from Proposition 3.5 and the exact computation of Proposition 2.1; note the scalar constant γ is specifically the coefficient computed in (3.8), so the logarithmic pole subtraction is exactly the expected one. Lemma 2.10 gives a permitted $u \in [T, T^2]$ with

$$R_{N_0}(e^u) - XR_1(e^u) \le e^{-T}.$$

Use the record with $x_1 = e^u$ and all derived levels, weights, spaces, lifts, and forms constructed from that record. Thus the hypotheses A–C compare identical data; no uniform pointwise bound in u has been substituted for the window argument.

The dimension and index estimates give q independent short forms among the translates, where $(1-4/b^2)n \le q \le n$, $n = \dim W > 0$. Their logarithmic squared covolume is bounded above by $q(H - m\mathcal{R} + L)$. The index cost is at most $5n\log(1+b^2)B$. In the lower estimate C it appears twice, while the leading height is subtracted three times by the allowance L. Because $H - m\mathcal{R} + L$ might have either sign, the relation $q \le n$ alone does not show that its multiple by q is at most its multiple by n. Use instead the crucial crude bound $-(H - m\mathcal{R} + L) \le b^2 x^{99}$ and $n - q \le 4n/b^2$. It gives

$$n(H - 3L) \le n(H - m\mathcal{R} + L) + 4nx^{99} + 10n\log(1+b^2)B.$$

After division by n the common H cancels. Since $B = faY^3$ and $L = BY$, divide the remaining error $4L + 4x^{99} + 10\log(1+b^2)B$ by BY^2. The quotient tends to zero: the linear terms in Y are $O(Y^{-1})$ after division, and $x^{99}/(BY^2) \to 0$ by $a \ge xb$ and $b = e^Y$. Hence for sufficiently large x, $m\mathcal{R} \le BY^2$.

Write the exact remainder identity in Chapter 27 in the form

$$R_{N_0}(x) = \frac{\mathcal{R}}{a|\mathcal{S}|} + R_{N_0}(x_1) - XR_1(x_1).$$

The two-cutoff term is bounded by x^{-10^6} by our choice of x_1. For the first term we have

$$\frac{\mathcal{R}}{a|\mathcal{S}|} \le \frac{fY^5}{m|\mathcal{S}|} \le \varphi(N_0)Y^6/m \le \varphi(N_0)x^{-1/2+6\varepsilon}.$$

The fixed factor $\varphi(N_0)$ and the much smaller window error are absorbed by $x^{2\varepsilon}$ for all sufficiently large x. Thus

$$R_{N_0}(x) \le x^{-1/2+8\varepsilon} \tag{3.19}$$

eventually. Every threshold in this deduction may depend on the fixed parameters N_0, d, ε.

To pass to an arbitrary character χ modulo q_*, suppose that $\widetilde{L}(\rho,\chi) = 0$ with $\Re\rho > 1/2$ and take $N_0 = 18480q_*$. The induced character modulo N_0 has a zero at the same point, by (3.2), since the additional Euler factors do not vanish in $\Re s > 0$. The class $N_0/2 - 1$ is a

unit modulo N_0, as its square is congruent to one modulo N_0 (use $4 \mid N_0$), and is 7 modulo 8 because $16 \mid N_0$. Choose $\varepsilon > 0$ with $\varepsilon < 1/100$ and $\beta = 1/2 + 8\varepsilon < \Re\rho$. By Proposition 2.7, whose classical dependencies have been proved here, the bad quadratic twists modulo N_0 with zeros in $\Re s \geq \beta$, $|\Im s| \leq |\Im\rho|$ have a convergent reciprocal sum. Proposition 3.6 supplies a prime d in the indicated class outside this set for all the characters modulo N_0 simultaneously. The hypothesis of this theorem and (3.19) apply to just that fixed d. Thus the residue at ρ of the meromorphic prime-error transform is nonzero, since there is a zero of one of the original characters and no zero of any of the twisted functions there. At every real point to the right of β its residue is nonnegative: in the interval $(\beta, 1)$ this is the sign of the positive integer multiplicity sum divided by the positive number $\varphi(N_0)\sigma(1-\sigma)$; beyond it the function is holomorphic by Proposition 3.5. Lemma 2.2, applied to (3.19), contradicts that nonzero residue at ρ. This proves the right-half-plane conclusion.

The symmetry of the completed primitive L-functions established above, followed by (3.2) for the imprimitive functions, proves the open-strip statement. Formula (3.12), nonvanishing on the line at one, and the explicit trivial-zero classification following that formula yield the full conclusion for ζ. □

We emphasize what the theorem does *not* prove by itself. The assumptions A, B, C contain the highly specialized arithmetic geometry and representation-theoretic estimates of the formal development. The conditional theorem removes none of their quantitative or compatibility requirements. It only replaces classical analytic black boxes and verifies that their exact conclusions, when proved, imply the stated result with no further analytic assumptions about zeros. In particular neither a uniform choice of x over the twisting primes, nor a simultaneous choice of quadratic twist avoiding zeros at arbitrarily large heights, is required. Such stronger claims are not used. The choice of a prime twist occurs after the hypothetical offending zero has been fixed; that choice is followed by the eventual limit in x for the chosen field.

The dependency order in the source is `exists_data`, the arithmetic conclusions from `New/Endgame/ShortAndEnough.lean` and `LongForms.lean`, `star_inequality`, `eventually_primeError_le_rpow`, and the final zero-free consequences in `New/Main.lean`. The analytic arguments of the present chapter and Chapter 2 supply their relevant analytic interfaces. Their proofs are independent of the remaining specialized input theorems in the arithmetic parts.

Chapter 4

The admissible data and coefficient spaces

Chapter 3 gives the analytic prime-window result used in this construction. Chapter 6 constructs the global Weil data, including the splitting character, the finite lattice and the signed Hermitian spaces, and establishes the integral unitary counts and mass-ratio estimate.

4.1 The CM extension and the different

We first record the precise arithmetic of the extension in the compact development. Fix $N_0 > 0$ with $18480 \mid N_0$, and a prime $d \equiv N_0/2 - 1 \pmod{N_0}$. Equivalently $N_0 \mid 2(d+1)$ and $N_0 \nmid d+1$. In particular $d \equiv 7 \pmod 8$, $d > 64$, and $(d, N_0) = 1$. Put $K = \mathbb{Q}(\sqrt{-d})$, of discriminant $-d$. For an integer $s \geq 1$ define

$$c_s = \lfloor \log_3 s \rfloor + 5, \quad T_s = 3^{c_s} 2^s, \quad N_1 = N_0 T_s, \quad E = K(\zeta_{N_1}), \quad M_\omega = (N_0 d)^{s+6}.$$

Here T_s is a number, not a set of places, and $T_s = |\mathcal{T}|$ below. For each prime $p \mid N_1$ the character $\chi_K(p) = 1$. Indeed for an odd such prime p, $p \mid (d+1)/4$ since all the primes of N_1 divide N_0; hence $-d \equiv 1 \pmod p$. For $p = 2$ the claim follows from $d \equiv 7 \pmod 8$. Consequently every nonsplit prime of K is at least 13, and d is prime to N_1. Since K and $\mathbb{Q}(\zeta_{N_1})$ are linearly disjoint (their discriminants have disjoint prime supports),

$$E \subset \mathbb{Q}(\zeta_{N_1 d}), \qquad [E : K] = \varphi(N_1) = T_s \varphi(N_0), \qquad \mathrm{Gal}(E/K) \simeq (\mathbb{Z}/N_1)^\times.$$

In the inclusion we use the quadratic Gauss sum of conductor d. The field E is totally imaginary and complex conjugation is central in its Galois group over $\mathbb{Q}$, so E is a CM field with real subfield E^+. The extension E/E^+ is unramified at the finite places: at every residue characteristic at least one of the primitive third and fourth roots of unity in E has order prime to that characteristic, and complex conjugation acts nontrivially on its reduction. Thus complex conjugation cannot belong to an inertia group of $E/\mathbb{Q}$. The roots of unity in E are exactly the N_1-th roots: the cyclotomic field of even order $N_1 d$ has no other roots of unity, and if E contained a primitive d-th root it would have degree over $\mathbb{Q}(\zeta_{N_1})$ at least $d - 1 > 2$. Also $N_1 d \mid M_\omega$, by inspecting the powers of 2, 3 and the remaining primes.

Let Φ be the set of archimedean places, each represented by the embedding extending our fixed embedding of K. We identify Φ with $\mathrm{Gal}(E/K)$ using one chosen embedding of E. Write $[j] \in \Phi$ for the residue class of j modulo N_1 if it is invertible, and for 1 otherwise (this latter convention occurs in the formal prime sums but contributes nothing when $\chi_K(p) = 1$).

Put $\mathcal{T} = \{\tau \in \Phi : \tau \equiv 1 \pmod{N_0}\}$. Then

$$|\mathcal{T}| = T_s, \qquad f = |\Phi| = \varphi(N_0)T_s.$$

For primes $p \nmid N_1 d$ the Frobenius in $\mathrm{Gal}(E/\mathbb{Q})$ has components $(\chi_K(p)p, \chi_K(p))$ in the identification $\Phi \times \{\pm 1\}$ in which the second factor is complex conjugation. This follows because the actual automorphism sends ζ_{N_1} to its p-th power and acts on K by $\chi_K(p)$. In particular a prime of E above p has residual degree the order of this pair. For inert p in K, the sign correction $-p^j$ at odd iterates is the origin of the place permutations in the local prime sums (Chapter 2). For these assertions see `New/Data/Different.lean` and `New/Data/WeilDataOfE.lean`.

The coprime discriminants imply $\mathcal{O}_E = \mathcal{O}_K[\zeta_{N_1}]$: the integral tensor product is maximal after localizing at each prime, since at least one factor is unramified there. Thus the relative different over K is generated by $P'(\zeta)$ for $P(X) = \Phi_{N_1}(X)$, the cyclotomic polynomial, and the absolute different is generated by $\sqrt{-d}\,P'(\zeta)$. The standard cyclotomic integer/different facts are recalled in [5]. The reciprocity relation $P(X) = X^D P(1/X)$, with $D = \deg P$ even, gives

$$P'(\zeta^{-1}) = -\zeta^{2-D}P'(\zeta).$$

Consequently multiplication by $i\zeta^{1-D/2}$ makes $P'(\zeta)$ real. We choose a purely imaginary generator ϖ of the absolute different by multiplying the resulting real element by $\sqrt{-d}$. All conjugates of the root-of-unity multiplier $i\zeta^{1-D/2}$ have modulus 1. In particular

$$|\tau(\varpi)| = \sqrt{d}|P'(\tau\zeta)|, \qquad \mathrm{rd}_E^f = \prod_{\tau\in\Phi} |\tau(\varpi)|, \tag{4.1}$$

where $\mathrm{rd}_E = |\operatorname{disc} E|^{1/[E:\mathbb{Q}]}$. The second equality follows by taking the square root of the absolute norm of the different. The identity

$$\Phi_{N_1}(X) = \Phi_{N_0}(X^{T_s})$$

follows either by listing roots or by successively applying $\Phi_{pN}(X) = \Phi_N(X^p)$ for $p \mid N$. Differentiating at ζ and using $|\tau\zeta| = 1$ shows that every $|P'(\tau\zeta)|$ is T_s times one of the fixed positive numbers $|\Phi'_{N_0}(\zeta_0)|$ for primitive N_0-th roots ζ_0. Hence $\mathrm{rd}_E \asymp_{N_0,d} T_s$. More specifically along the data below, for all sufficiently large x,

$$x^{10} \le |\tau(\varpi)|, \mathrm{rd}_E \le x^{12}. \tag{4.2}$$

This is the numerical content of `hasLogBound_varpi` and `hasLogBound_rootDisc` in `New/Seesaw/IndexTypes.lean`. The ensuing Weil data also choose a splitting character and compatible finite ideles; we construct these choices in Chapter 6, since different choices of theta normalizations can change the period and volume identities used below.

4.2 The scales, and the reason for so many parameters

The principal scale is x; the goal is to estimate $R_{N_0}(x)$. The proof uses another cutoff x_1 and a large order scale b, with

$$(\log b)^{100} \le x, \qquad x^{10^6} \le x_1, \qquad x_1^{10^6} \le b, \qquad (\log x_1)^6 \le \log b.$$

These inequalities are intentionally very far apart. They allow polynomial losses in the ranks or field degree to be negligible compared with the principal weights. For each fixed $0 < \varepsilon < 1/100$, the construction starts with $\log b = x^\varepsilon$ and $\log x_1 = x^{\varepsilon/6}$. The second cutoff will later be lowered within an admissible window. In the compact edition it is a real scale; no rounding of x is required.

Parameter	Meaning	Size or relation
m	Number of lifts	$\sqrt{x} \le m \le 2\sqrt{x}$
r	Rank of initial target group	$x^5 \le r \le 3x^5$
k	Rank of initial source group	$k = r + 1 - m$
t	Growth of row ranks per stage	$t = 5r$
h_2	Half-gap in the weights	$h_2 = r^8$
N_1	Cyclotomic order	$N_0 \mid N_1$
$\lvert\mathcal{T}\rvert$	Number of tops	$\asymp x^{11}$ up to log factors
$\lvert\mathcal{S}\rvert$	Number of thick places	$\ge \lvert\mathcal{T}\rvert / \log b$
a	Order of distinguished derivatives	$xb \le a < xb + N_1$
$g = g_0 g_1$	Product of two level primes	$b^{20} \le g \le b^{22}$
N	Extra level prime for source	$m \log b \le \log N \le 30m \log b$

The exact powers 10^6, 100 and 99 which occur in the source are generous margins; one cannot replace them without checking the many subsidiary inequalities, but they are not the source of the exponent $1/2$ in GRH. The significant power is the square-root length of the chain. The choice of m as a power of 2 and r as a power of 3 is compatible with the displayed intervals and serves the discrete combinatorial constructions in the proof.

The effective weights at the infinite places come in two sizes:

$$w_{\mathrm{nom}}(\sigma) = \begin{cases} a/x, & \sigma \in \mathcal{S}, \\ a/x_1, & \sigma \notin \mathcal{S}. \end{cases}$$

The integral weights actually used are rounded upwards to multiples of N_1 and shifted: if $w^-(\sigma)$ is the rounding of $w_{\mathrm{nom}}(\sigma)$ then $w_j(\sigma) = w^-(\sigma) + 2jh_2 + h_2$. Hence neighboring weights have a large gap. This separation permits quantitative archimedean and local estimates while the naive weights still record the two prime cutoffs x and x_1. It is important not to confuse the large weights with the logarithmic scale $\log b$.

4.3 Admissibility and construction of the data

It is useful to give the full existence proof rather than appeal to the entirety of the formal parameter-search module. The following is the portion of the definition of `Data` in `New/Data/AdmissibleData.lean` that is needed to begin the construction. Write $Y = \log b$, $U_1 = \log x_1$. There is one fixed positive constant c_* depending only on N_0, d (the source's `Parameters.c`, which includes the uniform CM generator constant), and the scales satisfy

$$10^{12} c_* \le \log Y, \quad 100 \log Y \le \log x, \quad 10^6 \log x \le U_1, \quad 10^6 U_1 \le Y, \quad 6 \log U_1 \le \log Y. \tag{4.3}$$

In addition to the rank intervals in the preceding table, require $m = 2^u$ with $u \ge 1$, $s = 22u$, $r = 3^v$ with $v \ge 1$, and $s \le 32 \log x$. The thick set lies in the coset -1 modulo $N_0 3^{c_s}$ and has the following property. If $\phi : \Phi \to (\mathbb{Z}/2^{s+1})^\times$ is reduction and $0 < i < j \le 324s$, then

$$\sigma, \tau, \rho \in \mathcal{S}, \quad \phi(\tau/\sigma)^j = \phi(\rho/\sigma)^i \implies \phi(\tau/\sigma) = 1. \tag{4.4}$$

The two distinct primes g_i lie in $[b^{10}, b^{11}]$ and are 1 modulo M_ω^2; there is $\pi_g \in \mathcal{O}_K$ with $\pi_g \bar{\pi}_g = g_0 g_1$. Finally the data include prime existence in the windows

$$\forall U \in [100Y, 2^s] \quad \exists p \text{ prime}, \quad p \equiv 1 \pmod{M_\omega^2}, \quad U \le \log p \le U + \sqrt{Y}. \tag{4.5}$$

The source defines the entire record before constructing any automorphic forms. All our subsequent spaces depend on this common record.

Lemma 4.1 (Behrend construction in a projected cyclic group)**.** *Let G and H be abelian groups, $\phi : G \to H$ a homomorphism, $u, z \in G$, and suppose the order of $\phi(u)$ is D. Let L_0, h, Q be positive integers with $(L_0Q)^h \le D$ and $L_0 \ge 2$, $Q \ge 2$. There is a subset S of $z\langle u\rangle$ of size at least $Q^h/(hQ^2)$ such that, for $s_0, s_1, s_2 \in S$ and $0 < i < j \le L_0$, the relation $\phi(s_1/s_0)^j = \phi(s_2/s_0)^i$ implies $s_1 = s_0$.*

Proof. Consider the Q^h integer vectors $v = (v_0, \dots, v_{h-1})$ with $0 \le v_a < Q$. The squared Euclidean lengths have at most $h(Q-1)^2 + 1 \le hQ^2$ values; choose a sphere containing at least the asserted number of vectors. Encode a vector by $e(v) = \sum_{a=0}^{h-1} v_a (L_0Q)^a$ and take $s(v) = zu^{e(v)}$. The map $v \mapsto \phi(s(v))$ is injective since $0 \le e(v) < D$ and the encoding is injective. The assumed multiplicative relation becomes

$$je(v_1) \equiv ie(v_2) + (j-i)e(v_0) \pmod{D}.$$

Both sides are integers with digits at most $L_0(Q-1) < L_0Q$ in base L_0Q, hence lie in $[0, (L_0Q)^h) \subset [0, D)$. The equality therefore holds digit by digit and $v_1 = (i/j)v_2 + (1 - i/j)v_0$. On a Euclidean sphere a strictly interior convex combination has strictly smaller norm unless the two endpoints are equal. Thus $v_0 = v_1 = v_2$. □

This is the combinatorial ingredient of `exists_three_term_free` in `New/Data/DataExist.lean`; it is a simultaneous Behrend-type construction (compare [4]).

Proposition 4.2 (Existence and flexibility of admissible data)**.** *For fixed N_0, d, $0 < \varepsilon < 1/100$, and every sufficiently large x, one can construct admissible data $\mathcal{D}$ satisfying*

$$Y = x^{\varepsilon}, \qquad U_1 = x^{\varepsilon/6} \ge (10^6 \log x)^2, \qquad |\mathcal{S}| \ge \frac{T_s}{Y}. \tag{4.6}$$

The same record is admissible if U_1 is replaced by any $u \in [10^6 \log x, (10^6 \log x)^2]$, interpreting every derived quantity with the new U_1.

Proof. Take m and r to be the least powers of 2 and 3 which are at least $\sqrt{x}$ and x^5, respectively, and choose s as above. Then $2^s = m^{22}$, $s \le 32 \log x$ eventually, and

$$81s < 3^{c_s} \le 243s, \qquad T_s = 3^{c_s} m^{22} \asymp x^{11} \log x.$$

In particular $x^{11} \le T_s \le x^{11} Y^{1/5}$ for sufficiently large x. The choice of Y, U_1 in (4.6) satisfies (4.3): the two relevant upper comparisons hold as $x^{\varepsilon/6} = o(x^{\varepsilon})$ and $6 \log U_1 = \log Y$; the remaining lower bounds follow from $\varepsilon < 1/100$. This proves the size comparisons used in (4.2).

To choose $\mathcal{S}$, apply Lemma 4.1 with $z = [-1]$, $u = [1 + M']$, $M' = N_0 3^{c_s}$ and $L_0 = 324s$. Put $v_0 = v_2(N_0) \ge 4$. Since $v_2(M') = v_0$ and the order of $1 + M'$ modulo 2^{s+1} is $D = 2^{s+1-v_0}$ when $s + 1 > v_0$, every power lies in the required coset after multiplication by z. The order formula follows inductively from $v_2((1+M')^{2^h} - 1) = v_0 + h$, by factoring successive differences of squares. Fix an integer $h > 44/\varepsilon$ and choose $Q = \lfloor D^{1/h}/L_0 \rfloor$, which is at least 2 for large x. Then

$$|\mathcal{S}| \ge \frac{Q^{h-2}}{h} \gg_{\varepsilon, N_0} \frac{2^{s(1-2/h)}}{(324s)^{h-2}}.$$

For the desired lower bound, compare with $T_s/Y \asymp s 2^s x^{-\varepsilon}$. Since $2^s \asymp x^{11}$, the ratio of the displayed lower bound to T_s/Y is bounded below by a positive constant times $x^{\varepsilon - 22/h}/(\log x)^{h-1}$, which tends to infinity. The sparsity property is immediate from the lemma.

It remains to find the primes. The restricted power-modulus prime-window theorem, proved in Corollary 3.11, applies to $Q_0 = M_\omega^2$ and to the indicated windows, uniformly in

U. It establishes (4.5), and also the existence of distinct primes $p_i \equiv 1 \pmod{M_\omega^2}$ in disjoint windows in $[10Y, 11Y]$, with $i = 1, \dots, h_K + 1$ where h_K denotes the class number of K. For example use the windows $[10Y + 2i\sqrt{Y}, 10Y + (2i+1)\sqrt{Y}]$. All p_i split in K, since $p_i \equiv 1 \pmod{d}$. Choose a prime ideal $\mathfrak{p}_i$ of K above each. By pigeonhole there exist distinct i, j with $[\mathfrak{p}_i] = [\mathfrak{p}_j]$ in the ideal class group. Then $\mathfrak{p}_i\bar{\mathfrak{p}}_j$ is principal; an integral generator π_g satisfies $\pi_g\bar{\pi}_g = p_ip_j$. Put $g_0 = p_i, g_1 = p_j$. This completes the construction of the record.

Lastly, any replacement value of U_1 in the asserted interval is positive, no larger than $x^{\varepsilon/6}$ and at least $10^6 \log x$. It preserves all inequalities in (4.3), while the other fields of the record do not depend on U_1. This proves flexibility. In the code this is used by `eventually_prime Error_le_rpow` to replace the scale in the structure before constructing any derived weights or levels. □

This supplies a conventional proof of `exists_data` using the restricted prime-window lemma `exists_prime_modEq_one_log_mem` in `Lib/Analytic/Primes/AuxiliaryPrimes.lean`, reconstructed in Chapter 3. Neither a uniform prime number theorem nor a pointwise prime number theorem at the second cutoff is required.

4.4 Derived weights, level and integral structures

In the actual integral weights write $a = N_1\lceil xb/N_1\rceil$, $a_1 = N_1\lceil (a/x)/N_1\rceil$ and $a_2 = N_1\lceil (a/x_1)/N_1\rceil$ (denoted b_1, b_2 in the formal source and in the local chapters). Thus $w_0(\sigma) - h_2$ is a_1 or a_2, with $w_j(\sigma) = w_0(\sigma) + 2jh_2$. For $0 \le j \le 2m$ we have

$$w_{\mathrm{nom}}(\sigma) \le w_j(\sigma) \le w_{\mathrm{nom}}(\sigma) + x^{42}, \quad w_j(\sigma) \le 2b, \quad \log w_j(\sigma) \le \log w_{\mathrm{nom}}(\sigma) + 1/x. \tag{4.7}$$

For the first upper bound use $N_1 + (4m+1)r^8 = O(x^{81/2}) + O(x^{11}\log x) \le x^{42}$ eventually. Next $w_{\mathrm{nom}} \le a/x \le b + N_1/x$, and $b \gg x^A$ for every fixed A. Finally $w_{\mathrm{nom}} \ge a/x_1 \ge xb/x_1 \gg x^{43}$ by the scale assumptions, and $\log(1 + x^{42}/w_{\mathrm{nom}}) \le x^{42}/w_{\mathrm{nom}} \le 1/x$. These translate `w_le_nominalWeight_add`, `w_le_two_mul_b` and `log_w_le` in `New/Data/Weights.lean`.

Write $w_* = f^{-1}\sum_\sigma \log w_{\mathrm{nom}}(\sigma)$ and $G_* = \log(g\,\mathrm{rd}_E/(2\pi))$, where $g = g_0g_1$. The extra prime N is selected as in `Data.N`, using the logarithmic target

$$\lambda_N = m(w_* + G_*) - G_* - \tfrac{15}{2}\log b, \quad \lambda_N \le \log N \le \lambda_N + \sqrt{\log b}, \quad N \equiv 1 \pmod{M_\omega^2}. \tag{4.8}$$

The large ranges give $w_* = (1 + o(1))Y$ and $G_* \in [(20 - o(1))Y, (22 + o(1))Y]$, uniformly after replacement of U_1. Hence eventually $mY \le \lambda_N \le 29mY$, $100Y \le mY$ and $29mY \le 2^s = m^{22}$; this uses $Y = x^\varepsilon$. We may therefore use (4.5), giving $mY \le \log N \le 30mY$. Since $4 \mid N_1 \mid M_\omega$, Fermat's two-square theorem furnishes $\pi_N = u + iv \in \mathcal{O}_E$ with $\pi_N\bar{\pi}_N = N$. Every prime above N has residual degree one, and the support of (π_N) contains exactly one member of each conjugate pair (because N is split and unramified and the exponents sum to 1). Likewise the chosen places above g_0, g_1 are the prime divisors of $\bar{\pi}_g$; they are conjugate-disjoint and there are f above each. These observations justify the finite groups and the determinant normalizations used later.

4.5 Two kinds of automorphic forms

If V is a nondegenerate Hermitian space over E with respect to the involution of E/E^+, $U(V)$ denotes its unitary group over the totally real field E^+. Automorphic forms on such a group are functions (with the specified transformation law at infinity) on an adelic quotient. For a totally definite Hermitian space, the group is compact at infinity, and spaces of algebraic

modular forms at a fixed level can be described by values in an algebraic representation on finitely many double cosets. In particular their coefficient integrality and Hermitian pairings can be studied by actual lattice calculations. For the indefinite groups, one instead uses cusp forms of prescribed discrete-series type at infinity. The source imposes suitable finite level and controls their behavior in the cusp.

The proof uses both types. Here are the main spaces; indices on Hermitian spaces should not be confused with indices of vectors.

- The source Hermitian spaces X_j are definite of rank $k + j(t+1)$, $0 \le j \le m$. Let C denote the space of algebraic modular forms on $U(X_0)$ of weight $(a, 0, \dots, 0)$ at tops and zero at the other infinite places, with the prescribed level at g and N. Let $C^{\mathrm{ram}} \subset C$ be the subspace of forms with vanishing averages under the prescribed local scalings at the places of N. These are the "ramified columns."
- The row spaces are forms on Y_j of rank $r + jt$. The starting Y_0 is indefinite, and Y_m is definite. Let W be the selected subspace of forms on $U(Y_0)$, at its prescribed archimedean type and at level g. (We use the letter W only for this vector space; the source also uses W_j for auxiliary rank-r Hermitian spaces.) The spaces selected by a finite-group projector are called good forms.
- For the determinant computation there are also single-lift auxiliary pairs of Hermitian spaces W_j, V_j of ranks $r, r + t$. These are used to replace the norm calculation for a whole chain by the calculations associated to individual lifts and their comparisons.

At the chosen finite places v above g, the principal-level form spaces carry an action of the finite group

$$\mathcal{G} = \prod_v \mathrm{GL}_r(\mathbb{F}_v).$$

The good projector e_0 is the sum of specified idempotents in $\mathbb{C}[\mathcal{G}]$. The good labels are defined through squarefree polynomials over the residue fields, with exclusions at roots of unity and primitive-polynomial conditions. This is more than a harmless way to cut down the dimension: it is designed to guarantee local representation types favorable both to the filling argument and to the Galois-theoretic analysis of congruences. The precise conditions are discussed in the subsequent sections. We write $n = \dim W$ and $c = \dim C$; neither of these is a Hermitian rank.

4.6 Precise weights and levels of the source and target

For a weight at a fixed infinite place we use column multiplicities $\kappa_\tau(h)$, meaning the number of Young-diagram columns of height h. Write $h_0 = 2h_2$ for the weight gap to distinguish it from the dummy height h, and put $n_j^X = k + j(t+1)$. The compact source uses the same multiplicity function at every stage, restricted to heights strictly below the rank of the stage:

$$\kappa_\tau(h) = \begin{cases} a\mathbf{1}_{\tau\in\mathcal{T}}, & h = 1, \\ w_1(\tau) - \lfloor (k+t+1)/2 \rfloor, & h = k, \\ h_0 - (t+1), & h > k,\ h \equiv k \pmod{t+1}, \\ 0, & \text{otherwise.} \end{cases}$$

The determinant exponents are $e_0(\tau) = 0$ and $e_j(\tau) = \lfloor n_{j-1}^X/2 \rfloor - w_j(\tau)$ for $j \ge 1$. These are `Data.` followed by the Greek letter κ, and `Data.lastEntry` in `New/Curve/LiveSum.lean`

(the source uses Greek letters in the identifiers). At X_0 the only column heights below k are the height-one columns supported on the tops, and this recovers the source weight described above. At the endpoint, the branching to Y_m retains

$$u(h) = \begin{cases} h_2 - \lfloor (t+1)/2 \rfloor, & h \geq k,\ h \equiv k \pmod{t+1}, \\ 0, & \text{otherwise}, \end{cases}$$

columns of their former height; its new multiplicity is $u(h) + \kappa_\tau(h+1) - u(h+1)$. Its determinant normalization comes from `restrictToY` in `New/Columns/ColumnDefs.lean`. This weight convention matters: it is not permissible to use the ordinary weight w_{nom} in the representation instead of the rounded weights.

Put $F = E^+$. For a standard definite unitary group of rank q', the coefficient space $V_{\kappa,e}$ is the product-place tensor of the complex Weyl modules with multiplicities $\kappa_\tau(h)$, tensored with the determinant of exponent $e(\tau)$. The basis in the code is its integral Weyl-coordinate basis, and its Gram form has an integral, positive-definite matrix G_κ; the determinant twist is unitary at infinity. The action of $\gamma \in U(q')(F)$ uses $\tau(\overline{\gamma})$ in the algebraic weight representation. For an open compact $K' \subset U(q')(\mathbb{A}_{F,f})$, the algebraic forms of fixed level are functions v on the finite adelic group with values in $V_{\kappa,e}$ satisfying

$$v(\gamma g k') = \rho_{\kappa,e}(\gamma) v(g), \qquad \gamma \in U(q')(F),\ k' \in K'.$$

There are finitely many adelic classes. Choosing a representative g of each class gives the direct sum of invariant coefficient spaces under the groups $\Gamma_g = U(q')(F) \cap gK'g^{-1}$. This follows at once from the equivariance condition. It is the model `AlgForm` of `Lib/GLn/SchurModule/WeightRepBasic.lean`. The class-sum Hermitian form is

$$\langle v, w \rangle_{K'} = \sum_{[g]} \frac{v(g)^* G_\kappa w(g)}{|\Gamma_g|}.$$

When the level K' has Haar volume one, this class sum is the integral of the invariant scalar pairing over the finite adelic quotient. For the correspondence with highest-weight matrix coefficients of the automorphic form at infinity, the source uses the factor $\dim V_{\kappa,e}$ from Schur orthogonality and the factor $[K_0 : K']$ from its finite measure convention. The precise comparison is `integral_mul_eq_classIntegral`, `Space.algFormEquiv` in `Lib/Adelic/CuspForms/AlgebraicForms.lean`.

Let K_g denote the level which is principal at the chosen places above g and maximal integral elsewhere. For X_j the column level is K_g intersected with the condition that the last row of the matrix at places dividing π_N is $(0, \ldots, 0, 1)$ modulo that prime. For each $w \mid \pi_N$, a scalar finite adelic unit z_w represents a generator of $\mathbb{F}_N^\times$ in the compatible standard split model. The ramified part of C is exactly

$$C^{\mathrm{ram}} = \mathrm{im} \prod_{w \mid \pi_N} (1 - \rho(z_w)).$$

The scalar operators commute and are of finite order on this level, so this is equivalently the intersection of the kernels of their cyclic averaging projections. It follows that for each fixed w, every ramified column has the form $(1 - \rho(z_w))v$; this elementary observation is used in the ramified kernel integrals. The codimension bound is established in Chapter 14 via the precise local level and class count. Source: `New/Columns/ChainValuation.lean`, `Data.ramifiedColumns` and `exists_eq_sub_centralGen`.

For the row representation, number the basis of Y_0 by $i = 0, \ldots, r-1$, using the numbering compatible with the Bessel flag in the source. At τ the (undoubled) Harish-Chandra parameter

is

$$\varsigma_\tau(i) = \begin{cases} k/2 - 1 - i, & i+1 < k, \\ h_2 - w_{i+2-k}(\tau), & i+1 \geq k. \end{cases}$$

Here k is even. The Harish-Chandra coordinates of all the row spaces, together with the signs specifying their real form, are chosen so that the parameter gap between coordinates of opposite signs is at least $2h_2$ in the doubled coordinate convention of the code. In particular the rank-r auxiliary space W_j is specified by the first $k+j-1$ positive and remaining $m-j$ negative coordinates, while Y_j has rank $r+jt$ and $\lfloor (m-j)/2 \rfloor$ negative signs at each place. On Y_j the doubled new blocks have coordinates $t-1-2w_l(\tau)-2b'$ $(0 \leq b' < t$, $1 \leq l \leq j)$; all doubled coordinates, including the base coordinates, are reduced by one if j is odd. The starting Y_0 is indefinite and the endpoint Y_m definite. The form of the row spaces and their adelic frames is part of the construction: the code gives `Data.exists_spaceY` in `New/Trace/GoodIdempotent.lean`, using the explicit glued Gram matrix and an integral frame. We regard the existence of this frame with its specified congruences as a formal-source input, and retain the flags and identifications of that construction (see Chapter 10).

The row form space before the good projection consists of automorphic functions of the indicated discrete-series type, right K_g-fixed. More precisely the source uses the subspace fixed by convolution with the reproducing kernel of the prescribed real type and demands rapid decay; its finite-dimensionality is the assertion `finite_levelSpace` for the cusp kernel. The fixed-level convolution representation of K_0/K_g at the split chosen places factors through

$$\mathcal{G} = \prod_{i=0,1} \prod_{w \mid \bar{\pi}_g,\ w \mid g_i} \mathrm{GL}_r(\mathbb{F}_{g_i}).$$

The good projector e_0 is a central idempotent of its group algebra, and $W = e_0 V$ with V the row form space at the principal level. The conditions and dimension estimates for e_0 are handled in the spanning chapter. We will write $C^{\mathrm{ram}}_{\mathrm{int}}$ for the forms in C^{ram} integral at every embedding and use $n = \dim_{\mathbb{C}} W$, $c = \dim_{\mathbb{C}} C$ throughout.

4.7 Integral gauges and the precise uniform index

For each rational prime p and field isomorphism $\iota : \mathbb{C} \xrightarrow{\sim} \overline{\mathbb{Q}}_p$, let $\mathscr{O}_\iota$ be the valuation ring pulled back by ι. If v is an algebraic form, the gauge of its value $v(g)$, with respect to the lattice at g, is

$$\|v(g)\|_{\iota,g} = \left\| \rho_{\kappa,e}\left((g^{-1})_{\iota\tau \circ \mathrm{conj}} \right)_{\tau \in \Phi} \iota(v(g)) \right\|_{\sup,p}.$$

The component matrix is obtained from the finite adele by extension to $\overline{\mathbb{Q}}_p$ via $\iota\tau \circ \mathrm{conj}$; coefficients of the Weyl matrix are integers, and the sup norm is taken in its integral coordinate lattice. This is `weylGauge` in `Lib/GLn/SchurModule/WeightRepBasic.lean`. Integral means every such gauge is at most one. Because the representation action respects the transported lattice, one may test on the finite class representatives. Notice that integrality is not defined merely by integrality of the complex coordinates at $g = 1$. This is significant when local Hecke modifications change the lattice or the branching map transports a different lattice at N.

For a family $y = (y_i)_{i=1}^{q}$ of row forms, the index inequality $\mathrm{IndexLe}(y, d_{\mathrm{ind}})$ uses real numbers e_p such that $\sum_{p \in S} e_p \leq d_{\mathrm{ind}}$ for every finite set S of primes. For each p, ι and every allowed m-fold twist θ of conductor prime to p, there must be a $q \times q$ matrix U over $\mathscr{O}_\iota$ satisfying

$$|\iota(\det U)|_p \geq e^{-e_p}, \qquad \Theta_\theta\Big(\sum_j U_{ij} y_j \Big) \text{ integral at } \iota \quad (1 \leq i \leq q).$$

The matrix is not required to be the same for different twists, but the budgets e_p are fixed before them. This is precisely `Data.IndexLe` in `New/Endgame/ShortAndEnough.lean` (there the field isomorphisms are supplied as instances). These conventions explain the factor $2d_{\mathrm{ind}}$ in the squared-determinant estimate. There is a stronger local assertion on column lattices, before applying the projected Poincaré map: the corresponding submodule is independent of the twist. The short-form argument uses this stronger assertion to choose change-of-family matrices with a global index bound.

4.8 The theta chains and their relation

Let θ be a weight-zero automorphic character of $U(1)$, with tame conductor at a set of places disjoint from its complex conjugate and prime to gN. The successive powers through θ^m must retain the required ramification at that conductor (the trivial character is also allowed). A theta lift is a linear transformation defined by pairing an automorphic form on one group with a theta kernel associated to a Weil representation on a dual pair. Schematic notation for a lift is

$$(\mathcal{L}F)(h) = \int_{[U(V)]} F(u)\,\overline{\vartheta(u,h)}\,du,$$

where the precise conjugations, kernels, measures and characters depend on the convention. The source fixes these choices by its definitions of the spaces and lifts; this schematic integral is a description of the idea, not a replacement for those definitions. For definite groups, such integrals reduce to sums over finitely many classes.

Denote by Θ_θ the composite lift along the row chain beginning in W. The column chain begins in C, runs along the definite X_j, and afterwards restricts from X_m to Y_m. The ranks have been chosen so that

$$\dim_E X_m = k + m(t+1) = r + mt + 1 = \dim_E Y_m + 1.$$

The one-dimensional complement is used in the branching and level normalization. The m extra rank increments on the two sides make the iterated seesaw identity possible. Informally the relationship is

$$\begin{array}{ccccc} C^{\mathrm{ram}} & \longrightarrow & \text{definite chain} & \longrightarrow & X_m\text{-forms} \\ \downarrow P & & & & \downarrow \text{restriction} \\ \text{initial row forms} & \longrightarrow & \text{row chain} & \longrightarrow & Y_m\text{-forms.} \end{array}$$

Here P is a Poincaré series constructed from a source form by Bessel induction. The actual identity involves a finite combination of definite chains, nonzero normalization factors and finite-place restrictions, not a literal equality between arbitrary individual theta kernels. The projected translates P_s in the upper bound, for $s \in \mathcal{G}$, are obtained from the Poincaré map by translating and projecting; we include the fixed normalizing factors from `projectedPoincare` in this notation.

The two rows of the diagram serve different purposes. The definite chain gives access to coefficient lattices and the sizes and prime divisibilities of derivatives of theta functions. The indefinite row chain gives a large automorphic space, a spanning theorem for Poincaré series, and a determinant calculation. Neither side alone yields the main inequality.

4.9 Norms, squared covolume and index

For each field automorphism σ of $\mathbb{C}$, the target definite space has a conjugated coefficient pairing, with squared norm denoted $\|\cdot\|_\sigma^2$. We use the pairing obtained by conjugating

the coefficient data as specified in `Lib/Adelic/CuspForms/LevelForms.lean`, not the absolute value of an arbitrary conjugate of a real analytic integral. For independent vectors $y_1, \dots, y_q \in W$ define

$$\mathsf{G}_\sigma(y) = \left(\langle \Theta_1 y_i, \Theta_1 y_j \rangle_\sigma\right)_{i,j}, \qquad \operatorname{covol}_\sigma^2(y) = |\det \mathsf{G}_\sigma(y)|.$$

The determinant may be zero for some family before an injectivity statement is available; the upper estimate does not presuppose it is positive. We call V a logarithmic upper bound if $\operatorname{covol}_\sigma^2(y) \le e^V$ for all σ. By the Gram determinant inequality,

$$\operatorname{covol}_\sigma^2(y) \le \prod_{i=1}^{q} \|\Theta_1 y_i\|_\sigma^2.$$

Thus bounds on individual lifts imply a bound on the determinant of their span. The source's definition uses the notation `covolumeSq`.

The paired finite-place structure and the index bound were fixed in Section 4.7.

4.10 The common leading term and the prime remainder

Here we introduce the quantities that occur in all three theorems. Let $\lambda_{d,1}(\nu) = \Lambda(\nu)(1-\chi_K(\nu))$ be the weight without a congruence restriction, and let s_ν be the place represented by $-\nu$ modulo N_1 when ν is invertible there (with the source's convention of the identity place otherwise). Define the finite prime sum at $\sigma \in \Phi$ by

$$A(y, \sigma) = \sum_{1 \le \nu \le a} \lambda_{d,1}(\nu) \left(\frac{y}{\nu} - w_{\mathrm{nom}}(s_\nu \sigma) \right)_+ .$$

Write $\kappa_K = \gamma - \Re(L'(1,\chi_K)/L(1,\chi_K))$ for the Chowla–Selberg constant used in the proof, and

$$\Omega = \frac{2\pi}{\mathrm{rd}_E} e^{\kappa_K}, \qquad J(j) = \log(j!) - j\log(\Omega/g), \qquad \overline{w} = \frac{1}{f}\sum_{\sigma \in \Phi} \log w_{\mathrm{nom}}(\sigma).$$

The quantity $J(j)$ is the archimedean cost of a derivative of order j in the normalization of the proof. Its dependence on the period of the CM curve is the key to matching archimedean terms with the weighted prime sum. An expected main term for $A(a,\sigma)$, averaged appropriately, is

$$A_0 = a(\log a - 1 - \kappa_K - \overline{w}).$$

The exact quantities in the inequalities are

$$\begin{aligned} \mathcal{R} &= \sum_{\sigma \in \mathcal{T}} \left(A(a,\sigma) - A_0\right), \\ \mathcal{U} &= \sum_{\sigma \in \Phi} \left(J(w_0(\sigma)) - A(w_{\mathrm{nom}}(\sigma), \sigma)\right), \\ H &= mk\mathcal{U}, \qquad B = fa(\log b)^3, \qquad L = B\log b = fa(\log b)^4. \end{aligned}$$

These definitions are `Data.primeSum`, `Data.jetSize`, `Data.dirCost`, `Data.mainHeight`, `Data.mainAsk`, and `Data.remainder`. The quantity $\mathcal{U}$ represents the net cost of an ordinary differentiation direction, obtained by subtracting the finite-place divisibility gain from the archimedean derivative size. Both the short-form argument and the determinant argument are calibrated to this quantity. It is not necessary to determine its precise sign or asymptotic value.

The prime remainder has the following exact interpretation, which will be important in the final deduction:

$$\mathcal{R} = a\{|\mathcal{S}|(R_{N_0}(x) - R_{N_0}(x_1)) + |\mathcal{T}|R_1(x_1)\}. \tag{4.9}$$

This is `Data.remainder_eq`. We explain its mechanism in Chapter 27: the two weights replace the finite prime sums by R_{N_0} at two scales and a remainder R_1 without a congruence condition. The role of the analytic choice of x_1 is to control the unwanted second-cutoff expression. Notice that the normalized error $R_M(z)$ of the preceding section is not the same as the unnormalized, large remainder $\mathcal{R}$.

Chapter 5

Adelic coefficient models and integral linear algebra

The proof uses automorphic forms on unitary groups in two rather different ways. An indefinite group supplies a spectral and integration problem; a definite group supplies a finite class set equipped with a coefficient lattice. The passage between the two is subtle. This chapter explains the elementary foundations of the definite model, the local integrality tests and the exact meaning of the norms used later. It also establishes several facts about these models directly. It does not assume a general comparison between an automorphic integral lattice and a Galois-representation lattice. The latter is a particularly substantial issue in Chapters 20 and 23.

The reader familiar with algebraic modular forms may skip most of this chapter, but should retain the normalization of the coefficient pairing in (5.6), the freeness criterion 5.2, and the distinction between the Weyl lattice and the lattice of polynomials in Lemma 5.7. These are significant for the quantitative statements of the proof. The finite-class-number framework is classical (see [6, 7]); the arguments below are given for the particular totally definite Hermitian spaces in use here.

5.1 Hermitian spaces, places and adelic lattices

Let E/F be a CM extension with $F = E^+$, equipped with conjugation $z \mapsto \bar{z}$. Its set of real places is denoted Φ after choosing one complex embedding of E above each real place. Let V be an E-vector space of dimension l with nondegenerate Hermitian form h; we use conjugate linearity in the first argument and linearity in the second. In a basis the form is $h(u, v) = u^* J v$ for a Hermitian matrix J. Define

$$U(V)(R) = \{A \in \mathrm{GL}_l(E \otimes_F R) : A^* J A = J\}$$

for an F-algebra R, with conjugation extended R-linearly. A totally definite space is positive definite at every complex embedding (changing the sign at a real place makes no difference to the compactness statement). Then $U(V)(F \otimes_{\mathbb{Q}} \mathbb{R}) \simeq \prod_{\tau \in \Phi} U(l)$ is compact. In the standard totally definite model of the proof $J = I$ and the lattice is $\Lambda_0 = \mathcal{O}_E^l$. Write $G = U(V)$.

For a finite place v of F, if v splits in E as $w\bar{w}$, then $E \otimes_F F_v \simeq F_v \times F_v$, with conjugation interchanging the two factors. A unitary matrix is determined by its matrix in the w factor; the other component is forced by the equation defining G. In the standard model this is

$$A \longmapsto (A, A^{-T}), \qquad A \in \mathrm{GL}_l(F_v). \tag{5.1}$$

At an inert finite place the group is genuinely unitary, not a copy of the general linear group; the residue field and Frobenius actions differ. The proof repeatedly uses the splitting at the selected primes g_0, g_1, N but cannot use such a local model at nonsplit primes in the jet argument. The precise choice of a member of each split pair matters for coefficient valuations: using the conjugate of the place changes the matrix A to A^{-T}. This is why the source records the finite-place embeddings of each CM-type factor explicitly.

The finite adelic group $G(\mathbb{A}_{F,f})$ is the restricted product of $G(F_v)$ with respect to the local stabilizers of $\Lambda_0 \otimes \mathcal{O}_{F_v}$; all but finitely many components of an element are integral invertible. Put $K_0 = \prod_v \mathrm{Stab}(\Lambda_{0,v})$, a compact open subgroup. A choice of finite adelic g defines a fractional $\mathcal{O}_E$-lattice

$$\Lambda_g = V \cap (g\widehat{\Lambda}_0), \qquad \widehat{\Lambda}_0 = \prod_w (\Lambda_0 \otimes_{\mathcal{O}_E} \mathcal{O}_{E,w}),$$

where V embeds diagonally in $V(\mathbb{A}_{E,f})$. The meaning of the intersection is that $x \in V$ satisfies $x \in g_w\Lambda_{0,w}$ at all finite places. Conversely a family of local lattices L_w equal to $\Lambda_{0,w}$ outside a finite set determines a lattice with these completions. One proof is as follows. Choose a nonzero ordinary integer c such that $c\Lambda_{0,w} \subseteq L_w \subseteq c^{-1}\Lambda_{0,w}$ at every w. The desired lattice is the inverse image in $c^{-1}\Lambda_0$ of the product of the corresponding subgroups in $c^{-1}\Lambda_0/c\Lambda_0$, using the Chinese remainder decomposition of this finite module. By construction its completions have exactly the desired images and contain $c\Lambda_{0,w}$, proving equality. In particular two lattices with the same completions coincide.

Proposition 5.1 (Adelic class interpretation)**.** *In the standard totally definite model the double quotient*

$$X_{K_0} = G(F)\backslash G(\mathbb{A}_{F,f})/K_0$$

is naturally the set of isometry classes of Hermitian lattices in V locally isometric to Λ_0 at all finite places. For every open subgroup $K' \subseteq K_0$, the class set $X_{K'}$ is finite and every stabilizer $\Gamma_g = G(F) \cap gK'g^{-1}$ is finite.

Proof. The assignment $g \mapsto \Lambda_g$ is surjective onto the indicated genus: choose a local isometry at each exceptional place and the identity where the lattices are equal. If $\gamma\Lambda_g = \Lambda_{g'}$, then at each completion $(g')^{-1}\gamma g$ preserves the standard local lattice, so it belongs to K_0. This proves the class interpretation. For stabilizers, the underlying positive trace form

$$Q(u, v) = \mathrm{Tr}_{E/\mathbb{Q}}\, h(u, v)$$

on a lattice Λ_g is a positive-definite symmetric integral form. Its integral isometries are finite: the images of any fixed basis under an isometry must lie among the finitely many lattice vectors of the prescribed lengths. Hence Γ_g is finite, even at level K_0.

Here is a proof of finiteness of X_{K_0}. The rank $D = l[E : \mathbb{Q}]$ and determinant of the trace form on Λ_g are fixed, the latter because all of its local completions are isometric to the fixed lattice. We first use the elementary consequence of Minkowski's second theorem ([8]) that there are only finitely many integral positive-definite lattices of fixed rank and fixed determinant up to isometry. Indeed, for the Euclidean length determined by Q, every nonzero vector has norm at least one. If $\lambda_1 \le \cdots \le \lambda_D$ are successive minima, Minkowski's bound on their product at fixed covolume gives an upper bound on every λ_i. We may choose D independent lattice vectors of uniformly bounded length. Their Gram matrix is an integral matrix with uniformly bounded entries; the index of their span in the original lattice is uniformly bounded by Hadamard's inequality for their parallelepiped divided by the fixed covolume. A lattice containing this span with index I is intermediate between that span and its scalar multiple

by $1/I$; for bounded I there are only finitely many such lattices. This proves the claimed finiteness of positive integral trace forms.

For each representative of the trace lattice, consider the possible $\mathcal{O}_E$-actions. Multiplication by each element u of a fixed integral basis of $\mathcal{O}_E$ has operator norm at most $\max_{\tau\in\Phi}|\tau(u)|$ on the trace form: the real orthogonal decomposition into CM factors shows this directly. Therefore there are only finitely many integral matrices for these actions on a fixed trace lattice. An isometry preserving the trace form and these actions preserves h, since the values of $\mathrm{Tr}_{E/\mathbb{Q}}(ah(x,y))$ for $a\in E$ determine $h(x,y)$ by nondegeneracy of the trace pairing. Indeed these values are $Q(x,ay)$; thus the Hermitian form is recovered from the trace form and the E-action. This reduces Hermitian lattice classes to a finite collection. Finally $[K_0:K']<\infty$ by compactness and openness, so each class at level K_0 splits into finitely many classes at level K', as asserted. □

This proof gives the qualitative finiteness assertion used in the algebraic modular form theory. It is not a quantitative mass estimate for varying fields and ranks; the counting inequalities in the main arithmetic construction use further results. In the code the applicable abstract setup is `HasFiniteType` in `Lib/Adelic/Hecke/ClassMean.lean`. The argument above suffices to justify it in the standard definite case, with maximal level and all its open subgroups.

5.2 Why the extra level removes stabilizer denominators

Although every stabilizer in Proposition 5.1 is finite, its order could divide a small prime and spoil integrality of a class sum. The selected primes are far larger than required to force triviality. The following trace argument gives the exact bound used in the compact source.

Lemma 5.2 (Large residue norm and free level)**.** *Let G be totally definite of rank l over E/F. Suppose $K'\subseteq K_0$ is a subgroup reducing to identity at distinct places $w_1,\dots,w_a$ of E. If*

$$(2l)^{[E:\mathbb{Q}]}<\prod_{i=1}^{a}\#k_{w_i}, \tag{5.2}$$

then $\Gamma_g=\{1\}$ for every $g\in G(\mathbb{A}_{F,f})$.

Proof. If $\gamma\in\Gamma_g$, its characteristic polynomial is monic with coefficients integral at every finite place, because it is locally conjugate to an integral matrix. In particular $z=l-\operatorname{tr}\gamma\in\mathcal{O}_E$. At every archimedean embedding all eigenvalues of γ have modulus one, so $|\tau(z)|\le 2l$. At the chosen finite places, conjugacy to an integral matrix reducing to the identity gives $z\in\mathfrak{p}_{w_i}$. If $\gamma\neq 1$, the trace cannot equal l: in a positive-definite factor, $\Re(\operatorname{tr}\gamma)=l$ forces every eigenvalue to be one, and a unitary matrix is diagonalizable. Thus $z\neq 0$. The absolute norm of a nonzero element in the product of the distinct prime ideals is at least the product of their norms. It is also at most $(2l)^{[E:\mathbb{Q}]}$ by the archimedean bound. This contradicts (5.2). □

In the proof's notation the chosen places occur over g_0,g_1, with f places over each prime, all of residue degree one. The right side of (5.2) is g^f, $g=g_0g_1$. As $[E:\mathbb{Q}]=2f$, it suffices that $(2l)^2<g$. For all the definite ranks used in the relevant class sums, $l\le 2(k+m(t+1))$, this follows from $g\ge b^{20}$ and the polynomial bounds in x for m,r,t. Therefore principal level at the chosen places is free. This proves the special case called `Data.classAut_eq_bot_of_le_chosenLevel` of `New/Data/ChosenPlaceLevels.lean`, which uses `classAut_eq_bot` in `Lib/Shimura/IntegralModel/NeatLevel.lean`. No general assertion about neat Shimura varieties is being imported here.

Corollary 5.3 (An integer clearing all stabilizer orders)**.** *Assume $l \le rt$ and that the large-residue condition holds for rank l. Put*

$$D_g = \prod_{w \text{ chosen}} |\mathrm{GL}_{rt}(k_w)|.$$

For every subgroup of maximal integral level in rank l and every adelic class, the order of its rational stabilizer divides D_g.

Proof. Conjugate the finite stabilizer by the representative g and reduce at the chosen places. Its kernel is a stabilizer at the principal level, hence trivial by the lemma. The reduction thus embeds the stabilizer in $\prod_w \mathrm{GL}_l(k_w)$. The latter group embeds by block diagonal matrices in $\prod_w \mathrm{GL}_{rt}(k_w)$, so Lagrange's theorem proves the assertion. □

This is the integer denoted `Data.chosenGroupOrder`. Notice that it is shared across every lower rank and every level in the specified range; using a different clearing integer for each class would be inadequate for the simultaneous determinant estimates. Its size can be bounded elementarily: $|\mathrm{GL}_L(k)| = \prod_{j=0}^{L-1}(|k|^L - |k|^j) < |k|^{L^2}$. Thus

$$\log D_g \le f(rt)^2 \log g.$$

In particular $\log D_g = O(fr^4 \log b)$ since $t = 5r$. With the scales $f = O(x^{11} \log x)$, $r = O(x^5)$, $\log b = x^{\varepsilon}$, this is $O(x^{31+\varepsilon} \log x)$ and hence is at most x^{32} eventually. This supplies a direct explanation of the apparently large polynomial bound appearing in the row budgets. It is a bound for the logarithm of the integer, not for the order itself.

5.3 Normalization of class sums and change of level

Let $G_f = G(\mathbb{A}_{F,f})$, $\Gamma = G(F)$, $K' \subseteq K_0$ open compact. For a left Γ-invariant and right K'-invariant function F taking values in a characteristic-zero vector space define

$$\mathcal{C}_{K'}(F) = \sum_{[g] \in \Gamma \backslash G_f / K'} \frac{F(g)}{|\Gamma_g|}, \qquad \Gamma_g = \Gamma \cap gK'g^{-1}.$$

This is a finite weighted sum, independent of the choice of representatives. The subscript does not mean that the measure of the whole adelic quotient is one. If Haar measure on G_f has $\mathrm{vol}(K_0) = 1$, then the quotient integral is $\mathrm{vol}(K')\mathcal{C}_{K'}(F) = \mathcal{C}_{K'}(F)/[K_0 : K']$. Indeed, the preimage of the class $[g]$ in G_f is the union of the translates $\gamma g K'$; such translates coincide precisely modulo Γ_g. For rational scalar expressions the class sum can be read in an arbitrary $\mathbb{Q}$-vector space without passing through real-valued integration. This is useful in comparing coefficient conjugates or applying p-adic valuations.

Lemma 5.4 (Level change)**.** *If $K'' \subset K'$ is open and F is right K'-invariant, then*

$$\mathcal{C}_{K''}(F) = [K' : K'']\mathcal{C}_{K'}(F).$$

Consequently $\mathcal{C}_{K'}(F)/[K_0 : K']$ is independent of the chosen level on which F is fixed.

Proof. Inside a class $[g]$ at level K' the finer classes are the orbits of the finite group $g^{-1}\Gamma_g g$ acting on the finite set K'/K'' on the left. Each finer stabilizer is conjugate to the stabilizer of the corresponding orbit representative. By the orbit-stabilizer formula, the sum over its orbits of the reciprocal stabilizer orders equals $[K' : K'']/|\Gamma_g|$. Multiply by the common value $F(g)$ and sum over the original classes. □

If the base level K_0 is fixed but a function is only invariant at some smaller level, its normalized integral is well-defined by the same lemma, using intersections of two levels. Translation by a finite adele need not preserve K' but preserves the normalized integral on smooth left-invariant functions, by Haar invariance; one may prove it algebraically by refining to $K' \cap aK'a^{-1}$. The unitary groups are unimodular, so left and right Haar measures agree. The compact source proves the needed class-sum identities in `Lib/Adelic/Hecke/ClassMean.lean`, notably `classSum_eq_relIndex_smul`. In a finite sum with coefficients in a valuation ring O, $D_g \mathcal{C}_{K'}(F)$ remains in O whenever F is O-valued and D_g is divisible by all stabilizer orders. At a free level D_g is unnecessary.

One should keep these three normalizations separate:

- the counting measure on a finite free class set, i.e. $\mathcal{C}_{K'}$;
- the Haar quotient measure with normalized maximal level (or a specified fixed level);
- the normalization of the real representation matrix coefficient by formal degree or coefficient dimension.

A ratio of theta periods computed with the first convention may involve extra factors of finite indexes when inserted into the second. The determinant height computation uses exactly these factors when moving between levels.

5.4 A useful elementary form of weak approximation

The proof requires a finite-place test of the coefficient lattice. It is not necessary to assume that the rational unitary group is dense in all finite adeles, which is in general false. Weak approximation at finitely many places is sufficient.

Lemma 5.5 (Weak approximation for the unitary group)**.** *Let J be any nondegenerate Hermitian form over the CM extension E/F. For a finite set of places of F the diagonal $U(J)(F)$ is dense in the product of the corresponding local unitary groups.*

Proof. Consider the F-linear space $\mathfrak{u}_J$ of matrices satisfying $X^*J + JX = 0$. On the open subset $\det(I + X) \neq 0$, the Cayley transform

$$X \longmapsto (I - X)(I + X)^{-1}$$

is unitary. This is verified by expanding $(I - X)^*J(I - X) = (I + X)^*J(I + X)$, and its inverse on the open subset $\det(I + A) \neq 0$ is $A \mapsto (I - A)(I + A)^{-1}$. The verification also holds after completion at each chosen place. To put a given finite tuple of local matrices A_v into these opens, multiply them by a common rational scalar from $U(1)(F)$. Indeed, choose any nonzero $\eta \in E$ with $\bar{\eta} = -\eta$; the elements $(1 + t\eta)/(1 - t\eta)$ for rational t are distinct elements of $U(1)(F)$, and for each fixed local A_v only finitely many scalar values make $I + cA_v$ singular. Choose c outside the finite union. For these local matrices the inverse Cayley coordinates are elements of $\mathfrak{u}_J(F_v)$. Approximate them simultaneously by a rational point of the finite-dimensional F-vector space $\mathfrak{u}_J$, using ordinary weak approximation in F. By continuity the rational Cayley transform approximates cA_v arbitrarily closely; multiplying by c^{-1} gives the required approximation. □

This gives the consequence of `Space.weak_approximation` in `Lib/Adelic/Landherr/WeakApprox.lean` used to pick p-integral representatives. In particular if g is a finite adele and p is rational, take a rational γ so that the components of γg at all places over p lie in the local

standard maximal integral groups. The required conditions are open and contain the tuple g_p^{-1} before translation. This makes it possible to compute integral coordinates of each class at a given p without selecting one representative which is integral at every prime.

5.5 Algebraic modular forms and the finite-dimensional model

Let (R, ρ) be a finite-dimensional complex representation of $G(F)$ obtained from an algebraic representation of the archimedean group. A smooth algebraic modular form of weight ρ is a function $\phi : G_f \to R$ such that $\phi(\gamma g) = \rho(\gamma)\phi(g)$ for $\gamma \in \Gamma$, with open right stabilizer. At a fixed level K', we require right invariance under K'. There is a minor variant used throughout the proof: one may require $\phi(gk) = \chi(k)\phi(g)$ for a finite-order character χ of K'. This is the usual character level, and the trivial character gives ordinary level forms. The representation $\rho(\gamma)$ in the source is the Weyl representation evaluated at the conjugate of the chosen embeddings, as in `Space.weightRep` of `Lib/GLn/SchurModule/WeightRepBasic.lean`.

Choose one representative g_i of each class. Evaluation gives an isomorphism between the space of forms and a direct sum of subspaces of R, one for each class. At a trivial-character level this subspace is

$$R^{\Gamma_{g_i}} = \{u : \rho(\gamma)u = u \text{ for all } \gamma \in \Gamma_{g_i}\}.$$

At a character level it is the subspace defined by $\rho(\gamma)u = \chi(g_i^{-1}\gamma g_i)u$. Indeed, extending the value u_i at the representative by equivariance gives $\phi(\gamma g_i k) = \chi(k)\rho(\gamma)u_i$; the asserted relation is precisely what makes this extension well-defined (compare two such expressions for the same element). Conversely all values of a form are obtained in this way. The finite class set and stabilizers therefore give a finite-dimensional space.

Corollary 5.6 (Dimension and free level)**.** *At a free level, the dimension of the space of algebraic modular forms of any fixed coefficient weight and finite-order character is* $|X_{K'}| \dim R$. *For two free levels* $K'' \subset K'$ *(with ordinary invariance) the ratio of the dimensions is* $[K' : K'']$. *More generally every level space has dimension at most* $|X_{K'}| \dim R$.

Proof. At a free level each of the indicated subspaces is all of R. The ratio of the class numbers is the index by Lemma 5.4 applied to the constant function 1 at free levels. The general bound follows by taking dimensions of the subspaces. □

In particular when a single scalar unit z at a prime of the column level generates a cyclic group of order $N-1$ modulo that level, the z-fixed part has dimension $c/(N-1)$ if the enlarged level remains free. This is the class-number reason behind the estimate for the kernel of $1 - z$ (see the ramifying operator in Chapter 10). For several commuting cyclic actions, iterating the operator $1 - z$ loses at most the sum of the dimensions of their individual kernels. Indeed the codimension of an intersection is bounded by the sum of codimensions, and finite-order commuting operators are simultaneously diagonalizable over $\mathbb{C}$. The ramification condition is thus a dimension-controlled way to eliminate the invariant vectors, not a qualitative cuspidality condition on the original finite group.

5.6 The integral Weyl model

We set out the coefficient construction because a polynomial representation over $\mathbb{C}$ does not by itself determine an integral lattice. Let l be a positive integer, and let $\kappa(h)$ for $1 \le h < l$ be the nonnegative integer number of columns of height h of a Young diagram. Write

$$\lambda_i = \sum_{h \ge i} \kappa(h), \quad 1 \le i \le l, \qquad D = \sum_i \lambda_i = \sum_h h\kappa(h), \qquad \lambda_l = 0.$$

Let $X = (X_{ij})$ be an $l \times l$ array of variables. For a subset I of rows of size h let $[I]_h(X)$ be the determinant of the submatrix with rows I (in increasing order) and the first h columns. Put

$$b_{\mathfrak{t}}(X) = \prod_{h=1}^{l-1} \prod_{j=1}^{\kappa(h)} [I_{h,j}]_h(X).$$

These products of minors are homogeneous of total degree D. Under right multiplication of X by an upper triangular matrix T they transform by the character $\prod_i T_{ii}^{\lambda_i}$, since its first h columns transform among themselves by the upper triangular leading block. Left row substitution by a matrix A sends such a product to an integral linear combination of such products by Cauchy–Binet. These two observations explain the representation and the row substitution used in the source.

The standard integral basis theorem for Schur modules gives a basis b_i, consisting of designated products of these minors, of the semi-invariant polynomial module. It exists over $\mathbb{Z}$ and commutes with extension of scalars; its rank is given by the Weyl dimension formula. This standard-basis assertion is the classical input from the theory of straightening laws (compare [9, 10]), in the precise form `PolyRep.exists_minorProduct_basis` in `Lib/GLn/SchurModule/MinorBasis.lean`. We use it here as a formal-source input for the exact chosen integral basis, not as a claim that an arbitrary chosen complex highest-weight basis has integral coordinates. In characteristic zero the resulting module, with the corresponding convention for the row action, is the irreducible polynomial representation of highest weight λ or its dual (the two are distinguished by whether one acts on polynomials or on their evaluations).

For each array X set $\mathrm{ev}_X = (b_i(X))_i$. The vectors ev_X span the coordinate space over any infinite field over which the b_i remain independent: a functional zero on all of them would give a polynomial zero on every array, hence the zero polynomial. Left substitution defines a matrix $S_\kappa(A)$ by

$$S_\kappa(A)\,\mathrm{ev}_X = \mathrm{ev}_{AX}\,. \tag{5.3}$$

It has polynomial entries with integral coefficients and satisfies the representation identity, by the spanning observation and by substitution. Over the integers, the full-rank column determinants are separated as a determinant power, and a rational weight is represented by $\det(A)^e S_\kappa(A)$ for $e \in \mathbb{Z}$. For several infinite places one takes the exterior tensor product of these spaces and the product of the determinant twists; the total degree is the sum of the separate degrees. These are exactly the matrix and weight models called `schur` and `weylRep`.

We now derive the metric in these coordinates rather than appeal to an unspecified invariant pairing. Define the bilinear kernel

$$\mathcal{K}(Y, X) = \prod_{h=1}^{l-1} \det(Y_h^T X_h)^{\kappa(h)},$$

where X_h, Y_h are the first h columns of the two arrays. Cauchy–Binet gives

$$\mathcal{K}(Y, X) = \sum_{\mathfrak{t}} b_{\mathfrak{t}}(Y) b_{\mathfrak{t}}(X) = \mathrm{ev}_Y^T\, G_\kappa\, \mathrm{ev}_X\,.$$

The sum is over all choices of the ordered row subsets in the separate column factors, including repetitions when a height occurs more than once. To see the properties of G_κ write $b_{\mathfrak{t}} = \sum_i c_{\mathfrak{t}i} b_i$ with integral coefficients. Then $G_\kappa = \sum_{\mathfrak{t}} c_{\mathfrak{t}} c_{\mathfrak{t}}^T$, an integral symmetric matrix. Because the products span the polynomial module, these vectors $c_{\mathfrak{t}}$ span in characteristic zero. Hence G_κ is positive definite over $\mathbb{R}$ and defines the Hermitian form $\langle u, v\rangle_\kappa = \bar{u}^T G_\kappa v$ over $\mathbb{C}$. In the case of several places, use the product of these kernels, giving the tensor-product matrix.

The two polynomials $\mathcal{K}(Y, AX)$ and $\mathcal{K}(A^T Y, X)$ coincide. Evaluating and using (5.3) yields the contravariance identity

$$G_\kappa S_\kappa(A) = S_\kappa(A^T)^T G_\kappa. \tag{5.4}$$

For a unitary A, $\bar{A}^T = A^{-1}$ and the representation identity therefore imply invariance of $\langle \cdot, \cdot \rangle_\kappa$. A determinant twist of any integer exponent has modulus one on a unitary element and preserves invariance. This reconstructs the integral Gram form used in `Lib/GLn/SchurModule/Gram.lean`, since symmetries and polynomial expansion uniquely determine the matrix of the kernel in the basis of minor products. At an adelic class the coefficient norm is not an arbitrary Euclidean norm in the selected complex basis: it is this precise Gram form.

5.7 Fischer pairing and the highest-weight hook constant

For polynomials of fixed degree in variables X_{ij} the bilinear Fischer pairing is

$$\mathsf{F}(f, g) = f(\partial) g(0) = \sum_\alpha \alpha! f_\alpha g_\alpha.$$

Its associated Hermitian pairing replaces f_α with $\bar{f}_\alpha$ and is positive definite. The row action by a unitary matrix preserves the Hermitian pairing: on a homogeneous degree-D component one can identify this form with $D!$ times the pairing induced from the D-fold tensor power of the standard Hilbert space, projecting to the symmetric tensors; the monomial basis has squared norms $\alpha!$. Equivalently the chain rule gives the bilinear adjoint of substitution by A as substitution by A^T. This remains true on the semi-invariant polynomial subspace because the subspace is stable under row substitutions.

Write H_κ for the matrix $\left(\mathsf{F}(b_i, b_j)\right)_{ij}$ and let $f_\lambda(X) = \prod_h \det(X_{[1,h],[1,h]})^{\kappa(h)}$. Let $h_\kappa = \mathsf{F}(f_\lambda, f_\lambda)$, a positive integer, the hook constant in the code. It is not in general the plain product of the factorials of the row lengths. Its relation with the Gram matrix is the one needed to turn polynomial differentiation into a statement about coordinate norms.

Lemma 5.7 (The two integral pairings)**.** *One has*

$$G_\kappa H_\kappa = H_\kappa G_\kappa = h_\kappa I. \tag{5.5}$$

In particular for a column vector u the polynomial $P_u = \sum_i (G_\kappa u)_i b_i$ satisfies

$$\mathsf{F}(P_u, b_j) = h_\kappa u_j, \qquad \mathsf{F}(P_u, P_v) = h_\kappa u^T G_\kappa v.$$

Proof. Let $S = S_\kappa(A)$ denote the row action on evaluation vectors. For substitution by A the action on the column vector of coefficients of a polynomial is S^T. The adjoint formula gives $S H_\kappa = H_\kappa S_\kappa(A^T)^T$. By (5.4), it follows that $H_\kappa G_\kappa$ commutes with every $S_\kappa(A)$. In characteristic zero the representation is irreducible, so Schur's lemma says that the product is scalar. (Equivalently, the irreducibility assertion can be replaced by the spanning orbit of the highest vector and the same weight calculation below.) To determine the scalar use $e = \mathrm{ev}_{I_l}$. The kernel $\mathcal{K}(I_l, X)$ equals $f_\lambda(X)$, so $P_e = f_\lambda$. A minor product of highest row weight λ is f_λ itself: the sum of occurrences of the first s rows is at most $\sum_h \kappa(h) \min(s, h)$, with equality for all s only if each column of height h has exactly the rows $1, \dots, h$. Consequently whenever b_i has that weight its evaluation at I_l is 1 and it equals f_λ; all other basis products have zero evaluation at I_l and, being of a different weight, have zero Fischer pairing with f_λ (orthogonality under the diagonal torus, or disjoint monomial weights). Thus $H_\kappa G_\kappa e = h_\kappa e$. As $e \neq 0$, the scalar is h_κ. Symmetry gives both product identities, and the remaining assertions follow by matrix multiplication. □

This proves the identities `fockPairing_toPoly_nablaBasis` and `weylGram_mul_fockPairing` of the same file. In particular inversion of the Gram matrix introduces at most a fixed integral hook denominator: $G_\kappa^{-1} = H_\kappa/h_\kappa$. The exact value often cancels from a ratio, but its size is needed if one works integrally.

For $D \geq 1$ the elementary polynomial expansion gives $h_\kappa \leq D^{3D}$. Indeed the sum of absolute values of the coefficients of f_λ is at most $\prod_h (h!)^{\kappa(h)} \leq D! \leq D^D$; every multi-index α of degree D satisfies $\alpha! \leq D!$. Therefore

$$h_\kappa = \sum_\alpha \alpha! |(f_\lambda)_\alpha|^2 \leq D! \left(\sum_\alpha |(f_\lambda)_\alpha| \right)^2 \leq D^{3D}.$$

At degree zero the constant is 1. For a collection of places the polynomial is a product in disjoint variables, and the same bound applies with total degree. This is `hookNumber_le` in `Lib/GLn/SchurModule/Hook.lean`.

5.8 Examples of the coefficient normalizations

A useful example is the diagram consisting of a columns of height one. Its polynomial basis is $b_\alpha = X_1^{\alpha_1} \cdots X_l^{\alpha_l}$ with $|\alpha| = a$. The kernel is $(Y^T X)^a$, so

$$G_\kappa = \operatorname{diag}\left(\frac{a!}{\alpha!}\right), \qquad H_\kappa = \operatorname{diag}(\alpha!), \qquad h_\kappa = a!.$$

Thus the polynomial lattice and the Weyl evaluation lattice differ: the map $u \mapsto P_u$ multiplies each coordinate by $a!/\alpha!$. Even in this simplest case, replacing the integral evaluation coordinates by coefficients of a polynomial changes the denominators at primes dividing factorials. This is precisely the diagram occurring on the tops at stage X_0. For $l = 2, a = 2$, the three diagonal entries in the order $(2,0), (1,1), (0,2)$ are $1, 2, 1$, and $H_\kappa = \operatorname{diag}(2,1,2)$. Over a field of characteristic different from two both give equivalent integral models; at 2 the distinction is real. The formal proof similarly tracks the gauges under Hecke operators using the actual evaluation lattice, not just a complex vector-space isomorphism to a symmetric power.

For a single column of height h, $b_I = [I]_h$, so $G_\kappa = I$ by Cauchy–Binet, whereas the Fischer pairing matrix is $h!I$, as a determinant expansion has $h!$ distinct monomials with coefficients ± 1. This is the exterior power example. With several columns of varied heights, products of minors satisfy straightening relations: they cannot all be used as independent coordinates. For example in three row variables the two-by-two minors of the first two columns satisfy the ordinary alternating relation

$$X_{11}[23]_2 - X_{21}[13]_2 + X_{31}[12]_2 = 0.$$

It follows by expanding all three determinants. The standard integral basis is needed to remove such relations without inverting integers; otherwise an estimate on an overcomplete tensor model would not imply the desired statement about the actual coefficient lattice.

5.9 The class-sum metric and the real matrix-coefficient model

Returning to the definite adelic space, use the Hermitian form G_κ at the coefficient module of an algebraic form, with the indicated determinant twist. The rational unitary action preserves

it. Thus for two forms of the same fixed level and character,

$$\langle \phi, \psi \rangle_{K'} = \sum_{[g]} \frac{\overline{\phi(g)}^{T} G_\kappa \psi(g)}{|\Gamma_g|} \tag{5.6}$$

is well-defined and positive definite. At every $\sigma \in \mathrm{Aut}(\mathbb{C})$ define instead the conjugate coefficient norm $\|\phi\|_\sigma^2$ by applying σ to the two coefficient vectors before taking the displayed Hermitian square. It is generally wrong to write $\|\phi\|_\sigma^2 = \sigma(\|\phi\|^2)$, since σ need not commute with conjugation. The mixed Gram formula and determinant inequality in Chapter 17 correct this issue.

We give the real comparison of the class-sum model with matrix coefficients. Let $K_\infty = \prod_{\tau \in \Phi} U(l)$ and let ρ_∞ be the irreducible unitary coefficient representation with a fixed highest vector v_+ of norm one. Let d_ρ be its dimension. With probability Haar measure on K_∞, averaging the rank-one projection P_{v_+} over the group gives

$$\int_{K_\infty} \rho_\infty(k) P_{v_+} \rho_\infty(k)^{-1} dk = d_\rho^{-1} I.$$

Indeed the average commutes with the irreducible representation and has trace one; this proves the formula by Schur's lemma. Consequently for any two vectors u, v the integral of the product of the corresponding matrix coefficients with v_+ is $d_\rho^{-1}\langle u, v\rangle$. This is the form of Schur orthogonality needed here.

For ϕ as above define on the adelic group

$$F_\phi(k, g) = \langle v_+, \rho_\infty(k^{-1})\phi(g)\rangle.$$

The left rational invariance follows by equivariance of ϕ and the identification of the algebraic weight with its real counterpart. Right finite-level transformation follows from that of ϕ. Applying Schur orthogonality at each class shows

$$\int_{G(F)\backslash G(\mathbb{A}_F)} \overline{F_\phi} F_\psi = \frac{1}{d_\rho [K_0 : K']} \langle \phi, \psi \rangle_{K'},$$

with the fixed maximal finite level and compact real measures normalized to one. Conversely, a scalar function whose real component lies in this matrix coefficient model determines the coefficient vector by integrating against $d_\rho \rho_\infty(k) v_+$, by the same projection identity. Thus highest-vector matrix coefficients and the coefficient-valued forms are mutually inverse in this model. These calculations explain the factors in `Space.integral_mul_eq_classIntegral` and `Space.algFormEquiv` of `Lib/Adelic/CuspForms/AlgebraicForms.lean`. They also explain why a formal-degree factor is needed when replacing an indefinite discrete-series model by a definite one; in the indefinite case the real group is not compact and the relevant normalization instead uses its formal degree and explicit reproducing coefficient (Chapter 14).

5.10 Transported lattices and gauges

Let p be a rational prime, and let ι be a field isomorphism from $\mathbb{C}$ to the algebraically closed p-adic coefficient field used in the source. In the notation below one can use the completed algebraic closure $\mathbb{C}_p$; the formal model is an algebraic closure of $\mathbb{Q}_p$ with its extended absolute value. Only coefficients and field embeddings, not a comparison of the complex and p-adic topologies, are involved. The valuation ring on the complex coefficient field is

$$O_\iota = \{z : |\iota(z)|_p \le 1\}.$$

The local matrix obtained from g^{-1} at the embedding $\iota\circ\tau\circ\mathrm{conj}$ acts in the weight representation on the corresponding factor. We denote the resulting product matrix by $T_\iota(g^{-1})$. If u is a coefficient vector, the actual gauge is

$$|u|_{\iota,g} = \|T_\iota(g^{-1})\iota(u)\|_{\sup}.$$

This expression is meaningful because the local factor extends the map $E \to \mathbb{C}_p$ continuously at its induced place. Notice that applying ι directly to infinitely many adelic components without choosing a place would be meaningless. The definition is `weylGauge` in `Lib/GLn/ScHurModule/WeightRepBasic.lean`.

Lemma 5.8 (Elementary invariances of the gauge)**.** *The following assertions hold for the standard definite model.*

1. *If $\gamma \in G(F)$, then $|\rho(\gamma)u|_{\iota,\gamma g} = |u|_{\iota,g}$.*
2. *If k is integral and invertible at all the places above p, then $|u|_{\iota,gk} = |u|_{\iota,g}$.*
3. *The gauge is ultrametric and homogeneous: $|u+v|_{\iota,g} \le \max(|u|_{\iota,g}, |v|_{\iota,g})$ and $|cu|_{\iota,g} = |\iota(c)|_p|u|_{\iota,g}$.*
4. *If g is integral invertible over p, then the gauge equals the coefficient sup norm. In particular to prove that an algebraic form is integral it is enough to test one p-integral representative of each class, using Lemma 5.5.*

Proof. The first follows from the representation identity and compatibility of the embeddings: $T_\iota((\gamma g)^{-1})\iota(\rho(\gamma)) = T_\iota(g^{-1})$. For the second, multiplication by $T_\iota(k^{-1})$ on the left preserves the sup norm. Indeed both the matrix and its inverse have integral entries, because the Weyl coefficients are integral polynomials and determinant powers of units are units. Each gives an inequality in one direction. The third is immediate from the linearity of the transported coordinate map. Finally the same observation applies to $T_\iota(g^{-1})$ when g itself is integral. Any value at an arbitrary representative of the class is related by the left rational covariance and the right level covariance, and finite-order scalars have absolute value one, giving the claimed test. □

These are the statements `weylGauge_embFin`, `weylGauge_mul`, `weylGauge_smul`, and `weylGauge_le_one_of_integralAbove` in `Lib/Adelic/CuspForms/ClassReps.lean`. The gauge is always attached to a specified embedding and often to a modified lattice at gd, with d a suitable adele. The latter is used in the treatment of the level prime N. It is not legitimate to use integrality at one representative of each class for every prime without considering that representative's transported lattice.

For a fixed level the integral forms at a fixed ι form a module over O_ι by the third assertion. At a free level this module is a direct sum of transported free coefficient lattices, so it is free. More generally the local constructions may intersect finitely generated modules or restrict to complex subspaces. A module over a valuation ring need not be finitely generated, and that is a reason for using finite-family determinant inequalities in the source rather than automatically speaking of a finite index. The following basic facts illustrate the distinction.

Lemma 5.9 (Finite matrices over a valuation ring)**.** *Let O be a valuation ring in a valued field with multiplicative valuation, and let M be a finite rectangular matrix over O. Then elementary row and column operations in the corresponding general linear groups over O put M in diagonal form with nonzero diagonal entries $a_1, \dots, a_j$ satisfying $a_i \mid a_{i+1}$, followed by zeros. In particular a full-rank finitely generated submodule $L \subseteq O^q$ admits a basis with logarithmic determinant cost $-\log|\det A|$ for a matrix A whose columns form the basis.*

Proof. If the matrix has a nonzero entry, choose one of maximal absolute value (equivalently least valuation), possible because there are finitely many entries. In a valuation ring divisibility is total; thus the pivot divides every entry. Move the pivot to the upper left by permutations and clear its column and row by elementary operations over O. All entries remaining are divisible by the pivot, so induction proves divisibility of successive pivots. The operations are invertible over O, and in the full-rank square-block output this gives the claimed basis and determinant. No uniformizer or discrete value group is needed. □

For a family $y = (y_i)_{i=1}^q$ of forms, integrality after taking combinations by a matrix $U \in M_q(O_\iota)$ says that the coefficient module of integral relations contains the rows of U. If $\det U$ is nonzero, this gives a finite determinant-cost witness even when the full integral relation module has no preferred finite presentation. Determinant costs add under composition: $-\log|\det UV| = -\log|\det U| - \log|\det V|$. The actual definition `Data.IndexLe` uses such witnesses with a common budget at each rational prime for every embedding and allowable twist. The matrix may depend on the embedding and twist; the bound must not. The stronger submodule witnesses used in the local Hecke computations are explained in Chapter 10.

5.11 Finite-place Hecke operators on explicit lattices

Hecke operators at the split places are central to the integral repairs. We recall the elementary double-coset model and its lattice action. Let G_0 be a locally compact group, K compact open, and M a representation over a commutative coefficient ring. For a K-bi-invariant function f of finite support on right K-cosets define on M^K

$$T_f(v) = \sum_{gK} f(g)\rho(g)v.$$

This is independent of the representatives since v is K-fixed. Left K-invariance of f implies the output is K-fixed as well, by reindexing. For two such functions f, f' the identity $T_{f'}T_f = T_{f'*f}$ follows by multiplying the two sums, where the convolution is

$$(f' * f)(h) = \sum_{gK} f'(g)f(g^{-1}h).$$

All equalities make sense over the coefficient ring; no analytic convergence and no division by the volume of a group is involved. Intertwiners of representations carry these sums to the same sums in the target. More generally, if $f(kg) = \chi(k)^{-1}f(g)$, the same reindexing shows that $\rho(k)T_f(v) = \chi(k)T_f(v)$; the inverse on the function is necessary for this convention of the operator. This is the model in `Lib/Adelic/Hecke/HeckeBasic.lean`, especially `heckeFun_eq_sum`, `heckeFun_heckeFun`, and `heckeFun_mem_levelSpace`.

Now take $G_0 = \mathrm{GL}_l(F_v)$ with maximal integral subgroup $K = \mathrm{GL}_l(\mathcal{O}_v)$, residue field k of order Q, and uniformizer ϖ_v. Let $d_j = \mathrm{diag}(\varpi_v I_j, I_{l-j})$. A right coset in Kd_jK/K is a lattice L with

$$\varpi_v\mathcal{O}_v^l \subseteq L \subseteq \mathcal{O}_v^l, \qquad \dim_k(\mathcal{O}_v^l/L) = j.$$

Indeed right multiplication does not change the lattice, and by the diagonal normal form every such lattice lies in the double coset. These lattices are in bijection with the subspaces of codimension j of k^l, or equivalently (after duality) the j-dimensional subspaces. Counting ordered independent j-tuples and dividing by the number of ordered bases in dimension j gives the Gaussian binomial coefficient

$$\#(Kd_jK/K) = \binom{l}{j}_Q = \prod_{a=0}^{j-1} \frac{Q^l - Q^a}{Q^j - Q^a}.$$

For $j = 1$ this is $(Q^l - 1)/(Q - 1)$; in particular in rank two the neighbor graph has $Q + 1$ outgoing edges at an unramified vertex, counted with multiplicity. This count is an exact finite identity, not a local asymptotic.

The split-unitary identification (5.1) transports these operators to unitary groups. In particular inverse transpose switches the role of a sublattice and a scaled superlattice. In the unitary model one must choose the correct factor before translating the gauge; this is the same sign convention that appears in the Frobenius-twisted depth bounds of Chapter 10. If T is a sum of translates whose coefficients are p-integral and all its input gauges at the translated lattices are bounded by A, the ultrametric inequality bounds the output gauge by A. There is no factor counting the number of cosets. By contrast the triangle inequality at an archimedean norm introduces such a factor; this difference is used throughout the proof.

A small but instructive example is the neighbor relation on local rank-two lattices: for the operator T_1 above one has in the unnormalized integral Hecke algebra

$$T_1^2 = T_{(2,0)} + (Q + 1)T_{(1,1)}.$$

Here $T_{(a,b)}$ denotes the indicator of the double coset of $\mathrm{diag}(\varpi_v^a, \varpi_v^b)$. Indeed a chain of two index-Q sublattices between the original lattice and a colength-two lattice has exactly one possible intermediate lattice when its quotient is cyclic of length two, and $Q + 1$ choices when the quotient is a two-dimensional residue vector space. Thus the coefficient of the scalar double coset is genuinely $Q + 1$; normalizing one operator by a power of ϖ_v or conjugating it by an integral weight changes its appearance. In the proof a considerably more involved system of commuting operators is normalized so that its local monomial matrix has entries like p^{depth}. A bound on the complex Hecke polynomial alone would not prove an integral elementary-divisor bound for this matrix.

5.12 Class correspondences and adjoints

The same integral-Hecke operator has a useful adjoint property for the normalized metric. Suppose G_f has a unimodular Haar measure and a representation on a scalar or vector-valued automorphic L^2 space is unitary under right translation. With base level K of volume one, the operator on the fixed space associated with f is $\int_{G_f} f(g)\rho(g)\,dg$, a finite sum when applied to fixed vectors. Its adjoint is associated with $f^\vee(g) = \overline{f(g^{-1})}$ by changing variables $g \mapsto g^{-1}$. This holds for different levels after inserting their volume factors or after restricting each side to the indicated fixed subspaces. In particular for a double-coset indicator the adjoint has the inverse double coset, using probability Haar on the level. If we instead work with the two class-sum metrics at different levels, the quotient-of-index factors in Lemma 5.4 must be retained.

More concretely let $L = K \cap aKa^{-1}$ for a finite adelic a. The two maps from classes at this common smaller level to classes at level K are $[g] \mapsto [g]$ and $[g] \mapsto [ga]$. Their fibers, with stabilizer weights, have relative sizes $[K : L]$ and $[K : a^{-1}La]$ respectively; these are equal by unimodularity. Pullback and finite weighted pushforward are adjoint for the corresponding normalized measure. This is a finite-class explanation of the correspondence description of Hecke operators. It extends to algebraic forms with the appropriate transports of coefficient vectors along the correspondence. In particular if a matrix-entry estimate on one side is established using the coefficient gauge, duality at the other side demands the actual contragredient coefficient structure, rather than an assumption that a chosen complex orthogonal projector has integral entries.

5.13 Where this elementary model stops

There are several boundaries to the argument just given. First, it provides finiteness but not the asymptotic estimates for the mass of the definite genus as the field and ranks grow. Second, the finite-level dimension formula does not identify the multiplicities of the different local representations on an indefinite group. Third, the Hecke double-coset description and the Gram matrix alone do not give the integral theta transfer or the small Selmer denominator estimates of the main argument. The first two issues reappear in the quantitative spanning argument (Chapter 14); the last takes up Chapters 10, 17, and 20. The elementary computations in this chapter isolate the exact places where their more advanced statements enter.

Chapter 6

The splitting character, integral unitary models and the definite mass

The same field, self-duality convention, and finite integral group occur on both sides of the seesaw, in the filling construction, and in the lower-bound periods. A small change to a local character or to the volume of a lattice cannot in general be tolerated by the final comparison. It is useful, therefore, to construct these data once and to distinguish their algebraic consequences from the subsequent analysis of the Weil representation. In this chapter we give the construction of the Weil data of the compact source, including the existence of its special unramified splitting character; the integral transitivity and group counts needed by the comparisons; and a derivation of the quantitative consecutive-mass estimate from the classical Hermitian mass formula. The principal classical global results retained in this chapter are the reciprocity norm theorem of class field theory, Landherr's local–global theorem and the mass formula. None involves a zero-free region of the strength sought in this book.

6.1 Notation and the local norm theorem

Write $E = K(\zeta)$, $K = \mathbb{Q}(\sqrt{-d})$, $F = E^{+}$, $L = N_1$, $f = [F : \mathbb{Q}]$, and $\zeta = \zeta_L$. We use the CM type Φ consisting of the embeddings of E which extend the chosen embedding $\sqrt{-d} \mapsto i\sqrt{d}$. For clarity, v in this chapter is a finite place of F, whereas w is one of the places of E over v; q_v denotes the order of the residue field of F_v. Ordinary additive valuations ord_w are normalized by $\mathrm{ord}_w(\pi_w) = 1$. Thus the source's multiplicative exponential valuation is $\exp(-\mathrm{ord}_w)$.

The extension E/F is unramified at every finite place, not the extension $E/\mathbb{Q}$. To establish this property, observe that i and the primitive third roots of unity belong to E; their orders, 4 and 3, are relatively prime. Suppose that conjugation belonged to an inertia group of $E/\mathbb{Q}$ at residue characteristic p. Select one of the two roots, of order n relatively prime to p. Reduction is injective on the n-torsion of units: if $u^n = 1$ and $u \equiv 1$ modulo the maximal ideal, then $(u-1)(1 + u + \cdots + u^{n-1}) = 0$ and the second factor is a unit. But conjugation acts as inversion on the root selected, which has order greater than two. This contradicts trivial inertia action. The inertia group of E/F is the intersection with $\{1, c\}$, where c is complex conjugation. The contradiction proves the assertion, and the local alternative at v is either

$$E \otimes_F F_v = F_v \times F_v \quad \text{or} \quad E \otimes_F F_v/F_v \text{ unramified quadratic.}$$

The primitive 5040-th roots that appear in the definition `StablyUnram` in `Lib/Adelic/Ade`

`les/CoprimeValuation.lean` also belong to the field. In fact 18480 need not itself contain 5040, but the factors $2^s 3^{\lfloor \log_3 s \rfloor + 5}$ ensure $5040 \mid L$ for the admissible $s = 22 \log_2 m$. We shall use a weaker property than the full root-of-unity condition: the local norm of any unit of F_v is attained by a unit of the quadratic algebra above it.

Lemma 6.1 (Norms and traces of integral elements)**.** *Let R be the ring of integers of a non-Archimedean local field, with finite residue field k of order q. Let $S = R \times R$, with involution exchanging factors, or let S be the integers of its unramified quadratic extension. Denote the involution by $z \mapsto \bar{z}$. Then*

$$\mathrm{Tr}_{S/R}(S) = R, \qquad \mathrm{N}_{S/R}(S^\times) = R^\times.$$

In the unramified field case an element of the base local field is a norm if and only if its valuation is even. In the split case the norm on the full group of units of the quadratic algebra of fields is surjective.

Proof. Both claims are immediate for the product algebra by choosing $(a, 0)$ for the trace and $(a, 1)$ for the norm. In the unramified field case the trace of the degree-two residue-field extension k_2/k is surjective. In characteristic two its being nonzero follows from the existence of $z \notin k$ (then $z^q + z \neq 0$); in odd characteristic the trace of 1 is 2. Lift an element of residue trace one to an element of S and divide it by its unit trace. This proves trace surjectivity over R. The map on residue units is $z \mapsto z^{q+1}$, surjective because $k_2^\times$ is cyclic. Given $a \in R^\times$, choose a lift $z_1 \in S^\times$ of a residue preimage. If $z_j \bar{z}_j \equiv a \pmod{\pi^j}$ and $a/(z_j \bar{z}_j) = 1 + \pi^j b_j$, choose $u_j \in S$ with $u_j + \bar{u}_j \equiv b_j \pmod{\pi}$. Then

$$z_{j+1} = z_j(1 + \pi^j u_j), \qquad z_{j+1}\overline{z_{j+1}} \equiv a \pmod{\pi^{j+1}}.$$

The product converges in the complete ring S to a unit of norm a. A uniformizer of R is also a uniformizer of S, and its norm is its square. Arbitrary elements are powers of that uniformizer times units. This proves the valuation assertion, including at residue characteristic two. □

This lemma is the reason we do not need a separately quoted dyadic norm theorem. In particular each conjugation-invariant finite integral unit idele is the product $z\bar{z}$ for an integral unit idele z. Indeed at a conjugate pair one chooses the first coordinate and then the second, and at a fixed w one uses the lemma. All coordinates are units, so the choices define a restricted product.

6.2 A character extension with exactly the required sign

We first isolate the small but decisive arithmetic condition on roots of unity that makes the desired character possible. If δ is a unit of E, write

$$P(\delta) = \prod_{\tau \in \Phi} \tau(\delta).$$

The right side is the image in $\mathbb{C}$ of $N_{E/K}(\delta)$. The following lemma verifies the hypothesis used in `nonempty_splittingChar`, not just an infinity-type existence condition.

Another small consequence of the choice of K is used in the normalized CM reciprocity argument. Since its discriminant is the negative of a prime, the class number of K is odd. Here is a direct proof in place of invoking the general genus theorem.

Lemma 6.2 (Odd class number of the quadratic base)**.** *The ideal class group of $K = \mathbb{Q}(\sqrt{-d})$ has no nonidentity element of order two. In particular the squaring map on it is bijective.*

Proof. Conjugation acts by inversion on ideal classes, since $\mathfrak{a}\bar{\mathfrak{a}}$ is a principal rational ideal. If the class of $\mathfrak{a}$ has order dividing two, then $\bar{\mathfrak{a}} = \lambda\mathfrak{a}$ for some $\lambda \in K^\times$. Applying conjugation twice gives $(\lambda\bar{\lambda})\mathfrak{a} = \mathfrak{a}$; hence $\lambda\bar{\lambda}$ is a positive rational algebraic unit, therefore 1. If $\lambda \neq -1$, let $u = 1 + \lambda$; then $u/\bar{u} = \lambda$. If $\lambda = -1$ use $u = \sqrt{-d}$. In both cases $\mathfrak{b} = u\mathfrak{a}$ is invariant under conjugation. Factor this fractional ideal into primes. At each split rational prime the exponents of the two conjugate factors are equal, and the paired product is principal. At an inert prime the prime ideal is generated by the rational prime. The only ramified prime lies above d, and it is the principal ideal generated by $\sqrt{-d}$. Thus $\mathfrak{b}$ is principal. Since the class group is finite, the assertion about squaring follows from the structure of a finite abelian group. □

The last assertion is stronger than merely saying that the order of the ramified prime class is odd. It allows one to choose an ideal-class square root of any desired CM ideal class in the normalized reciprocity calculation in Chapter 12.

Lemma 6.3. *For every root of unity $\xi \in E$ there exists a unit $\delta \in \mathcal{O}_E^\times$ with*

$$\delta/\bar{\delta} = \xi, \qquad P(\delta) = 1.$$

Proof. As in Chapter 4, the entire group of roots of unity of E is generated by ζ. The element $\delta_0 = i(\zeta - 1)$ satisfies $\bar{\delta}_0 = \delta_0/\zeta$. The norm of $\zeta - 1$ from $\mathbb{Q}(\zeta)$ to $\mathbb{Q}$ is $\Phi_L(1) = 1$, since L has more than one prime divisor. For completeness, the identity $X^n - 1 = \prod_{a|n} \Phi_a(X)$, divided by $X - 1$ and evaluated at 1, gives

$$n = \prod_{1<a|n} \Phi_a(1).$$

Induction on n shows that $\Phi_{p^e}(1) = p$ and that every other factor is 1. Thus $\zeta - 1$ is a unit. Its norm over K is the same rational number as its cyclotomic norm, since K and $\mathbb{Q}(\zeta)$ are linearly disjoint. The norm of i is 1 too: the cyclotomic automorphisms pair each embedding with its conjugate, and each product of the two conjugates of i is 1. Hence $P(\delta_0) = 1$. For $\xi = \zeta^a$ take $\delta = \delta_0^a$, allowing negative a because δ_0 is a unit. □

Let $J_E = \mathbb{A}_E^\times$, $U_f = \prod_{w<\infty} \mathcal{O}_w^\times$, and $H = E_\infty^\times U_f$, an open subgroup of J_E. Define the continuous unitary character

$$S(x) = \prod_{\tau\in\Phi} \frac{x_\tau}{|x_\tau|} \qquad (x \in H).$$

One would like to prescribe this character on H, require triviality on $E^\times$, and also require $\upsilon(\bar{x}) = \upsilon(x)^{-1}$. These requirements are compatible precisely when S is trivial on

$$H \cap E^\times\{z\bar{z} : z \in J_E\}.$$

Checking this last intersection is important; choosing an arbitrary unramified Hecke character of the right infinity type does not prove the conjugation requirement.

We use the following standard consequence of global class field theory: for the quadratic extension E/F the global Artin character $\eta : J_F \to \{\pm 1\}$ has kernel $F^\times N_{E/F}(J_E)$; its local restriction is the usual norm character, at real places the sign, at unramified inert finite places $(-1)^{\mathrm{ord}_v}$, and at split places trivial. This is the quadratic norm subgroup theorem (e.g. Neukirch, [11], global reciprocity). These assertions apply to the full ideles and do not presume GRH.

Proposition 6.4 (Existence of the normalized splitting character)**.** *There is a continuous character $\upsilon : J_E \to \{z \in \mathbb{C} : |z| = 1\}$ such that*

1. *$\upsilon(a) = 1$ for $a \in E^\times$;*
2. *$\upsilon|_H = S$;*
3. *$\upsilon(\bar{x}) = \upsilon(x)^{-1}$ for $x \in J_E$;*
4. *its restriction along $J_F \hookrightarrow J_E$ is η.*

In particular its value on every finite uniformizer idele supported at a conjugation-fixed place is -1.

Proof. Set $D = E^\times\{z\bar{z} : z \in J_E\}$. Suppose $x = az\bar{z} \in H \cap D$. At every finite w, the equality $\mathrm{ord}_w(x) = 0$ implies

$$\mathrm{ord}_w(a) = -\mathrm{ord}_w(z) - \mathrm{ord}_{\bar{w}}(z),$$

also when $w = \bar{w}$. Thus the fractional ideals of a and $\bar{a}$ coincide, so $\bar{a}/a$ is an algebraic unit. At every complex embedding it has modulus 1. Such a unit is a root of unity: the polynomials of its successive powers are monic integral of bounded degree with bounded coefficients, because their roots have modulus 1. Only finitely many such powers can occur, as each polynomial has finitely many roots. Lemma 6.3 supplies a unit δ with $\delta/\bar{\delta} = \bar{a}/a$ and $P(\delta) = 1$. Then $A = a\delta$ lies in $F^\times$.

The finite idele $Az_f\bar{z}_f = \delta x_f$ is an integral unit and conjugation invariant. By Lemma 6.1 it is a norm of an integral unit idele u_f of E. Consequently A^{-1} is a finite idele norm. Global reciprocity for the diagonal element A implies

$$\prod_{\tau \in \Phi} \mathrm{sgn}(\tau(A)) = 1.$$

The archimedean coordinates of $z\bar{z}$ are positive, and $P(\delta) = 1$ implies $S(\delta) = 1$. Therefore $S(x) = S(a) = S(A) = 1$. This establishes the intersection condition.

On HD define $\upsilon_0(hd) = S(h)$. The intersection condition makes the definition independent of the representation. A homomorphism from a subgroup of an abelian group to the circle extends to the full group. Here is the short algebraic proof of this assertion. Order extensions by inclusion of their domains and take a maximal extension using the union along chains. If g is outside its domain B, either no positive power of g lies in B, in which case its value can be chosen freely, or there is a least positive m with $g^m \in B$. Choose an m-th root of the already prescribed value of g^m in the circle. Defining the values on $B\langle g\rangle$ multiplicatively is well-defined by division with remainder modulo m. Either case contradicts maximality unless every element is in the domain. Apply this to υ_0 to obtain υ on J_E. Since it is continuous on the open subgroup H containing the identity, the extension is continuous everywhere. Triviality on D proves the principal and conjugation assertions.

The restriction to J_F is trivial on $F^\times$ and on idele norms from E. The global norm subgroup has index two, and the restriction is not trivial: on the idele with one real coordinate -1, all others $+1$ and finite part 1, its value is -1. It follows that this restriction equals η. If $w = \bar{w}$, the local extension is unramified quadratic, a uniformizer of F_v is a uniformizer of E_w, and η_v sends it to -1. An arbitrary uniformizer differs from it by a unit, on which υ is trivial. This proves the last assertion. □

Notice that the proof used the full conjugation condition, not merely the identity $\upsilon(a) = 1$ on rational a. The possible choices of the extension are immaterial to its restrictions used here.

The existence of one such choice is exactly what the source's record demands. Its construction is the classical proof by extending a character to a divisible group, with the special compatibility just checked. The source anchor is `Lib/ClassField/HeckeCharacter/SplittingCharExists.lean`; the auxiliary norm decomposition in `Lib/ClassField/Global/CMNorms.lean` follows from the same global Artin character:

Corollary 6.5 (Signs of norms and decomposition of finite ideles)**.** *Let $a \in J_{F,f}$. It can be written as $AN(z_f)$ with $A \in F^\times$. If $a = 1$, the sign vector of A may be any vector of product $+1$, and no sign vector of product -1 is possible.*

Proof. Extend the finite idele a to a full idele of F whose real sign vector has product equal to the value of the finite quadratic Artin character on a. Such a sign vector always exists. The full idele then has trivial Artin character, and global reciprocity writes it as $AN(z)$ for $A \in F^\times$ and $z \in J_E$. Restriction to its finite part gives the desired decomposition. When $a = 1$ take the requested real sign vector of product $+1$; positivity of the real local norms implies that the resulting A has exactly those signs. Necessity follows from reciprocity for A. □

Finite ideal notation. For an unramified character on J_E, trivial on U_f, fix a finite idele y and put $m_w = \operatorname{ord}_w(y_w)$. Since $y(\prod_w \pi_w^{-m_w})$ is an integral unit idele,

$$\upsilon(x_\infty, y) = \left(\prod_{\tau \in \Phi} x_\tau / |x_\tau| \right) \prod_w \upsilon(\pi_w)^{m_w}. \tag{6.1}$$

At a split pair the two finite generator values are inverses; at a fixed place the generator value is -1. These are `upsilon_eq_prod_mul_of_v_eq`, `upsilon_eq_atUnif_zpow` and `upsilon_atUnif_mul_star` in `Lib/General/NumberTheory/UpsilonValues.lean`. There is no choice of a logarithm in (6.1).

6.3 The self-dual lattice and half of the different

Let $p_L(X) = \Phi_L(X)$ and $D_K = p_L'(\zeta)$. These denote the polynomial and a generator of the relative different, not a discriminant. The integral ring of E equals $\mathcal{O}_K[\zeta]$, as established in Chapter 4, and the absolute different $\mathfrak{d}_E$ is the principal ideal $(\sqrt{-d}D_K)$. Define $c_0 = (1 + \sqrt{-d})/2$, $\jmath = -1/\sqrt{-d}$, and the finite idele α with

$$\alpha_w = \begin{cases} D_K^{-1}, & c_0 \in w, \\ 1, & c_0 \notin w. \end{cases}$$

The equality $\operatorname{Im} \tau(\jmath) = 1/\sqrt{d}$ for every $\tau \in \Phi$ fixes the real sign, and $(1 - \jmath^{-1})/2 = c_0$ is integral. In the finite part, the assertion to be verified is

$$\operatorname{ord}_w(\jmath \alpha_w \overline{\alpha}_w) = -\operatorname{ord}_w(\mathfrak{d}_E). \tag{6.2}$$

Here $\overline{\alpha}_w$ means the conjugate of the coordinate at $\bar{w}$, not an automorphism of the w-coordinate if that place is split.

Lemma 6.6. *The idele α satisfies (6.2) at every w. Moreover α^{-1} and $L\alpha$ belong to $\widehat{\mathcal{O}}_E$.*

Proof. First, D_K divides L in $\mathcal{O}_E$: differentiate $X^L - 1 = p_L(X)q(X)$ at ζ, obtaining $L\zeta^{L-1} = D_K q(\zeta)$. Thus the valuations of D_K are supported at primes over divisors of L.

Also $\bar{D}_K$ differs from D_K by a root of unity and sign, by the reciprocal polynomial identity. In particular $\mathrm{ord}_w(D_K) = \mathrm{ord}_{\bar{w}}(D_K)$.

If $p \mid L$, then $p \mid N_0$, and $c_0\bar{c}_0 = (d+1)/4$ is divisible by p (also for $p = 2$, since $d \equiv 7 \pmod 8$). Since $c_0 + \bar{c}_0 = 1$, at each w over p exactly one of these elements lies in the maximal ideal. Conjugation exchanges the two alternatives. It follows that

$$\mathrm{ord}_w(\alpha\bar{\alpha}) = -\,\mathrm{ord}_w(D_K)$$

at such a place. At every other place both sides are zero. Multiplying by $\jmath$ gives (6.2) using the product formula for the different in the tower. The final two assertions are immediate from the definition and $D_K \mid L$. □

For reference the source calls α `Parameters.a`, and the preceding lemma is `v_alpha_mul_relDifferentGen` together with `inv_a_mem` and the divisibility by N_1. The construction of $\jmath$, α and υ has now established the full record `Parameters.instWeilDataE`, using global reciprocity as indicated.

What self-duality means. For a finite local field E_w, let ψ_w be the rational additive character composed with $\mathrm{Tr}_{E_w/\mathbb{Q}_p}$. Its annihilator on $\mathcal{O}_w$ is $\mathfrak{d}^{-1}_{E_w/\mathbb{Q}_p}$. Indeed the character is trivial on all $\psi_w(za)$, $a \in \mathcal{O}_w$, exactly when $\mathrm{Tr}(z\mathcal{O}_w) \subset \mathbb{Z}_p$, which is the definition of the inverse different. Therefore the bilinear annihilator of $\widehat{\mathcal{O}}_E$ for

$$(x, y) \longmapsto \psi_f(\jmath\alpha\bar{\alpha}x\bar{y})$$

is precisely $\widehat{\mathcal{O}}_E$, by (6.2). In the split case the annihilator condition at a pair is tested by varying the components of y independently, so the reasoning applies to the whole restricted product. With Haar measure giving the integral finite adelic lattice volume one, the Fourier transform of its indicator for this pairing is again that indicator. To see this, if x lies in the annihilator the character is identically one on the lattice; otherwise translation by an element on which the character is nontrivial multiplies the integral by a scalar different from one, forcing it to vanish. The square of Fourier transform is reflection, so this is the self-dual measure for the scaled pairing.

The corresponding real scale is

$$s_\tau = (2\pi \operatorname{Im}\tau(\jmath))^{-1/2}, \qquad s_f = \alpha.$$

This is `WeilData.scaleIdele`. At a complex place the alternating form $\mathrm{Tr}_{\mathbb{C}/\mathbb{R}}(\jmath z\bar{w})$ is $-2\operatorname{Im}\jmath\,\operatorname{Im}(z\bar{w})$. Under the scaling $z \mapsto s_\tau z$, its additive-character exponent becomes the standard symplectic exponent. This calculation explains why the factor 2π is present rather than a separate adjustable scalar. The finite lattice Fourier identity just proved does not establish by itself a splitting of the metaplectic representation or a global theta identity; those are treated separately in Chapter 12.

6.4 Reduction of finite ideals and their adelic generators

We record one auxiliary construction used repeatedly when a level or denominator ideal is inserted; it is `exists_reduction_and_idele_generator` in `Lib/General/NumberTheory/UpsilonValues.lean`. For a nonzero integral ideal $I \subset \mathcal{O}_E$, let

$$e_w = \mathrm{ord}_w(I), \qquad t_w = \begin{cases} \pi_w^{e_w}, & e_w > 0, \\ 1, & e_w = 0. \end{cases}$$

Then t is an integral finite idele. The natural map from $\mathcal{O}_E/I$ to $\prod_{w|I} \mathcal{O}_w/\pi_w^{e_w}\mathcal{O}_w$ is an isomorphism: apply the Chinese remainder theorem and note that completion does not change a quotient by a finite power of the maximal ideal. Composing its inverse with the projections defines

$$\rho_I : \widehat{\mathcal{O}}_E \longrightarrow \mathcal{O}_E/I, \qquad \ker \rho_I = I\widehat{\mathcal{O}}_E = t\widehat{\mathcal{O}}_E.$$

It restricts to ordinary reduction of algebraic integers. For any finite adele y, one has

$$(\forall c \in I,\ cy \in \widehat{\mathcal{O}}_E) \quad \Longleftrightarrow \quad ty \in \widehat{\mathcal{O}}_E.$$

One implication uses $c_w \in t_w\mathcal{O}_w$ for every c. For the other, at each w choose $c \in I$ with valuation exactly e_w; it exists because $I\mathcal{O}_w = \pi_w^{e_w}\mathcal{O}_w$ and the localization of I is dense in this ideal. The hypothesized integrality of cy forces $\mathrm{ord}_w(y) \geq -e_w$. Since this works at all w, ty is integral. This proof uses I as a nonzero ideal, not a chosen principal generator in the number field. This distinction matters in the unitary class sums, where the global class need not be principal.

6.5 Completions of the coefficient primes and the choice of a side

Here we give a direct proof of the decomposition-group and chosen-place statements of `New/Data/WeilDataOfE.lean` that are used to label the finite factors of the groups. The Galois group over $\mathbb{Q}$ can be written as

$$(\mathbb{Z}/L)^\times \times \{\pm 1\}, \qquad (a,\epsilon):\ \sqrt{-d} \mapsto \epsilon\sqrt{-d}, \quad \zeta \mapsto \zeta^{\epsilon a}.$$

It is important to include the factor ϵ in the second displayed action. It arises because the first coordinate refers to the CM type.

Proposition 6.7. *For $p \nmid Ld$, put $\epsilon = \chi_K(p)$ and $a = \epsilon p \pmod L$. Then the decomposition group of every place above p in $E/\mathbb{Q}$ is cyclic generated by (a,ϵ). Its order is the residue degree. If $\epsilon = -1$, the elements of $\mathrm{Gal}(E/K)$ which fix a prescribed place are the even powers of $-p$; those mapping it to its conjugate are the odd powers of $-p$.*

Proof. The prime is unramified in $E/\mathbb{Q}$ because this extension is contained in the cyclotomic extension of conductor Ld. Arithmetic Frobenius acts as p-th power on reduction of ζ, and as the Legendre-symbol action on $\sqrt{-d}$. The reduction of L-torsion is injective, and reduction of $\sqrt{-d}$ is nonzero. Consequently the congruence characterizes the stated Galois element. The decomposition group of an unramified prime is the cyclic group generated by Frobenius; the reduction map identifies it with the Galois group of the residue fields. An element of $\mathrm{Gal}(E/K)$ is $(b,1)$. It fixes a place if it is in the decomposition group, and maps the place to its conjugate if multiplying by $(1,-1)$ belongs to the decomposition group. For the inert case these conditions give respectively the even and odd powers. For example, when the order of $-p$ is odd, an odd power can equal the identity in $(\mathbb{Z}/L)^\times$, and the place can be its own conjugate. No set of chosen representatives should identify this situation with a split prime. □

For $p \equiv 1 \pmod{Ld}$ decomposition groups are trivial and all residues have cardinal p. Suppose $b \in \mathcal{O}_E$ and $b\bar{b} = pc$ with $p \nmid c$. At a place $w \mid p$, the valuations satisfy

$$\mathrm{ord}_w(b) + \mathrm{ord}_w(\bar{b}) = 1.$$

Both are nonnegative integers, and conjugation exchanges the two valuations. Thus exactly one place in each conjugate pair divides b; its valuation of b is one. The unconjugated set of places in a Galois orbit under $\mathrm{Gal}(E/K)$ contains one place in each conjugate pair, though not necessarily the place dividing b. Starting from one place and for each $g \in \mathrm{Gal}(E/K)$ selecting gw or $\overline{gw}$ according to divisibility by b therefore gives a bijection

$$\Phi \simeq \{w \mid p : b \in w\}.$$

The domain has f elements, and injectivity follows from the trivial decomposition group and the fact that no element of $\mathrm{Gal}(E/K)$ equals conjugation. This proves the assertions `star_mem_iff`, `intValuation_eq_of_mul_star` and `nonempty_infinitePlace_equiv_finPlace`. In the case $p = g_i$, b is the element π_g (or its conjugate, according to the chosen side), and the cofactor is g_{1-i}. In the case $p = N$ the cofactor is one. This argument is the precise reason reduction at the selected places can be modeled by independent general linear factors rather than by a quadratic unitary group.

6.6 Unimodular lattices over split and unramified places

The local integral group calculation begins with a simple form of Gram–Schmidt. We use R, S as in Lemma 6.1. A Hermitian form on an S-module M is conjugate-linear in the first argument and S-linear in the second. It is unimodular if its Gram matrix in an S-basis is invertible over S.

Proposition 6.8. *Every free unimodular Hermitian S-lattice of rank n admits an orthonormal basis. If $v_1, \dots, v_l$ are already orthonormal vectors, they extend to one. Consequently the unitary group of the standard unimodular lattice acts transitively on orthonormal ordered l-frames over S.*

Proof. In the unramified field case, reduce the Gram matrix to the residue field k_2. A nondegenerate Hermitian form over k_2 has a vector with nonzero norm. To prove this, if every norm is zero, then for every x, y, b,

$$h(x + by, x + by) = \mathrm{Tr}_{k_2/k}(bh(x, y)) = 0.$$

The trace duality on the finite separable field extension is nondegenerate; indeed for $h(x, y) \neq 0$ multiply its inverse by a trace-nonzero element. It follows that the form is zero, a contradiction. Lift a vector of nonzero norm and use norm surjectivity on units to rescale it to norm one. Its projection $x \mapsto vh(v, x)$ splits an orthogonal summand, and the orthogonal complement is a finite free unimodular lattice since it is a projective module over a local principal ideal domain. Induct on rank. The same projection formula applied to the given orthonormal vectors proves the extension assertion.

In the split case, write $M = M_1 \times M_2$, where both M_i are free R-modules of rank n. A Hermitian form is precisely a perfect pairing $M_2 \times M_1 \to R$ together with its transpose. Choose a basis of M_1 and its dual basis of M_2; this gives an orthonormal basis of M. Given an orthonormal sequence, the corresponding projections again split both modules, and the complementary pairing is perfect. The group sends one chosen orthonormal basis to the other, proving transitivity. □

This proof of integral Witt extension applies at dyadic places too: the trace calculation rather than division by two is its essential step. Over a field the usual Witt extension theorem gives transitivity on norm-one vectors in a fixed nonsingular Hermitian space. In the definite

global standard space we can also see it directly. If u, v have norm one and $u \neq v$, then $w = u - v$ has nonzero norm because the form is positive at infinity. The transformation $T(x) = x - wh(w,x)/h(w,u)$ sends u to v and preserves the form. Indeed $h(w,u) \neq 0$ (otherwise $h(v,u) = 1$ and $h(w,w) = 0$), and

$$h(w,w) = h(w,u) + \overline{h(w,u)}.$$

Expansion of $h(Tx, Ty)$ gives the original pairing. Thus T is invertible and unitary. The identity transformation treats $u = v$. This elementary reflection avoids needing a more general version of Witt's theorem for the norm-one mass calculation below.

6.7 Existence and equivalence of the framed global spaces

The group used at the finite places has the identity Gram matrix, while at the real places the same group need not be compact. We now show that these local conventions can be imposed simultaneously by an actual Hermitian space over E. In the source a `Space` with sign pattern s is given by a Hermitian matrix J over E and a matrix $x \in \mathrm{GL}_n(\mathbb{A}_E)$ such that

$$x^* J x = J_s, qquad (J_s)_\tau = \mathrm{diag}(s_1, \dots, s_n), \quad (J_s)_f = I_n.$$

Here $s_i \in \{\pm 1\}$ and the same pattern is used at every real place of F; the notation $*$ means the transpose of conjugation, evaluated on the appropriately paired local factors. The rational unitary group acts in these coordinates by

$$\gamma \longmapsto x^{-1} \gamma x.$$

In particular a rational matrix need not be a rational point of the group whose finite coordinates have the identity matrix. The conjugating matrix x is part of the data; it determines the integral lattices and levels.

Proposition 6.9 (Existence of the normalized signed space)**.** *For every sign pattern s of length n there exists such a pair (J, x). More generally, for each coordinate one can prescribe a sign at each real place provided the product of its signs is positive.*

Proof. In the present field $f = \varphi(L)$ is even. Consequently a sign vector that is constantly -1 has product $+1$; so does the constant positive vector. For each i apply Corollary 6.5 with $a = 1$ and requested signs s_i (or the stipulated varying signs). It gives $a_i \in F^\times$ with these signs and $a_i N(z_{i,f}) = 1$. Put

$$J = \mathrm{diag}(a_i), \qquad x_f = \mathrm{diag}(z_{i,f}), \qquad x_\tau = \mathrm{diag}(|\tau(a_i)|^{-1/2}).$$

These matrices satisfy the required equality. The finite entries $z_{i,f}$ are ideles, not merely independently chosen nonzero elements with uncontrolled denominators. Thus x is an element of the restricted adelic general linear group. This gives the sought global realization and proves the additional sign assertion. □

For equivalence we invoke the classical Landherr theorem over a number field: nonsingular Hermitian forms for a quadratic extension are isometric if and only if they are isometric at all places of the base field. This is a Hasse principle for all signatures, not a strong approximation theorem for a definite group. A reference is Landherr, [13], and Kneser's lectures on Galois cohomology, Section 5.5. The hypothesis of this theorem is especially transparent here: the adelic frames explicitly exhibit the local isometries, including at the places over 2.

Proposition 6.10 (Equivalence with movement of the adelic frame)**.** *Suppose (J, x) and (J', x') realize the same sign pattern. There exist $a \in \mathrm{GL}_n(E)$ and $g \in G_s(\mathbb{A}_F)$, the common adelic group of J_s, with*

$$a^* J a = J', \qquad x' = a^{-1} x g.$$

If γ belongs to the unitary group of J', then

$$x^{-1}(a\gamma a^{-1})x = g\,(x'^{-1}\gamma x')\,g^{-1}.$$

Proof. At every place the matrix $x(x')^{-1}$ is a local isometry from J' to J. By the classical Landherr theorem these local isometries have a rational isometry a (not necessarily approximating the given local matrices). Put $g = x^{-1}ax'$. Then $g^* J_s g = (x')^* a^* J a x' = J_s$, so g lies in the prescribed adelic group. Rearranging its definition proves the second equality, and substituting into $x'^{-1}\gamma x'$ gives the conjugation formula. □

This is the content of `Space.landherr` and `Space.exists_eq_move` in `Lib/Adelic/Landherr/Landherr.lean`. The special signed realization does not require the full theorem; it follows directly from the finite norm subgroup. Equivalence is a standard classical application. The quantitative comparison of the actual frames in the nontrivial theta kernel is a stronger assertion: the theorem only gives a rational isometry and imposes no bound on its archimedean entries or on the chosen lattice denominator. In particular this equivalence cannot by itself complete a quantitative upper bound or transfer an integral algebraic form without a finite-lattice calculation. Later comparisons shall keep the chosen x and g and specify those additional bounds.

6.8 Counting the local finite groups and their congruence quotients

The standard finite integral unitary group is hyperspecial in our unramified setting; what we need is the following explicit description, which can be proved without that terminology. Let $\mathcal{U}_n(S) = \{g \in \mathrm{GL}_n(S) : \bar{g}^t g = 1\}$ and π be a uniformizer of R. The reduction maps

$$\mathcal{U}_n(S) \longrightarrow \mathcal{U}_n(S/\pi^e S)$$

are surjective. We first show this directly in the unramified case. If g_j is a unitary lift modulo π^j, put $C = (\bar{g}_j^t g_j - I)/\pi^j$ modulo π. It is Hermitian. Choose X modulo π with $X + \bar{X}^t = -C$: for $i < l$, set $X_{il} = -C_{il}$ and $X_{li} = 0$; on the diagonal choose an element of residue trace $-C_{ii}$. Then $g_{j+1} = g_j(1 + \pi^j X)$ satisfies the equation modulo π^{j+1}. These lifts converge. The solution space for X in one step is an affine translate of the skew-Hermitian matrices. Over k this space has dimension n^2: each of $n(n-1)/2$ off-diagonal pairs contributes two dimensions, and each diagonal contributes one. Hence each fiber of reduction by one exponent has cardinal q^{n^2}. In the split case the group identifies with $\mathrm{GL}_n(R)$ via the pair $(g, (g^t)^{-1})$. Reduction and these fiber sizes are then immediate.

Lemma 6.11. *Over the residue field k of cardinal q, the split and nonsplit group orders are respectively*

$$|\,\mathrm{GL}_n(k)| = q^{n^2} \prod_{i=1}^{n} (1 - q^{-i}), \tag{6.3}$$

$$|\mathcal{U}_n(k_2)| = q^{n^2} \prod_{i=1}^{n} (1 - (-1)^i q^{-i}). \tag{6.4}$$

At exponent $e \geq 1$ multiply these orders by $q^{(e-1)n^2}$.

Proof. For (6.3), count successive basis vectors: $q^n - q^j$ choices for vector $j+1$. In the nonsplit case we count the vectors of norm one. Fix a nontrivial additive character $\psi : k \to \mathbb{C}^\times$. For $a \neq 0$,

$$\sum_{z \in k_2} \psi(az\bar{z}) = 1 + (q+1) \sum_{b \in k^\times} \psi(ab) = -q.$$

By additive orthogonality, the number of vectors $z \in k_2^n$ of norm one is

$$\frac{1}{q} \sum_{a \in k} \psi(-a) \left(\sum_{z \in k_2} \psi(az\bar{z}) \right)^n = q^{2n-1} - (-1)^n q^{n-1}.$$

Every such vector extends to an orthonormal basis by the residue-field part of the proof of Proposition 6.8; its stabilizer is $\mathcal{U}_{n-1}(k_2)$. Thus the order is the displayed count times the order in rank $n-1$. Starting with rank zero gives (6.4). The fiber calculation above proves the general reduction-order assertion. □

At each chosen $w \mid g_i$, the other conjugate place gives the inverse transpose, so the finite quotient K_0/K_g for rank n is

$$\prod_{i=0,1} \prod_{w \text{ chosen} \mid g_i} \mathrm{GL}_n(\mathbb{F}_{g_i}).$$

Its order is at most g^{fn^2}, using f places above each g_i and (6.3). The block inclusion of a group of rank n in rank rt proves divisibility into `Data.chosenGroupOrder` whenever $n \leq rt$. This verifies both the size and divisibility ingredients of `New/Data/ChosenPlaceLevels.lean`; the proof that stabilizers inject into this quotient is the neatness argument in Chapter 5. We stress that using only an upper bound for the stabilizer order is insufficient for denominator clearing: divisibility of its order into the common integer is what is required.

6.9 An exact norm-one class-sum identity

Before applying a mass formula it is worth understanding why consecutive ranks appear naturally. Let $G_n = U(I_n)$, viewed as a group over F, and $K_n \subset G_n(\mathbb{A}_{F,f})$ the stabilizer of $\widehat{\mathcal{O}}_E^n$. Denote the genus mass by

$$\mathfrak{m}_n = \sum_{[g] \in G_n(F) \backslash G_n(\mathbb{A}_{F,f}) / K_n} \frac{1}{|\Gamma_g|}, \qquad \Gamma_g = G_n(F) \cap g K_n g^{-1}.$$

The finiteness and class-sum convention are those of Chapter 5. At rank zero put $\mathfrak{m}_0 = 1$; in positive ranks the sum is positive.

Proposition 6.12 (The norm-one identity)**.** *Let $n \geq 1$ and let $c(g)$ count $y \in E^n$ with $h(y,y) = 1$ and $g^{-1}y \in \widehat{\mathcal{O}}_E^n$. Then*

$$\sum_{[g]} \frac{c(g)}{|\Gamma_g|} = \mathfrak{m}_{n-1}. \tag{6.5}$$

Proof. The set counted by $c(g)$ is finite: the finite integrality condition defines a lattice in the real product E_∞^n, and the norm condition imposes a compact real subset. Consider pairs (gK_n, y) with the stated conditions modulo the diagonal action of $G_n(F)$. The left hand side of (6.5) is the sum over orbits of reciprocal stabilizer orders. This follows from the finite orbit-stabilizer formula: for fixed g, $\sum_{O \in \Gamma_g \backslash X_g} |\mathrm{Stab}(O)|^{-1} = |X_g|/|\Gamma_g|$.

Every rational norm-one vector can be sent to e_n by a rational unitary transformation, by the reflection in Proposition 6.8. Its stabilizer is the naturally embedded $G_{n-1}(F)$. For a finite adelic element g with $g^{-1}e_n$ integral, Proposition 6.8 at every finite place gives $k \in K_n$ with $ke_n = g^{-1}e_n$. Then gk fixes e_n, so belongs to $G_{n-1}(\mathbb{A}_{F,f})$. All such pairs are obtained in this way. Two representatives with vector e_n lie in the same rational orbit exactly when their finite adelic G_{n-1} cosets lie in the same double coset modulo $G_{n-1}(F)$ and K_{n-1}, because any rational map between them must fix e_n and $K_n \cap G_{n-1}(\mathbb{A}_{F,f}) = K_{n-1}$. Stabilizer groups in that description agree. Thus the right side of (6.5) counts exactly the same orbits with the same weights. □

This gives `classSum_card_norm_one_eq_mass` in `Lib/Adelic/Transference/MassBasic.lean`. It is not merely a comparison between two class numbers. A mass has a reciprocal automorphism weight, and the exact identity uses that weight even in the presence of a non-free level. At a level with trivial stabilizer these weights disappear; this useful special case does not justify replacing them at maximal level.

6.10 The classical mass formula in the present normalization

There are two routes to the mass estimate. The compact code proves it by combining (6.5) with the one-vector Siegel–Weil formula and a uniform Fourier estimate. For a classical exposition we can instead use the standard exact Hermitian mass formula directly, since the lattice is unimodular and E/F is unramified at all finite places. This bypasses the Fourier estimate for this particular purpose only; it does not remove the separate Siegel–Weil period identities used in the lower bound.

We state carefully the specialization being used. For a CM quadratic extension E/F with relative discriminant ideal 1, put $D_F = |\operatorname{disc}(F)|$ and let η be the associated quadratic character of F. For the definite unimodular lattice of positive rank n the formula is

$$\mathfrak{m}_n = 2D_F^{n^2/2} \left(\prod_{j=1}^{n} \frac{(j-1)!}{(2\pi)^j} \right)^f \prod_{j=1}^{n} L_F(j, \eta^j). \tag{6.6}$$

In the product η^j is trivial for even j; $L_F(1, \eta)$ is understood as the usual nontrivial quadratic value. This is Proposition 4.2.7(3) of Kirschmer, [12], with the local factors of Lemma 4.2.9 there all equal to 1 (every prime is good in the sense of Definition 3.5.1, since there is no ramified quadratic extension at a dyadic place): at a split or inert unramified prime the entire lattice is the single unimodular Jordan block ($m_1 = 0$). The archimedean factor in that reference is $\gamma_n^{-1} = \prod_{j=1}^{n}(j-1)!/(2\pi)^j$. The mass is the sum of reciprocal orders of the full unitary automorphism groups, not of the special unitary groups, and the leading factor 2 is the one for this convention. In particular the two groups do not differ by a harmless constant if the formula is used at rank zero.

One can check the Euler factors local by local. At an unramified split place the volume-normalized reduction order of Lemma 6.11 is $\prod_{j=1}^{n}(1 - q_v^{-j})$. At an inert place it is $\prod_{j=1}^{n}(1 - (-1)^j q_v^{-j})$. Their formal reciprocals are the Euler factors in (6.6); the exceptional first factor is the nontrivial character value, as required by the Siegel formula rather than an unregularized absolute product at 1. This confirms the convention. The classical formula, not just these finite counts, gives the equality of the global mass with the completed product; it is the explicit external theorem in this subsection.

Theorem 6.13 (Consecutive-mass estimate)**.** *Let* $r_E = |\operatorname{disc}(E)|^{1/(2f)}$*. If* $8f \le 2^u$ *and* $j \ge 4u + 12$*, then*

$$\left|\log \mathfrak{m}_{j+1} - \log \mathfrak{m}_j - f\left((j+\tfrac{1}{2})\log r_E + \log j! - (j+1)\log(2\pi)\right)\right| \le 2\,2^{-u}. \tag{6.7}$$

Proof. The discriminant tower formula gives $|\operatorname{disc}(E)| = D_F^2$, hence $D_F = r_E^f$. Taking the quotient of (6.6) at $j+1$ and j cancels all terms with index at most j, including the delicate $L_F(1,\eta)$ term. With $n = j+1$ the exact quotient is

$$\frac{\mathfrak{m}_{j+1}}{\mathfrak{m}_j} = \left(r_E^{j+1/2}\frac{j!}{(2\pi)^{j+1}}\right)^f L_F(n,\eta^n).$$

For $n \ge 3$, the Euler factors are positive (including when $\eta^n(v) = -1$), so the logarithm is real and absolutely convergent. For the relative unramified character η,

$$\begin{aligned}|\log L_F(n,\eta^n)| &\le \sum_{v<\infty}\sum_{h\ge 1}\frac{q_v^{-nh}}{h}\\ &\le f\log\zeta(n) \le f(\zeta(n)-1) \le 2f\,2^{-n}.\end{aligned}$$

The middle bound follows by grouping places over each rational prime: there are at most f places and $q_v \ge p$. For the last inequality use $\sum_{a\ge 2} a^{-n} \le 2^{-n} + \int_2^\infty t^{-n}dt = 2^{-n}(1 + 2/(n-1)) \le 2\,2^{-n}$. Finally $2f2^{-n} \le 2^u 2^{-n}/4 \le 2\,2^{-u}$ for the stated bound on n. Substituting the exact quotient proves (6.7). □

The conclusion is exactly `abs_log_mass_ratio_sub_le` in `Lib/SiegelWeil/Mass/MassRatio.lean`. This independent route shows why the exceptionally large rank in the source makes the mass correction uniform even though f is growing. The small factor is an L-value at a large positive integer, not a value near the critical strip. No assertion about zeros of the Dirichlet L-functions under study has entered. The classical mass formula is cited rather than re-proved in this volume; the derivation of the required uniform inequality, and its justification for the exact lattice and conventions of the proof, are given here.

6.11 Useful consequences and the separation of normalizations

The exact formula supplies convenient telescoping identities. For integers $r \ge 1$, $s \ge 0$ such that $r \ge 4u + 12$ and $8f \le 2^u$, one obtains

$$\log\frac{\mathfrak{m}_{r+s}}{\mathfrak{m}_r} = f\sum_{j=r}^{r+s-1}\left((j+\tfrac{1}{2})\log r_E + \log j! - (j+1)\log(2\pi)\right) + \mathcal{R}, \tag{6.8}$$

$$|\mathcal{R}| \le 4f\,2^{-r-1}. \tag{6.9}$$

The latter, stronger than adding $2/2^u$ at each step, follows by summing the geometric bound $2f\,2^{-(j+1)}$ established in the proof. An empty sum has error zero. In particular the apparent growth of the number of consecutive ratios costs no extra factor of s. The identities

$$\sum_{j=r}^{r+s-1}(j+\tfrac{1}{2}) = rs + \tfrac{s^2}{2}, \qquad \sum_{j=r}^{r+s-1}(j+1) = rs + \tfrac{s(s+1)}{2}$$

are often useful when comparing dimensions of successive groups. They have the same weights that appear in the common leading terms in the final assembly.

We also explain an elementary normalization check relating the finite class count to a real coefficient space. If $K' \subset K_n$ is an open subgroup, the total mass at K' is $[K_n : K']\mathfrak{m}_n$: indeed over a K_n-class the fiber groupoid is the action of Γ_g on the $[K_n : K']$ cosets. By orbit-stabilizer its mass is $[K_n : K']/|\Gamma_g|$. If every stabilizer at K' is trivial this mass is the number of K'-classes. A coefficient representation of dimension d then supplies precisely d independent coordinates per class. Thus the scalar space of definite algebraic forms has dimension $d[K_n : K']\mathfrak{m}_n$. This argument does not show that the automorphic form space for an indefinite real group has the same dimension; that requires the transfer or filling argument. In particular one cannot replace a trace formula by this formula for the definite group without supplying the transfer.

This distinction also clarifies three different uses of measures:

1. The product of finite additive Haar measures used for the integral Weil lattice, normalized by self-duality as in (6.2).

2. The finite unitary quotient measure, for which the maximal integral unitary group has volume one and the total volume at a definite real place is the genus mass.

3. The real Haar and formal-degree factors, whose exact normalizations are required for the discrete-series projector and the periods. These are not determined solely by the first two choices.

The class sum comparison and local unitary orders establish the elementary links between the first two structures. The third, and its compatibility with the chosen theta models, is the substance of later chapters. Even given the mass formula, an arbitrary convention for the archimedean formal degree would change the quantitative comparison, so that the additional local calculation cannot be omitted.

Concordance. The assertions now reconstructed in this chapter from the named classical references are the Weil record `Parameters.instWeilDataE`; the conjugation-compatible `nonempty_splittingChar`; the finite ideal and uniformizer relations of `UpsilonValues`; the decomposition and chosen-place calculations of `WeilDataOfE`; the integral orthonormal frame and finite group calculations used by the level; the norm-one class sum; and the mass ratio. The norm-one proof and the finite matrix arguments are independent of the mass formula. A full derivation of the Siegel–Weil formula, of the global Weil representation and of the specialized uniform theta-period identities is not asserted here. Those are different obligations, even when their objects have now been normalized by an independent construction.

Chapter 7

Heisenberg operators and the finite vacuum multiplier

This chapter supplies common background for the definite and indefinite theta models. In particular the finite multiplier recording the change of a vacuum lattice is not a free scalar; its sign and its absolute value occur in the upper-bound and period calculations. We give a direct lattice-theoretic computation of that multiplier and explain its relation to the Heisenberg model. We prove that the phase-normalized finite intertwiners form a genuine representation by combining a finite Gauss sum with a sign argument modulo four; an absolute-value calculation alone could not establish this assertion.

7.1 The saturating lattice at a finite place

Put $\widehat{\mathcal{O}} = \widehat{\mathcal{O}}_E$, and let $b \in \mathrm{GL}_n(\mathbb{A}_{E,f})$. The lattice

$$L_b = \widehat{\mathcal{O}}^{\,n} \cap b^{-1}\widehat{\mathcal{O}}^{\,n}$$

is compact open in the additive finite adeles, and of finite index in $\widehat{\mathcal{O}}^{\,n}$. It is the intersection of the old lattice with the inverse image of the new one. Its determinant ideal is the quantity that controls the finite part of the vacuum multiplier.

Lemma 7.1 (Saturating lattice matrices)**.** *There exist $u \in \mathrm{GL}_n(\mathbb{A}_{E,f})$ and $l, m \in M_n(\widehat{\mathcal{O}})$ such that*

$$u, bu \in M_n(\widehat{\mathcal{O}}), \qquad u^{-1} = l + mb.$$

Every such u satisfies $u\widehat{\mathcal{O}}^{\,n} = L_b$. At each finite place w, if b has elementary divisors $\pi_w^{a_1}, \dots, \pi_w^{a_n}$ relative to the standard lattice, then

$$\mathrm{ord}_w(\det u) = s_w(b) := \sum_{i=1}^{n} \max(0, -a_i). \tag{7.1}$$

Proof. At a finite place write $b = kDk'$ by the elementary divisor theorem over the discrete valuation ring, where k, k' are integral invertible matrices and $D = \mathrm{diag}(\pi^{a_i})$. Set $P = \mathrm{diag}(\pi^{\max(0,-a_i)})$ and $u = (k')^{-1}P$. Then u and $bu = kDP$ are integral. To make the claimed inverse identity explicit, let A, B be the diagonal matrices with $A_{ii} = 1$ when $a_i \geq 0$ and 0 otherwise, and $B = I - A$. Since $P^{-1} = A + BD$, one has

$$u^{-1} = P^{-1}k' = Ak' + Bk^{-1}b.$$

Thus take $l = Ak'$ and $m = Bk^{-1}$. At every place where b is integral invertible we may instead take $u = l = I$ and $m = 0$. The choices are therefore finite-adelic. If $x \in \widehat{\mathcal{O}}^n$ with bx integral, the inverse identity gives $u^{-1}x = lx + mbx$ integral. Conversely integrality of u and bu implies that uy lies in L_b for any integral y. This proves the lattice equality, and the explicit diagonal computation proves (7.1). It is independent of the choice of u since two bases of the same integral local lattice differ by a matrix of determinant a unit. □

The source uses the predicate `Sat` and the selected matrix `saturatingElem` in `Lib/Weil/Construction/FiniteMultiplier.lean`. The proof just given is the direct elementary divisor version of that construction. In particular there is a simple expression for every local determinant valuation, even when neither b nor its inverse is integral.

7.2 The multiplier and its transformation laws

Write $|a|_f = \prod_w q_w^{-\operatorname{ord}_w(a_w)}$ for the idele modulus, and let υ be Proposition 6.4's splitting character. In this section its restriction to finite ideles is denoted by the same symbol. Define

$$\chi_f(a) = \upsilon(a)|a|_f^{-1/2}, \qquad \Delta_f(b) = \chi_f(\det u).$$

By Lemma 7.1 and triviality of υ on finite units, this is well-defined. In particular

$$\Delta_f(b) = \prod_w (\upsilon(\pi_w)\sqrt{q_w})^{s_w(b)}, \qquad |\Delta_f(b)|^2 = \prod_w q_w^{s_w(b)}. \tag{7.2}$$

The product is finite, and the choice of local uniformizer is immaterial. This is the `glMultiplier` of the compact development.

Proposition 7.2 (Lattice multiplier identities)**.** *The multiplier has the following properties.*

1. *It is invariant under multiplication on either side by matrices in* $\mathrm{GL}_n(\widehat{\mathcal{O}})$.
2. $\Delta_f(b^{-1}) = \chi_f(\det b)\Delta_f(b)$.
3. $\Delta_f(b^*) = \overline{\Delta_f(b)}$, *where* $*$ *denotes conjugate transpose of the adelic matrix, including exchange of conjugate places.*
4. *It is multiplicative on block diagonal matrices. It is also multiplicative on a product* bb' *if at every place one of the two matrices is the identity.*
5. *For Haar measure satisfying* $\mathrm{vol}(\widehat{\mathcal{O}}^n) = 1$ *one has*

$$\mathrm{vol}(L_b) = |\Delta_f(b)|^{-2}, \qquad [\widehat{\mathcal{O}}^n : L_b] = |\Delta_f(b)|^2.$$

Proof. The first assertion is equivalent to invariance of the elementary divisors. For the second, the elementary divisors of b^{-1} are the negatives of those of b, in reverse order. Therefore

$$s_w(b^{-1}) - s_w(b) = \sum_i a_i = \operatorname{ord}_w(\det b).$$

Use (7.2) and the definition of χ_f. Transposition by itself does not change the elementary divisors, and conjugation sends the list at w to the list at $\bar{w}$. Since the two residue cardinalities agree and $\upsilon(\pi_{\bar{w}}) = \upsilon(\pi_w)^{-1}$ by (6.1), the third assertion follows. The fourth follows by addition of the local sums (7.1), respectively by separating the local factors in (7.2). Finally the index of $\pi^a\mathcal{O}_w$ in $\mathcal{O}_w$ is q_w^a for $a \geq 0$. The determinant sizes in Lemma 7.1 give the index formula, and Haar measure divides by that finite index on passage to the smaller lattice. □

Several identities used by the seesaw and the exchanged periods are immediate consequences. First, for g in the finite unitary group, $g^{-1} = g^*$ and $|\det g|_f = 1$, the latter because the determinant times its conjugate is one. Thus

$$\Delta_f(g^{-1}) = \overline{\Delta_f(g)} = \upsilon(\det g)\Delta_f(g). \tag{7.3}$$

If S is a set of places disjoint from its conjugate and g is trivial outside $S \cup \bar{S}$, let b equal its matrix at S and the identity elsewhere. Then $g = b(b^*)^{-1}$ and the supports are disjoint. Hence

$$\Delta_f(g) = |\Delta_f(b)|^2 \overline{\chi_f(\det b)}. \tag{7.4}$$

Indeed apply (2) and (3) and multiply. This identity explains why a single choice of one member of a split pair is sufficient in a local calculation: the other factor is determined by the inverse transpose and contributes the displayed conjugate factor. Replacing it by an independent local modulus would change the character.

As a one-dimensional example, suppose at a split pair g has values π^a and $\bar{\pi}^{-a}$, with $a \geq 0$. Then $s_w(g) = 0$ and $s_{\bar{w}}(g) = a$, so

$$\Delta_f(g) = \upsilon(\pi_w)^{-a} q_w^{a/2}, \qquad |\Delta_f(g)|^2 = q_w^a.$$

The same result follows from (7.4). At a conjugation-fixed place the sign in (7.2) is $(-1)^{s_w(b)}$; it is not a sign introduced by a choice of square root. It follows from the inert uniformizer value of the splitting character.

Equations (7.2)–(7.4) prove the identities `glMultiplier_inv`, `glMultiplier_star`, `glMultiplier_fromBlocks`, `latticeMeasure_inter_eq_glMultiplier`, `norm_finMultiplier_sq`, `upsilon_mul_finMultiplier`, and `finMultiplier_eq_of_supportedOn` in `Lib/Weil/Construction/FiniteMultiplier.lean`.

7.3 The finite vacuum and the Heisenberg cocycle

Let $V_f = \mathbb{A}_{E,f}^n$ with its standard Hermitian form $h(x, y) = x^* y$. It is convenient to work with the scaled lattice $L = \alpha\widehat{\mathcal{O}}^n$. Let ψ_f be the finite additive character of E, extended from the rational character by the trace. Set

$$\beta(x, y) = \psi_f(\tfrac{1}{2}\jmath h(y, x)).$$

We do not assume that the scalar 2 is a unit in the finite integers. It is invertible in the local fields, which is enough for the displayed definition. Bilinearity of the additive character makes β a bicharacter on the underlying additive group. Conjugation invariance of ψ_f and $\bar{\jmath} = -\jmath$ give

$$\beta(y, x) = \beta(x, y)^{-1}.$$

Moreover $\beta(x, x) = 1$, because $h(x, x)$ is fixed by conjugation and the trace of a purely imaginary element is zero. The latter trace assertion holds componentwise over each local base place, including split pairs.

The corresponding central extension has product

$$(x, z)(y, t) = (x + y, zt\beta(x, y)).$$

Associativity follows from the bicharacter identities, and the commutator is $\beta(x, y)^2$. The annihilator of L for this commutator pairing is L itself, by Lemma 6.6. Thus L is a maximal compact isotropic subgroup for the commutator, although β restricted to $L \times L$ may have values -1. This is the dyadic issue that must be respected when normalizing the vacuum.

Define on L the function

$$\lambda(x) = \psi_f(\tfrac{1}{2}c_0 \jmath h(x,x)), \qquad c_0 = (1-\jmath^{-1})/2.$$

It is a quadratic refinement on the lattice.

Lemma 7.3. *For $x, y \in L$ one has*

$$\lambda(x+y) = \lambda(x)\lambda(y)\beta(x,y).$$

If g is unitary and $x, gx \in L$, then $\lambda(gx) = \lambda(x)$.

Proof. Put $a = h(y,x)$ so that $h(x,y) = \bar{a}$. By conjugation invariance of the character,

$$\frac{\lambda(x+y)}{\lambda(x)\lambda(y)} = \psi_f(\tfrac{1}{2}c_0\jmath(a+\bar{a})) = \psi_f(\tfrac{1}{2}(c_0-\bar{c}_0)\jmath a).$$

Divide by $\beta(x,y) = \psi_f(\frac{1}{2}\jmath a)$ and use $c_0 + \bar{c}_0 = 1$. The ratio is $\psi_f(-\bar{c}_0\jmath a)$. Since $x, y \in L$, the element a belongs to $\alpha\bar{\alpha}\widehat{\mathcal{O}}_E$ and $\bar{c}_0$ is integral. The argument lies in the inverse different by (6.2), so its character is trivial. This proves the refinement. The second assertion is the equality of the Hermitian norms under g. □

The finite part of the source's `vacuumFn` in `Lib/Weil/Construction/AdelicHeisenberg.lean` is $\lambda\mathbf{1}_L$. This proves its essential lattice identity directly from the choice of the different idele and the trace-one element. In particular it would not suffice to use the indicator alone in a presentation of the Heisenberg bicharacter with the factor $1/2$.

7.4 The lattice model and the absolute vacuum coefficient

We describe the finite part in a standard induced model, rather than assume an arbitrary Hilbert-space normalization. Let $\mathcal{H}_L$ be the Hilbert space of functions f on V_f satisfying

$$f(y+\ell) = \lambda(\ell)\beta(y,\ell)f(y) \qquad (\ell \in L), \tag{7.5}$$

with square-summable values on the countable discrete quotient V_f/L. Use counting measure on that quotient. The absolute value of an equivariant function is independent of the representative by the unit modulus of all factors; hence the quotient norm is well-defined. Consistency of two translations in (7.5) is precisely the identity in Lemma 7.3, together with $\beta(\ell_1,\ell_2)^2 = 1$ on L. Define

$$(\rho(x)f)(y) = \beta(x,y)f(y-x).$$

For $\ell \in L$ the bicharacter terms show that this formula preserves (7.5), and it is an isometry because translation permutes the quotient. A direct calculation gives $\rho(x)\rho(z) = \beta(x,z)\rho(x+z)$, the cocycle convention of the source's `Heisenberg.Model`.

On L, we also have $\lambda(x)^2 = 1$, by integrality of c_0, self-duality and the definition of the finite character. The function $e = \lambda\mathbf{1}_L$ satisfies (7.5) by the refinement and is of norm one. For $\ell \in L$, $\rho(\ell)e = \lambda(\ell)e$ (use the refinement on the common support); if $x \notin L$ then $\rho(x)e$ has support in a different coset and is orthogonal to e. The vectors $\rho(x)e$ indexed by representatives of the quotient form an orthonormal basis, up to the irrelevant phases of those representatives. An element in $\mathcal{H}_L$ transforming by λ under all $\rho(\ell)$ has support in L: at any other coset, the equivariance formula and the action formula force the commutator character $\beta(\ell,y)^2$ to be identically one, hence $y \in L$. Its values on L are determined by its value at zero, so the space of such vectors is the single line $\mathbb{C}e$.

Let g be in the finite unitary group. It acts symplectically on V_f. Suppose U_g is a unitary intertwiner implementing g on the Heisenberg operators,

$$U_g \rho(x) U_g^{-1} = \rho(gx).$$

Existence and uniqueness up to phase follow from the standard Stone–von Neumann theorem, or directly by comparing the induced models of the two maximal compact isotropic subgroups L and gL. Here we compute the exact modulus of its vacuum coefficient without making a phase choice.

There is a direct identification with the integrable reproducing-kernel model in `Lib/Weil/Construction/Model.lean`. With $e = \lambda \mathbf{1}_L$ and $\mathrm{vol}(L) = 1$, its defining equation is

$$f(x) = \int_{V_f} f(y)\beta(y,x)e(x-y)\,dy.$$

If f satisfies (7.5), set $y = x - \ell$ on the support of the integrand. Then the lattice equivariance of f and the identities $\lambda(-\ell) = \lambda(\ell)$, $\lambda(\ell)^2 = 1$, and $\beta(\ell, x) = \beta(x,\ell)^{-1}$ show that the integrand is exactly $f(x)$. The integral has volume one. Conversely, replace x by $x + \ell$ in the reproducing equation; use the lattice equivariance of $e(x-y)$ and the bicharacter identity $\beta(y,\ell)\beta(x-y,\ell) = \beta(x,\ell)$. This yields (7.5). Thus the source's integrable finite model and the integrable lattice-equivariant functions agree, and the finitely supported coset functions are dense in their Hilbert completion. Since both implement the same group action on the Heisenberg operators, equality of the normalized vacuum coefficient identifies the finite intertwiners, not only their absolute coefficients.

Proposition 7.4. *With e as above,*

$$|\langle e, U_g e\rangle|^2 = [L : L \cap gL]^{-1} = |\Delta_f(g)|^{-2}.$$

Proof. On L the operators $\lambda(\ell)^{-1}\rho(\ell)$ form a genuine unitary representation, since the refinement cancels the cocycle on L. Integrating them against the probability Haar measure of L gives the orthogonal projection P onto the vacuum line just identified. Taking the matrix coefficient of P at $U_g e$, one obtains

$$\begin{aligned} |\langle e, U_g e\rangle|^2 &= \langle U_g e, P U_g e\rangle \\ &= \int_L \lambda(\ell)^{-1}\langle e, \rho(g^{-1}\ell)e\rangle\,d\ell. \end{aligned}$$

Here the inner product is linear in the second variable. The integrand vanishes off $L \cap gL$ by the coset support of the translated vacuum. On that intersection it equals $\lambda(\ell)^{-1}\lambda(g^{-1}\ell) = 1$ by the norm invariance in Lemma 7.3. The integral is therefore the reciprocal lattice index. Scaling by α commutes with g, and Proposition 7.2 applied to g^{-1} gives the last equality, using (7.3). □

The phase-normalized operator U_g is specified by $\langle e, U_g e\rangle = 1/\overline{\Delta_f(g)}$. Its inverse has vacuum coefficient $1/\Delta_f(g)$, so (7.3) and uniqueness up to a scalar give $U_g^{-1} = U_{g^{-1}}$. Uniqueness of an intertwiner still implies only $U_g U_h = c(g,h)U_{gh}$ for some scalar of absolute value one. We next prove that scalar equals one.

7.5 The Gauss-sum phase and the finite splitting

Extend the expression for λ to all of V_f (the refinement identity was proved only on L). Write $f_g = U_g e$ and $c_g = 1/\overline{\Delta_f(g)}$. For $a \in gL$, intertwining gives $\rho(a)f_g = \lambda(g^{-1}a)f_g = \lambda(a)f_g$,

using preservation of the Hermitian norm. At zero the eigenrelation yields $f_g(-a) = c_g\lambda(a)$; since the expression for λ is even, $f_g(a) = c_g\lambda(a)$. At any x in the support of f_g, comparison with the lattice-model eigenrelation for $a \in L \cap gL$ forces $\beta(a,x)^2 = 1$. Both lattices are self-dual for this commutator pairing. Their intersection has annihilator equal to their sum: the annihilator of a sum is the intersection of annihilators, and the double-annihilator identity for compact open lattices follows from perfect trace duality on finite lattice quotients. Hence f_g is supported in $L + gL$.

Write $x = a + b$ with $a \in gL$ and $b \in L$. By (7.5),

$$f_g(x) = c_g\lambda(a)\lambda(b)\beta(a,b) = c_g\lambda(x)\psi_f(\bar{c}_0\jmath h(b,a)). \tag{7.6}$$

For the second equality, by conjugation invariance of ψ_f, the ratio $\lambda(x)/(\lambda(a)\lambda(b))$ is $\psi_f((c_0 - \bar{c}_0)\jmath h(b,a)/2)$ and $1 - c_0 + \bar{c}_0 = 2\bar{c}_0$. This does not require $a \in L$.

The common support in the inner product of $f_{g^{-1}}$ and f_h gives the finite $\widehat{\mathcal{O}}_E$-module

$$Q = ((L + g^{-1}L) \cap (L + hL))/L.$$

Formula (7.6) and cancellation of $\lambda(x)$ show that

$$\overline{f_{g^{-1}}(x)}f_h(x) = \overline{c_{g^{-1}}}c_h F(x \bmod L)$$

where F is well-defined by common lattice equivariance, $F(0) = 1$, and all its values are roots of unity. More precisely F is the additive character of the difference of the two cross terms in (7.6). Scaling the two decompositions of x by $t \in \widehat{\mathcal{O}}_E$ scales this difference by $\bar{t}t$, so

$$\begin{aligned} F(tq) &= F(q) && \text{if } \bar{t}t = 1, \\ F(tq) &= 1 && \text{if } \bar{t}t = 0, \\ F(tq) &= \overline{F(q)} && \text{if } \bar{t}t = -1. \end{aligned} \tag{7.7}$$

Compute the normalized vacuum coefficient:

$$\begin{aligned} \frac{c(g,h)}{\Delta_f(gh)} &= \langle e, U_g U_h e\rangle = \langle U_{g^{-1}}e, U_h e\rangle \\ &= \frac{\sum_{q\in Q} F(q)}{\overline{\Delta_f(g)}\,\overline{\Delta_f(h)}}. \end{aligned} \tag{7.8}$$

This is the finite Gauss-sum relation. All lattices, modules and additive characters factor over finite places of F, treating the two components of a split pair together. Only finitely many factors are nontrivial. It suffices to show that the corresponding local cocycle is identically one.

Lemma 7.5 (The orbit sign test)**.** *Suppose a finite abelian group Q has no nonzero element of order two, and an automorphism i has $i^2 = -1$. Let $F : Q \to \mathbb{C}^\times$ have root-of-unity values, with $F(0) = 1$, $F(iq) = F(q)$, and $S = \sum_q F(q)$ real. If an integer $r \equiv 1 \pmod 4$ is nonzero and*

$$cr^a = Sr^b, \qquad a, b \geq 0, \quad |c| = 1,$$

then $c = 1$.

Proof. Every nonzero orbit of i has size four: if its size divided two, then $q = -q$, hence $q = 0$. Thus $S = 1 + 4A$ for an algebraic integer A, and the assumed equality gives $c = 1$ or -1. If $c = -1$, it implies $(r^a + r^b)/4 = -r^b A$, a rational algebraic integer and hence an integer. This is impossible since the numerator is 2 modulo 4. □

Theorem 7.6 (The finite normalized action)**.** *For the data of Chapter 6, the phase-normalized Heisenberg intertwiners satisfy $U_g U_h = U_{gh}$.*

Proof. At a split place write h as $(a, (a^t)^{-1})$ over the two copies of F_v. First suppose a^{-1} is integral. Then the second component of hL is contained in L, so Q is supported on the first side. With t the projection idempotent of that side one has $tq = q$ and $\bar{t}t = 0$. Equation (7.7) implies $F(q) = 1$ for every q, giving a positive integer sum.

For a split unitary element with first component b, the elementary divisor lists on the two sides are opposite. The phase in (7.2) is $\upsilon(\pi_w)^{-\operatorname{ord}_w(\det b)}$, a multiplicative character of b. Hence the multiplier ratio in (7.8) is positive real. The unit-modulus scalar $c(g, h)$ is therefore one for every g when a^{-1} is integral. The elements h with $c(g, h) = 1$ for all g form a subgroup by the cocycle identity $c(g, h)c(gh, k) = c(g, hk)c(h, k)$ (also apply it with $k = h^{-1}$). Integral unitary elements and diagonal elements with inverse integral lie in that subgroup. By the elementary divisor theorem they and their inverses generate the split group.

At a conjugation-fixed place w the quadratic local extension is unramified and $q_w = r^2$ with $r = q_v$. The residue characteristic is odd: a fixed place of E would restrict to a fixed place of K, and 2 splits in K. The reduction of $i = \sqrt{-1}$ has order four and its conjugate is its negative, so the nontrivial residue automorphism $x \mapsto x^r$ gives $r \equiv -1 \pmod 4$. The inert uniformizer value of υ is -1, whence the local multiplier (7.2) is an integer power of $-r$.

The local group Q has odd prime-power order. Multiplication by i squares to negation and preserves F by (7.7). By Lemma 6.1 choose a unit of norm -1; multiplication by it permutes Q and conjugates every value of F. Thus the sum of F is real. Equation (7.8) is of the form in Lemma 7.5, with base $-r \equiv 1 \pmod 4$ and nonnegative exponents $s_w(g) + s_w(h)$ and $s_w(gh)$. This proves $c(g, h) = 1$, completing both local cases and hence the global finite proof. □

The sign lemma and finite-lattice sum supply the mathematical content of `eq_one_of_orbits` in `Lib/Weil/Construction/GaussSumSign.lean` and `stdLattice.mem_cocycleRadical` in `Lib/Weil/Construction/ThetaFunctional.lean` for the specified fields, in the lattice model. The proof uses the standard irreducible Heisenberg intertwiner but computes its finite normalization independently. An absolute value of a Gauss sum would not suffice: the order-four orbits determine the final inert sign.

7.6 The real harmonic model and its basic normalization

The real part of the same oscillator has a particularly convenient holomorphic model. For a positive Hermitian space of complex dimension n, put on holomorphic polynomials the Fischer–Bargmann pairing

$$\langle P, Q \rangle = \pi^{-n} \int_{\mathbb{C}^n} \overline{P(z)} Q(z) e^{-|z|^2} \, d^{2n} z.$$

The constant polynomial 1 has norm one. For multi-indices a, b the monomials satisfy

$$\langle z^a, z^b \rangle = \delta_{a,b} \prod_j a_j!,$$

because angular integration kills distinct exponents and the radial integral at each coordinate is $\int_0^\infty t^{a_j} e^{-t} dt = a_j!$. Thus, in the Hilbert completion, $z^a / \sqrt{a!}$ is an orthonormal basis. The reproducing kernel is

$$K(z, w) = \sum_a \frac{z^a \bar{w}^a}{a!} = e^{z \cdot \bar{w}},$$

with locally uniform convergence by the exponential series. For a unitary matrix u, the action on polynomials by $P(z) \mapsto P(u^{-1}z)$ preserves the pairing, because both the Gaussian and Lebesgue measure are invariant. Tensoring with a determinant character of modulus one likewise preserves it. These elementary facts underlie the integral Fischer pairing and the compact harmonic coefficient formulas in Chapter 5.

The scale $s_\tau = (2\pi \operatorname{Im} \tau(\jmath))^{-1/2}$ fixed in Chapter 6.3 gives the norm-one vacuum wave function a Gaussian of the form $e^{-|z|^2/2}$ in the original coefficient variable, since $\pi \operatorname{Im} \tau(\jmath)|s_\tau z|^2 = |z|^2/2$. Its square therefore has the Bargmann density $e^{-|z|^2}$. For an indefinite unitary group the holomorphic and antiholomorphic variables are chosen according to the signs, and mixing them is the noncompact oscillator action. The compact calculation alone does not determine the noncompact multiplier or the full Flensted–Jensen coefficient. Those are studied in Chapter 16 in the normalization required by the filling system.

7.7 The rational theta functional in a polarized model

Finally, to explain the relationship between the Heisenberg model and adelic theta sums, consider a symplectic F-space with a rational Lagrangian splitting $X \oplus X^\vee$. For a Schwartz–Bruhat function f on $X(\mathbb{A}_F)$ define

$$\theta(f) = \sum_{x \in X(F)} f(x).$$

The sum is absolutely convergent: the finite part of a Schwartz–Bruhat function is compactly supported and locally constant, reducing the possible rational points to a finite union of lattices in the real vector space; the real part and all its derivatives decay faster than any polynomial. The Poisson identity

$$\sum_{x \in X(F)} f(x) = \sum_{x \in X^\vee(F)} \widehat{f}(x)$$

uses the self-dual Haar measure for the standard rational additive character. One proof is to reduce to the ordinary real lattice Poisson formula by decomposition into finite congruence cosets, followed by the finite character orthogonality of the dual lattice. The inverse different is exactly the trace-dual lattice used in that proof.

The three basic Schrödinger operations are changes of coordinates, multiplication by quadratic characters, and Fourier transform. For a rational invertible coordinate change $A \in \mathrm{GL}(X)(F)$, the change of variables factor in the oscillator is $|\det A|_{\mathbb{A}_F}^{1/2} = 1$ by the product formula, and the rational summation indices are permuted. For a rational symmetric form B on X, the operation multiplying by the full additive character of $\frac{1}{2}B(x,x)$ fixes each rational summand, because the full additive character is trivial on F. Fourier transform preserves the sum by Poisson. These operations generate the rational symplectic action in its standard global splitting, which follows from the elementary symplectic row reduction using Lagrangian bases. This is the usual proof of rational theta invariance in a Schrödinger model. When a unitary dual pair is embedded in this symplectic group and equipped with the same rational splitting, the calculation applies to its rational points too.

There are two comparisons one must still supply in the nonpolar presentation. First, the compact source constructs the theta functional from invariance under rational Heisenberg translations and normalizes it by a vacuum coefficient; existence and one-dimensionality of its continuous invariant functional are established in `Lib/Weil/Construction/ThetaFunctiona l.lean`. Second, it proves that the exact phase-normalized Weil operators fix this functional, using a doubling argument, `Space.theta_omega_ratEmbed` in `Lib/Weil/DualPair/PairSp`

`litting.lean`. The classical polarized calculation above proves the rational invariance in the standard model and explains the role of Poisson and product formula. To deduce the assertion for the precise nonpolar normalization one must verify the unitary intertwining identification including the real phase and theta-functional normalization. The finite phase has now been calculated in Theorem 7.6, but the global identification is not a consequence of that calculation alone. Its normalized mixed-model formulas are the input used in Chapter 12, where the actual orbit reduction and degenerate and nonsingular finite lattice calculations are developed.

Classical background and precise scope. For the existence and splitting of the adelic Weil representation, the classical sources are Weil's construction and Kudla's unitary dual-pair splittings ([15, 16]). The proofs in this chapter establish the finite saturation, all multiplier and index identities, the lattice quadratic refinement, absolute vacuum overlap, and finite normalized group action, and recall the usual polarized theta invariance. The remaining real-phase and theta-model-identification assertions are kept separate. They cannot be removed by proving only that the operators are unitary.

Part II

Upper bounds: CM jets, theta chains and Hecke repair

Chapter 8

The CM differential, theta normalization and integral frame

The purpose of this chapter is to establish the arithmetic preliminaries of the local jet theorem. A theta kernel is not an arbitrary analytic function on a complex vector space. Its translation law, the rational structures of the Hermitian spaces, and the lattice in the invariant differential have to be compatible. If one normalizes an individual complex theta function independently at every prime, the product formula used later in the proof is lost. Here we exhibit the single normalization and explain why its finite and infinite sizes agree with the requirements of the jet estimate. We give proofs for the local theta identities, the period scaling conditional only on classical CM and limit theorems and the normalization of the Hecke character, and the arithmetic tight-frame construction. At the end we explain the remaining Galois comparison that is needed to pass between finite Weil representations. The discussion uses d, K, E, N_1, Φ, f as in the admissible data. Write $N_1 = N$ temporarily within the frame construction; this N is *not* the auxiliary level prime of the proof.

8.1 An elliptic curve normalized at a point of order two

Let d be the prime discriminant parameter, $d \equiv 7 \pmod 8$, and set $c_0 = (1 + \sqrt{-d})/2$, so that $\mathcal{O}_K = \mathbb{Z}[c_0]$ and $c_0(1 - c_0) = (d + 1)/4$. In particular 2 splits in K. If Λ is homothetic to a fractional ideal of K, the quotient $\mathbb{C}/\Lambda$ has its natural $\mathcal{O}_K$ action. The two idempotents in $\mathcal{O}_K/(2) \simeq \mathbb{F}_2 \times \mathbb{F}_2$ act on its 2-torsion. There is a distinguished nonzero 2-torsion point t whose two projections are nonzero. Thus $2t \in \Lambda$, but neither $c_0 t$ nor $(1 - c_0)t$ belongs to Λ. In the source, such a t is part of the datum `CMLattice`.

Denote the Weierstrass functions of Λ by $\wp$, $\wp'$. Since t is a nonzero point of order two, $\wp'(t) = 0$ and $\wp''(t) \neq 0$. The second fact follows either from smoothness of the cubic or from the differential equation $\wp'^2 = 4\wp^3 - g_2\wp - g_3$: the cubic has distinct roots. On scaling Λ and t by u, the value of $\wp''(t)$ is multiplied by u^{-4}. We may therefore arrange $\wp''(t) = 2$. This fixes the differential up to a fourth root of unity. Choose a_1 with $a_1^2 = 12\wp(t)$. The shifted functions

$$X(z) = \wp(z + t) - \wp(t), \qquad Y(z) = \tfrac{1}{2}\big(\wp'(z + t) - a_1 X(z)\big)$$

give the affine equation

$$Y^2 + a_1 XY = X^3 + X. \tag{8.1}$$

Indeed $g_2 = 12\wp(t)^2 - 4$, so shifting the cubic by its root $\wp(t)$ gives $(2Y + a_1X)^2 = 4X^3 + a_1^2X^2 + 4X$. Its Weierstrass discriminant is $\Delta = a_1^4 - 64$ and $c_4 = a_1^4 - 48$. This is checked from $b_2 = a_1^2$, $b_4 = 2$, $b_6 = 0$, and $b_8 = -1$. The point $z = 0$ corresponds to $(0, 0)$ and the

point at infinity to $z \equiv t \pmod{\Lambda}$; z agrees with the invariant differential coordinate after translation by the point of order two. In particular $X = z^2 + O(z^4)$ and $Y = z + O(z^2)$. Differentiating gives

$$X' = 2Y + a_1X, \qquad Y' = 1 + 3X^2 - a_1Y, \qquad X'' = 6X^2 + a_1^2X + 2. \tag{8.2}$$

Proposition 8.1 (Unit discriminant in the chosen normalization)**.** *For every prime p and embedding of coefficients into an algebraically closed p-adic field, the normalized cubic* (8.1) *satisfies $|a_1|_p \le 1$ and $|\Delta|_p = 1$.*

Proof. We use two classical elliptic-curve facts: a CM j-invariant is an algebraic integer, and integrality of j implies potential good reduction (Silverman, *Advanced Topics in the Arithmetic of Elliptic Curves*, II, §6, and *The Arithmetic of Elliptic Curves*, VII, Proposition 5.5). The normalized a_1 is algebraic, since $j = (a_1^4 - 48)^3/(a_1^4 - 64)$. Thus after a finite extension of the local field of definition we have a good Weierstrass model. All the endomorphisms and the point t can also be defined there, enlarging the extension if necessary. If $p \neq 2$, the nonzero point of order two has nonzero specialization, by finite etaleness of the 2-torsion. For $p = 2$ we use the distinguished choice of t. The complementary idempotents of $\mathcal{O}_K/(2)$ decompose the finite flat 2-torsion group scheme in the good model into two factors of rank two. The endomorphisms extend to the good model, and the induced idempotents act on its 2-torsion including the special fibre. On the one-dimensional tangent space of the special-fibre elliptic curve one of these two projectors is zero, the other the identity. The corresponding factor of rank two with zero tangent is etale. The projection of t onto this factor is a nonzero generic-fibre point, and specializes nontrivially. In particular t itself specializes nontrivially. This argument is the reason that choosing the point outside both CM kernels is essential: an arbitrary 2-torsion point on an ordinary curve might specialize to the identity.

We may work over the algebraically closed valuation field and translate the coordinates of our good Weierstrass model so that t is $(0, 0)$. Its affine coordinates are integral, since the reduction of the point is different from the point at infinity. Thus a good equation has the shape

$$y^2 + A_1xy = x^3 + A_2x^2 + A_4x$$

with integral coefficients; the vanishing of the constant term and of the term linear in y follow from the point and its order two. At $(0, 0)$ the partial derivative with respect to y vanishes, and the partial derivative with respect to x is $-A_4$. Smoothness of the integral model therefore implies that A_4 is a unit. A change $y = y' + sx$ eliminates the term in x^2 if $s^2 + A_1s = A_2$. All roots of this monic integral equation are integral. Scaling by a fourth root of A_4 then gives a good equation $y'^2 + A_1'x'y' = x'^3 + x'$. These operations are integral changes with unit leading scale. Finally the two normal forms of equation (8.1) over the algebraically closed local field are isomorphic preserving infinity and the distinguished point $(0, 0)$. A general change of Weierstrass coordinates with these two conditions has $x = u^2x'$ and $y = u^3y' + su^2x'$. Because the two coefficients of the term linear in x are both 1, one has $u^4 = 1$. Consequently the two discriminants differ by a unit, and $|a_1|_p \le 1$ follows from $a_1^4 = 64 + \Delta$. □

This proves the normalized integrality statement `CMLattice.norm_disc_eq_one` and `CMLattice.norm_`a_1`_le_one` in `Lib/CM/LocalHeight/IntegralParameter.lean` by a classical-model route instead of the formal source's explicit calculation with a chain of two-isogenies. We do not need the stronger assertion of reduction over any specified field of definition: the formal argument works after transferring coefficients to an algebraically closed p-adic field. The elementary good-model normalization also applies to all conjugate curves, since they carry the conjugate CM action and distinguished point.

8.2 Theta sections and their integral size

Put $A = \mathrm{area}(\Lambda)$. We must use the sigma function in the Hermitian normalization, not the unmodified classical Weierstrass product. If σ_W is the latter, with quasi-period η_l at $l \in \Lambda$, the usual quasi-period relation on an oriented basis of the lattice gives

$$\frac{\pi\bar{l}/A - \eta_l}{l} = c \qquad (l \in \Lambda \setminus \{0\})$$

for a single complex scalar c (this is the Legendre relation). Thus $\sigma(z) = e^{cz^2/2}\sigma_W(z)$ is an odd theta section with derivative one at zero, transforming at a period l by the sigma sign ($+1$ if $l \in 2\Lambda$, -1 otherwise) times $\exp(\pi(z\bar{l} + |l|^2/2)/A)$. This is precisely the choice of `PeriodPair.s` in `Lib/CM/Sigma/SigmaFunction.lean`. The classical addition and doubling identities are unchanged by multiplication by the indicated quadratic exponential; the exponential factors on their two sides agree. For a section of degree h use the Hermitian form $h\pi z\overline{w}/A$, and define the twisted translation

$$(U_y^h F)(z) = \exp(-h\pi(z\bar{y} + |y|^2/2)/A)F(z+y). \tag{8.3}$$

Let $\vartheta = U_t^1\sigma/(U_t^1\sigma)(0)$. It is even, satisfies $\vartheta(0) = 1$, and has its simple zeros at $t + \Lambda$. These facts follow from the sigma transformation law and $2t \in \Lambda$. The addition formula for the sigma function gives, off the indicated zeros and then by continuation,

$$\vartheta(z+w)\vartheta(z-w) = \vartheta(z)^2(\vartheta(w)^2 - X(z)\sigma(w)^2). \tag{8.4}$$

Another consequence is the identity $(\sigma(z)/\vartheta(z))^2 = X(z)$. To see this directly, use the addition identity at $w = t$ and the relation $2t \in \Lambda$ to move the translated arguments back by a period. The proportionality constant is one, as both sides begin with z^2. Write $G(z) = Y(z)\vartheta(z)^3$; the apparent pole at $t + \Lambda$ cancels because Y has a pole of order three there. Thus G is an entire section of degree three. The identities used below are

$$\vartheta(2z) = \vartheta(z)^4 - \sigma(z)^4, \tag{8.5}$$

$$\sigma(2z)^2 = \sigma(z)^2\vartheta(z)^2(4\vartheta(z)^4 + a_1^2\sigma(z)^2\vartheta(z)^2 + 4\sigma(z)^4), \tag{8.6}$$

$$G^2 + a_1\sigma^2\vartheta G = \sigma^6 + \sigma^2\vartheta^4. \tag{8.7}$$

For the first, specialize (8.4) to $w = z$; for the second, use $\sigma(2z)^2 = \wp'(z)^2\sigma(z)^8$ together with $\vartheta(z)^2 = (\wp(z) - \wp(t))\sigma(z)^2$ and the shifted cubic. The last is (8.1) multiplied by ϑ^6. As both sides are entire, the identities are valid even when a denominator used in their derivation vanishes. The sigma addition/translation identities are classical (Weil, *Elliptic Functions according to Eisenstein and Kronecker*, Springer, 1976); the exact formulas are also the declarations `theta_add_mul_theta_sub`, `sigma_two_mul_sq`, `phiThetaCube_sq_add`, `θ_two_mul` in `Lib/CM/Sigma/CurveTheta.lean`.

Write $H_p(z) = \max(|\sigma(z)|_p, |\vartheta(z)|_p)$ after applying one fixed embedding ι to the complex values. It might seem strange to use a p-adic norm of the value of a complex analytic function at every complex argument; no continuity of this operation is asserted. The point is that the following conclusions are algebraic identities between their values.

Lemma 8.2 (Exact doubling of the local size)**.** *One has $H_p(z) > 0$, $|G(z)|_p \le H_p(z)^3$ and $H_p(2z) = H_p(z)^4$ for all $z \in \mathbb{C}$.*

Proof. The zeros of σ and ϑ are disjoint, so the maximum is positive. Divide (8.7) by the sixth power of a value of modulus $H_p(z)$. The resulting equation is monic in the normalized

G, with integral coefficients by Proposition 8.1. It follows that $|G(z)|_p \le H_p(z)^3$. The two duplication formulas immediately give $H_p(2z) \le H_p(z)^4$. If the inequality were strict, choose a scalar C of modulus $H_p(z)$ and put $u = \sigma(z)/C$, $v = \vartheta(z)/C$, viewed through ι. Their maximum norm is one. Both right sides of (8.5) and (8.6), divided by C^4 and C^8 respectively, would reduce to zero in the residue field. The first gives $u^4 = v^4$; hence both are units. The second implies $8 + a_1^2(u/v)^2 = 0$ in that field. Squaring and using $(u/v)^4 = 1$ yields $a_1^4 = 64$ there. This contradicts the unit discriminant, even when $p = 2$ (then a_1 reduces to a unit). The opposite inequality follows. □

It is useful to compare this with the usual local canonical height. A generic Weierstrass model leads to a bounded-error doubling formula and to a limiting procedure. Here the carefully chosen model removes the error at every prime, exactly in the same normalization. For a section represented by a monomial $\vartheta^a\sigma^{2b}G^c$ of degree $h = a + 2b + 3c$, the bound is $|\vartheta^a\sigma^{2b}G^c|_p \le H_p^h$; this is why the distinguished coordinates use two units of degree for X and three for Y.

Lemma 8.3 (Spanning by monomial sections)**.** *The degree-h theta sections for the multiplier of ϑ are spanned by*

$$\vartheta^a\sigma^{2b}G^c, \qquad a, b, c \ge 0, \quad c \le 1, \quad a + 2b + 3c = h.$$

Proof. Divide a section by ϑ^h. The quotient is a meromorphic function on $\mathbb{C}/\Lambda$ with no pole except at t, of order at most h. In the coordinates (8.1), the punctured elliptic curve is affine, and its ring of regular functions is $\mathbb{C}[X, Y]/(Y^2 + a_1XY - X^3 - X)$. Every element is uniquely $P(X) + YQ(X)$. Alternatively this uniqueness and spanning follow by subtracting the leading Laurent coefficient at the pole successively, since X, Y have poles of orders two and three. The monomials X^b, YX^b have distinct pole orders $2b, 3 + 2b$, so the pole bound means all appearing monomials have order at most h. Multiplying by ϑ^h and using $X\vartheta^2 = \sigma^2$, $Y\vartheta^3 = G$ proves the assertion. This includes $h = 0, 1$, when there are only the constant and ϑ sections respectively. □

The analytic and algebraic arguments are interchangeable here. In Lean the spanning assertion is `CMLattice.sections_le_span_monomialSection`, in `Lib/CM/Sigma/PadicSeriesDefs.lean`, and is proved by dividing entire sections rather than invoking the affine description of the elliptic curve. We will use both the pole description and its degree bookkeeping.

8.3 Isogenies, their multipliers, and Tate's inequality

If $a\Lambda \subseteq \Lambda'$ is an isogeny of degree n, then $|a|^2A = nA'$. The principal polarizations on these curves are represented by the Hermitian forms above. Pullback of the second form is n times the first form. The multipliers in the source are chosen to be the canonical ones defined using c_0; on periods their formula is

$$\chi(l) = \exp(\pi i E(\bar{c}_0 l, l)),$$

where E is the integral Riemann form. The formula can be checked explicitly on $\Lambda/2\Lambda$. In a basis e_1, e_2 of its two split summands the symplectic pairing of the two basis vectors is one modulo two. The sigma sign of a class $ue_1 + ve_2$ is $(-1)^{u+v+uv}$, being one just for the zero class. Translation by the distinguished t introduces the sign $(-1)^{u+v}$, because its class after multiplication by two is $e_1 + e_2$. Their product is $(-1)^{uv}$. Projection by $\bar{c}_0$ onto one summand

followed by pairing with the original vector has value uv modulo two, giving the claimed formula. Because multiplication by a is complex-linear and commutes with $\bar{c}_0$, the pullback of this multiplier is χ^n. This is the elementary consequence `comp_mul_mem_sections` of `multiplier_eq_canonicalMultiplier`, in `Lib/CM/Sigma/CurveTheta.lean`. In particular it is important not to choose independently a different two-torsion point on the target curve. If F is a degree-h section on the target, $F(az)$ is a degree-hn section on the source.

Proposition 8.4 (Tate's local inequality)**.** *Let a be a CM isogeny as above and y a torsion point modulo Λ'. If $|F(w)|_p \le H'_p(w)^h$ for all w, then*

$$|(U_y^h F)(az)|_p \le H_p(z)^{hn} \quad (z \in \mathbb{C}). \tag{8.8}$$

Proof. All values are read through the same embedding. Put $A_y(w) = \exp(-\pi(w\bar{y}+|y|^2/2)/A')$ and define

$$R_a(z) = |A_y(az)|_p \frac{H'_p(az+y)}{H_p(z)^n}.$$

The function R_a is bounded. Indeed, choose $q > 0$ with $qy \in \Lambda'$. Raising each of the two functions $A_y(w)\sigma'(w+y)$ and $A_y(w)\vartheta'(w+y)$ to the power $2q$ (where primes here label the target curve's sigma and theta) makes them degree-$2q$ sections of the original multiplier: twisted translation by a point of order dividing q introduces a character of finite order dividing $2q$ on the period lattice. Pulling back by a yields degree-$2qn$ sections on the first curve. By Lemma 8.3 and the height bound for each monomial, each such section is bounded by a constant times H_p^{2qn}. Taking a maximum and a $2q$-th root proves the asserted boundedness. This uses finite-dimensional linear algebra, not continuity in the p-adic topology of the complex variable.

The identity $A_y(2w)A_y(2w+y) = A_y(w)^4$ follows by adding the exponents. Lemma 8.2 gives

$$R_a(z)^4 = R_a(2z)R_1(2az+y),$$

where R_1 is the same expression for the identity on the target curve. Let $M_a = \sup_z R_a(z)$; it is a finite positive real number, as is M_1. Then $M_1^4 \le M_1^2$ and hence $M_1 \le 1$. Likewise $M_a^4 \le M_a M_1$, so $M_a \le 1$. Raising $R_a(z) \le 1$ to the power h and using the hypothesis on F proves (8.8). □

This is the mechanism of `CMLattice.norm_translate_le_naiveHeight_pow` in `Lib/CM/LocalHeight/IntegralParameter.lean`. The proof only needs the boundedness of each ratio and the exact duplication identity. It is stronger than a bound allowing an unspecified factor depending on the order of torsion, which could not be paid uniformly when selecting tight frames at different primes.

8.4 The integral formal parameter and the Gauss principle

We use the translated coordinate $Y = \phi$ as local parameter at $z = 0$. It is important that $\mathfrak{o}[[\phi]]$ need not equal the integral series ring in the invariant differential z. In what follows $\mathfrak{o}$ is the valuation ring of the algebraically closed p-adic field, and all series are formal. Substitution $T \mapsto \phi(z)$ gives a subring $\mathfrak{R}_\phi \subset \mathbb{C}_p[[z]]$; because the linear term of ϕ is one, it is isomorphic to $\mathfrak{o}[[T]]$. From the equation

$$X = Y^2 + a_1 XY - X^3$$

there is a unique formal solution $X = Y^2 + \cdots$ over $\mathfrak{o}$ by coefficient recursion. Hence $X \in \mathfrak{R}_\phi$. Equation (8.2) shows that differentiation in the invariant differential preserves the ring: if

$F(T) \in \mathfrak{o}[[T]]$, then $(F(\phi(z)))' = F'(\phi(z))(1+3X^2-a_1\phi(z))$. All factors are in the ring. This gives the first part of Proposition 9.3.

Lemma 8.5 (Gauss principle for finitely many polynomials)**.** *For nonzero polynomials $P_1, \dots, P_j$ over an algebraically closed valued field with infinite residue field, one can find x integral such that $|P_i(x)|_p$ equals the maximum of the absolute values of the coefficients of P_i for every i simultaneously.*

Proof. Divide each P_i by a coefficient of maximal absolute value and reduce the polynomials. None of the resulting polynomials in the residue field is zero. Choose a residue point outside the finitely many roots of their product, and lift it to the valuation ring. For several variables the same proof uses that a nonzero polynomial over an infinite field cannot vanish everywhere (induct on the number of variables). □

Proposition 8.6 (Integral coefficients of a bounded section)**.** *Let F be a degree-h section of the above multiplier and suppose $|F(z)|_p \le H_p(z)^h$ everywhere. Then $F = \vartheta^h(P(X)+YQ(X))$ for polynomials P, Q all of whose coefficients are p-adically integral. In particular the transported formal Taylor series of F belongs to $\vartheta^h\mathfrak{R}_\phi$. If F is even then $Q = 0$. For even h, one has $F(t) = P_{h/2}\sigma(t)^h$ when F is even, where $P_{h/2}$ denotes the coefficient in that degree.*

Proof. Write the quotient by ϑ^h as in Lemma 8.3. Let $D(X) = 4X^3 + a_1^2X^2 + 4X$. This polynomial has Gauss norm one: for odd p the coefficient 4 is a unit, and for $p = 2$ Proposition 8.1 gives that a_1 is a unit. Choose an integral x at which $|D(x)|_p = 1$ and both P, Q attain their coefficient maxima, using Lemma 8.5 and neglecting identically zero polynomials. Such a point can be chosen in the image of the complex coefficient field: it suffices to choose its residue in $\overline{\mathbb{F}}_p$ and lift algebraically. The associated points of (8.1) above x are conjugate under elliptic negation with Y coordinates y and $-a_1x - y$. They are integral by the monic quadratic equation, and their difference has norm one since $(2y + a_1x)^2 = D(x)$. By elliptic uniformization (in the shifted coordinate), choose $z, -z$ giving these two points. Here ϑ is nonzero, and $|X(z)|_p \le 1$ gives $H_p(z) = |\vartheta(z)|_p$ by $X = (\sigma/\vartheta)^2$. The hypotheses show that both

$$P(x) + yQ(x), \qquad P(x) + (-a_1x - y)Q(x)$$

are integral. Subtracting and dividing by the unit difference proves integrality of $Q(x)$, hence also of $P(x)$. The way x was chosen proves integrality of all coefficients. This proof works even at $p = 2$: we divide by the norm-one difference of the roots, not by 2. The formal integrality statement follows from $X, Y \in \mathfrak{R}_\phi$.

When F is even, $\vartheta(-z) = \vartheta(z)$ and $Y(-z) = -a_1X(z) - Y(z)$ give $(2Y + a_1X)Q(X) = 0$ identically. Thus $Q = 0$. Finally $\vartheta^hX^b = \vartheta^{h-2b}\sigma^{2b}$. The pole-order bound permits only $2b \le h$; at $z = t$ all terms with $2b < h$ vanish, proving the claimed formula. □

This is a reconstruction of `CMLattice.exists_eq_theta_pow_mul_of_norm_le`, including the exact even-coefficient assertion, from `Lib/CM/LocalHeight/IsogenyIntegrality.lean`. Applying Proposition 8.4 to the three sections ϑ, σ^2, G gives their pullbacks under an isogeny as ϑ^n (or its corresponding power) times elements of $\mathfrak{R}_\phi$. Dividing the pullbacks of G' and ϑ'^3 is permitted on the formal disc since $\vartheta'(0) = 1$ (here ϑ' is the *target theta function*, not differentiation). The factor in $\mathfrak{R}_\phi$ has constant coefficient one and therefore an integral inverse in the Y coordinate. Thus the target formal parameter pulled back by the isogeny is integral in the source parameter. Comparing their first coefficients shows that the scalar a of any isogeny between these normalized curves is integral at every p-adic embedding. We recover `rescale_translatedWpDeriv_mem` and `norm_le_one_of_isogeny`, with no loss from the degree of the isogeny.

Corollary 8.7 (Integral isogeny translates)**.** *Let $\delta \in \mathcal{O}_K$ be nonzero, put $n = \delta\bar{\delta}$, and suppose $\delta y \in \Lambda$. Then the Taylor series of $(U_y^1\vartheta)(\bar{\delta}z)$ lies in $\vartheta(z)^n\mathfrak{R}_\phi$.*

Proof. Pullback by multiplication by $\bar{\delta}$ yields the original multiplier to power n. Twisted translation of this pullback by $y/\bar{\delta}$ preserves it, since $n(y/\bar{\delta}) = \delta y \in \Lambda$ and the dual lattice contains Λ. Its degree is therefore n. Formula (8.3) identifies it with the function in the statement. Proposition 8.4 bounds its size by H_p^n. Apply Proposition 8.6. □

This proves the assertion (9.5) in the local jet chapter, using the classical theta multiplier identities as above. Notice that this corollary is a statement about a formal expansion near zero; it does not assert that the theta function is a globally analytic p-adic function.

8.5 Logarithms, norm-pair limits, and the formal height

As $\vartheta(0) = 1$, define the formal logarithm $\mathscr{L} = -\log\vartheta$. Evenness gives $\mathscr{L}(0) = \mathscr{L}'(0) = 0$. From the addition law (8.4), take the coefficient of w^2 and use $\sigma'(0) = 1$. On the left the coefficient is $\vartheta\vartheta'' - (\vartheta')^2$; on the right it is $\vartheta^2(\vartheta''(0) - X)$. Consequently

$$\mathscr{L}'' = L_2 + X, \qquad L_2 = -\vartheta''(0). \tag{8.9}$$

For $\delta \in \mathcal{O}_K \setminus \{0\}$, put $n = \delta\bar{\delta}$. The isogeny pullback at $y = 0$ proves that the series

$$\phi(\delta z), \qquad F_\delta = \exp(n\mathscr{L}(z) - \mathscr{L}(\delta z))$$

are in $\mathfrak{R}_\phi$. Since $F_\delta(0) = 1$, its inverse belongs to the same ring; differentiation and the logarithmic derivative give $D'_\delta \in \mathfrak{R}_\phi$ for $D_\delta = n\mathscr{L}(z) - \mathscr{L}(\delta z)$. The same statements hold trivially for $\delta = 0$.

Differentiating once more and evaluating at zero shows that $(n - \delta^2)L_2$ is integral. For $\delta = c_0$ and $\bar{c}_0$ the norms n coincide. Their difference is $(\bar{c}_0^2 - c_0^2)L_2 = -\sqrt{-d}L_2$. This gives the desired bound $|\sqrt{-d}L_2|_p \le 1$ in (9.3), including at the ramified prime.

View $\mathfrak{R}_\phi$ with the coefficientwise topology, not with a topology coming from a complex disc. Let Z denote the formal inverse of ϕ in the variable z. A series belongs to $\mathfrak{R}_\phi$ exactly when every coefficient after composition with Z belongs to the valuation ring. In particular $\mathfrak{R}_\phi$ is closed in the topology of pointwise convergence on coefficients. Each coefficient of $\phi(cz)$ and of $\exp(D\mathscr{L}(z) - \mathscr{L}(cz))$ depends continuously (indeed polynomially, over the coefficient field) on (c, D); the same is true of the derivative of the exponent. Thus the three integrality assertions extend to the closure of the pairs $(\iota\delta, \iota(\delta\bar{\delta}))$ in the product of the two p-adic fields. This proves the entirety of Proposition 9.3 from the preceding propositions. It is the normalization in `CMLattice.exists_log_theta`, `Lib/CM/KroneckerLimit/LogTheta.lean`, together with `rescale_translate_eq_theta_pow_mul`.

Lemma 8.8 (Prime-to-p root of unity dilation)**.** *If p is inert in K, let $q = p^2$; otherwise, let $q = p$. Let ζ be a $(q-1)$-st root of unity in the algebraically closed local field. Then (ζ, ζ^{p+1}) belongs to the norm-pair closure just described.*

Proof. In the inert case, choose an algebraic integer $a \in \mathcal{O}_K$ whose reduction at the given embedding equals that of ζ in $\mathbb{F}_{p^2}$. This is possible because $\mathcal{O}_K/(p) = \mathbb{F}_{p^2}$, and conjugation on this field is the p-power map. Thus the reduction of $\bar{a}$ is that of ζ^p. In the split or ramified case $q = p$, and one may choose a to be an ordinary integer with the specified reduction in $\mathbb{F}_p$; the two reductions then agree. Since q is a power of p, one has $\iota(a)^{q^h} \to \zeta$ and $\iota(\bar{a})^{q^h} \to \zeta^p$ as $h \to \infty$. Indeed if a unit is $\zeta(1 + w)$ with $v_p(w) > 0$, then the binomial theorem shows

$v_p((1+w)^p-1) \geq \min(v_p(w)+1, pv_p(w))$, and iteration tends to infinity. The same argument applies to the conjugate unit. The integers a^{q^h} yield the required limiting norm pairs. □

This proves `mem_normPairClosure_of_pow_eq` in `Lib/CM/KroneckerLimit/LogTheta.lean`. Note that in the inert case conjugation is genuinely a nontrivial Frobenius on the residue field. Treating it as the identity would produce the wrong sign in the theta couplings. This is also why using an arbitrary prime-to-p root of unity, rather than one of order dividing $q-1$, would not suffice.

The remaining height input in Chapter 9 may be expressed in conventional elliptic-curve terms. Assume that $a \in \mathcal{O}_K$ is nonzero, $a\bar{a} = q$ is odd, and denote multiplication by a on the curve by $[a]$. With the distinguished point t of order two, $[a](t) = t$: multiplication by a on $\mathcal{O}_K/(2) \simeq \mathbb{F}_2^2$ is invertible because its norm is odd, so it is the identity there. We consider the germ of the isogeny at the point t, as in the invariant coordinate z of the local jet theorem.

Proposition 8.9 (The degree upper bound for the first nonzero reduced coefficient)**.** *Let T be any integral formal parameter tangent to the identity in $\mathfrak{R}_\phi$ and let $Z = T^{-1}$ in the invariant differential coordinate. The first nonzero coefficient of the reduction of $g(u) = T(\iota(a)Z(u))$ has degree at most q, including when $p \mid q$.*

Proof. The equation (8.1) is a good Weierstrass model by Proposition 8.1; the endomorphism extends to the good integral elliptic curve. Its reduction is an isogeny of the same degree q. Here we use the classical good-reduction property that reduction preserves degrees of isogenies (e.g. Silverman, *Advanced Topics in the Arithmetic of Elliptic Curves*, II, Proposition 4.4). The local ring of the reduced curve at the point t is a smooth one-dimensional local ring; its completed local coordinate Y is a uniformizer, as $\partial(Y^2 + a_1XY - X^3 - X)/\partial X = -1$ there. The fibre of the isogeny above the reduced t has scheme-theoretic length q. At the point t its local length is the order of the image of the target uniformizer in the source formal coordinate; hence the order is finite and at most q. Since $T = \phi + O(\phi^2)$ is an integral change of uniformizer, the same is true of the order computed with T. Equivalently, g has a unit coefficient of degree at most q. □

This supplies the needed case of `CMLattice.exists_unit_coeff_degree_le` (`Lib/CM/FormalGroup/FormalHeight.lean`) by the ordinary formal-group interpretation of the local degree. It is essential that the bound is for a reduction modulo the maximal ideal, not for the first nonzero coefficient over the local field. The characteristic-zero linear coefficient is a, but it is a nonunit at the primes of interest. For instance an inert prime gives $a = p$, $q = p^2$, whereas the ramified prime admits $a = \sqrt{-d}$ and $q = p = d$. In Chapter 9 the symmetry under the group of order $q-1$ leaves only the first and the q-th coefficient through degree q; the first is a nonunit. The present bound therefore proves that the q-th coefficient is a unit, which is exactly the input used to normalize the leading Frobenius term of the inverse parameter.

Lemma 8.10 (A coefficient form of the Hasse invariant)**.** *Let p be odd. Suppose every $F \in \mathfrak{R}_\phi$ satisfies $|[z^{p-1}]F'|_p < 1$. Then $|[z^{p-1}]X|_p \geq 1$. Moreover, if $\chi_K(p) = 1$, the displayed hypothesis cannot hold.*

Proof. Write $p - 1 = 2k$. Assume for a contradiction that the coefficient in X has norm less than one. Since $1, X, X'$ belong to $\mathfrak{R}_\phi$, the same assertion for 1 is trivial (its positive-degree coefficient is zero). For each $n \geq 0$ the differential equation gives

$$(X^nX')' = (4n+6)X^{n+2} + (n+1)a_1^2X^{n+1} + (4n+2)X^n. \tag{8.10}$$

Indeed use $(X')^2 = 4X^3 + a_1^2X^2 + 4X$ and (8.2). For $0 \le n \le k-2$, the integer $4n+6$ is positive, even, and strictly less than $2p$, so it is prime to p. The hypothesis on derivatives and induction, using integrality of a_1, show that $[z^{p-1}]X^j$ has norm less than one for all $0 \le j \le k$. But $X = z^2 + O(z^4)$ has $[z^{p-1}]X^k = 1$. When $p = 3$ the supposed assertion for X contradicts this directly. Thus $|[z^{p-1}]X|_p \ge 1$.

To see the final statement, take the dilation $c = \sqrt{-d}$ with $n = d$. The series $(d\mathscr{L}(z) - \mathscr{L}(cz))'$ lies in the integral parameter ring, so its derivative has z^{p-1} coefficient of norm less than one. By (8.9) this coefficient is $(d - c^{p+1})[z^{p-1}]X$. If p splits in K, the reduction of c is in $\mathbb{F}_p^\times$, so $c^{p+1} \equiv c^2 = -d \pmod{\mathfrak{m}}$. The factor $d - c^{p+1}$ is then a unit as p is odd and $p \neq d$, a contradiction. □

The even prime is not nonsplit in these data, so Lemma 8.10 is precisely the relevant part of `CMLattice.one_le_norm_coeff_translatedWp_nonsplit` in `Lib/CM/FormalGroup/FormalHeight.lean`. This proof avoids an identification of the actual coefficient with a modular form. In the nonsplit calculation the gap in the invariant coordinate of the root-of-unity equivariant parameter makes the hypothesis automatic at an inert prime, giving the norm-one first exceptional coupling coefficient of the local normal form. Its sharpness is needed for the comparison in the determinant lower bound; for the upper bound it would not be sufficient just to know that all the indicated coefficients were integral.

8.6 A unit-adjusted Chowla–Selberg formula

An archimedean period for an elliptic curve generally carries no absolute p-adic coefficient normalization. The following logarithmic version of Chowla–Selberg is particularly useful because the normalizing factor is an algebraic unit, so it contributes nothing to finite valuations. We explain why it has this stronger property.

For a fractional ideal $\mathfrak{b}$ of K denote its norm by $\mathcal{N}(\mathfrak{b}) > 0$, with the usual multiplicativity under fractional ideal products. The area of $\lambda\mathfrak{b}$ is $|\lambda|^2\mathcal{N}(\mathfrak{b})\sqrt{d}/2$. For $h_K = |\mathrm{Cl}(K)|$, the class group has odd order by the classical genus theory of imaginary quadratic fields of prime discriminant: its quotient by squares has cardinality one. Thus squaring permutes its elements. Choose representatives $\mathfrak{a}_c$ for each class c. The ideals $\mathfrak{b}_c = \mathfrak{a}_{c^2}/\mathfrak{a}_c$ again represent every class once, and $\prod_c \mathfrak{b}_c = \mathcal{O}_K$ exactly. This last equality of fractional ideals, not merely of their ideal classes, is why the quotients of the chosen representatives are used.

Fix a normalized CM curve C of period lattice Λ. For every c choose a normalized curve C_c with period lattice $u_c\mathfrak{b}_c\Lambda$, where $u_c \in \mathbb{C}^\times$. Such a choice is possible by the normalization at t above. Every element of $u_c\mathfrak{b}_c$ gives an isogeny (or zero) from C to C_c, and every element of $u_c^{-1}\mathfrak{b}_c^{-1}$ gives one in the reverse direction. The isogeny integrality proved after Proposition 8.6 implies that every such multiplier is integral at every prime. The product of these fractional modules over all c is $u_*\mathcal{O}_K$, with $u_* = \prod_c u_c$. The product of the inverse modules is $u_*^{-1}\mathcal{O}_K$. Since products and sums of integral elements are integral, both u_* and its inverse are integral at every prime and every embedding. The u_c themselves are algebraic: a nonzero multiplier gives an algebraic isogeny between elliptic curves over $\overline{\mathbb{Q}}$, and its effect on the invariant differential belongs to $\overline{\mathbb{Q}}$. Dividing by an element of $\mathfrak{b}_c \setminus \{0\}$ proves this for u_c. Here we use the usual algebraicity of homomorphisms and torsion points of elliptic curves, or equivalently the description of isogenies via finite subgroups and quotients. Consequently u_* is an algebraic unit.

Write A_C^γ and A_c^γ for the period areas of the normalized curves obtained by conjugating their algebraic equations by a field automorphism γ of $\mathbb{C}$ over K. Algebraic isogenies and their differential multipliers are carried to the corresponding conjugate isogenies. Hence the

conjugate lattices satisfy $\Lambda_c^\gamma = \gamma(u_c)\mathfrak{b}_c\Lambda^\gamma$: for each individual multiplier the containment follows from the corresponding isogeny, and the reverse containment follows from the inverse module of multipliers, since $\mathfrak{b}_c\mathfrak{b}_c^{-1} = \mathcal{O}_K$. It follows that

$$\prod_c A_c^\gamma = (A_C^\gamma)^{h_K}|\gamma(u_*)|^2. \tag{8.11}$$

The curves C_c^γ have the same underlying ideal-class family, with their prescribed individual differential scalings. In particular the product formula (8.11) does not require choosing a period scaling separately for each embedding.

Put $k_0 = 2/\sqrt{d}$. For a complex lattice L of area A_L and discriminant $\Delta_L = g_2^3 - 27g_3^2$, the classical first Kronecker limit formula gives

$$(s-1)\sum_{l\in L\setminus\{0\}} \left(|l|^2/(k_0A_L)\right)^{-s} = \pi k_0 + \pi k_0(s-1)\left(2\gamma_{\mathrm{E}} - \log\frac{A_L}{k_0\pi^2} - \tfrac{1}{6}\log|\Delta_L|\right) + O((s-1)^2). \tag{8.12}$$

Here γ_{E} denotes Euler's constant. To check the normalization, write $L = \lambda(\mathbb{Z}+\tau\mathbb{Z})$ with $y = \Im\tau > 0$. Its area is $|\lambda|^2 y$ and its discriminant has modulus $(2\pi)^{12}|\lambda|^{-12}|\eta(\tau)|^{24}$. The usual constant term in the nonholomorphic Eisenstein sum of weight zero is $2\pi(\gamma_{\mathrm{E}} - \log 2 - \log(\sqrt{y}|\eta(\tau)|^2))$ and its residue is π. Multiplication by k_0^s adds $\pi k_0 \log k_0$ to the constant term and gives (8.12). We use the classical Kronecker limit formula here (Siegel, *Advanced Analytic Number Theory*, I, Theorem 1), whose exact lattice formulation is `kronecker_limit` in `Lib/CM/KroneckerLimit/LogTheta.lean`. Changing the discriminant by a factor of $(2\pi)^{12}$ without compensating in this formula would disrupt the later level cancellation.

For a fractional ideal $\mathfrak{b}$ the map $y \mapsto (y)\mathfrak{b}^{-1}$ identifies its nonzero elements, modulo the finite unit group, with integral ideals in class $[\mathfrak{b}]^{-1}$. If $L = \lambda\mathfrak{b}$, then $|\lambda y|^2/(k_0A_L) = |N_{K/\mathbb{Q}}(y)|/\mathcal{N}(\mathfrak{b})$, the norm of this integral ideal. Consequently the sum of the left sides of (8.12) over the lattices Λ_c^γ is

$$2(s-1)\zeta_K(s) = 2(s-1)\zeta(s)L(s,\chi_K),$$

as the unit group has order two ($d > 3$). This is an exact identity for $s > 1$. Taking the constant and first coefficients at $s = 1$ gives $h_K\pi k_0 = 2L(1,\chi_K)$ and

$$h_K\pi k_0\left(2\gamma_{\mathrm{E}} - \frac{1}{h_K}\sum_c \log A_c^\gamma + \log(k_0\pi^2) - \frac{1}{6h_K}\sum_c \log|\gamma(\Delta_c)|\right) = 2\big(\gamma_{\mathrm{E}}L(1,\chi_K) + L'(1,\chi_K)\big).$$

Divide by the common positive residue. With $\kappa_K = \gamma_{\mathrm{E}} - \Re L'(1,\chi_K)/L(1,\chi_K)$ this yields

$$\exp\left(\frac{1}{h_K}\sum_c \log A_c^\gamma + \frac{1}{6h_K}\sum_c \log|\gamma(\Delta_c)|\right) = \frac{2\pi^2}{\sqrt{d}}\exp(\kappa_K). \tag{8.13}$$

Proposition 8.11 (CM normalization by a single algebraic unit)**.** *There exists an algebraic unit u_0 such that for every $\gamma \in \mathrm{Aut}(\mathbb{C}/K)$,*

$$A_C^\gamma|\gamma(u_0)|^2 = \frac{2\pi^2}{\sqrt{d}}\exp(\kappa_K). \tag{8.14}$$

Proof. By Proposition 8.1, each Δ_c is an algebraic unit. Choose a nonzero root

$$u_0^{12h_K} = u_*^{12}\prod_c \Delta_c.$$

A root of an algebraic unit and its inverse are integral at all primes, so u_0 is an algebraic unit. By (8.11), the expression on the left in (8.13) is $A_C^\gamma|\gamma(u_*)|^{2/h_K}\prod_c|\gamma(\Delta_c)|^{1/(6h_K)} = A_C^\gamma|\gamma(u_0)|^2$. This proves (8.14). □

This is the exact assertion of `CMLattice.chowla_selberg` in `Lib/CM/FormalGroup/FormalHeight.lean`. Notice the two elements of the argument: the analytic limit computes the mean archimedean area, while the ideal-class ratios and the isogeny integral scales show that this mean is reached using an actual algebraic unit. A formula merely describing the product of the areas of conjugates, without showing the unit-adjustment property, would not establish the individual simultaneous norm comparison of the local jet theorem.

8.7 The finite idele and the period scale

Let ζ be a primitive N_1-st root of unity, $D = \Phi'_{N_1}(\zeta)$. We use the distinction between the relative different of E/K, generated by D, and the absolute different, generated up to a unit by $\sqrt{-d}D$. Define the finite idele α by

$$\alpha_w = \begin{cases} D^{-1} & c_0 \in w, \\ 1 & c_0 \notin w. \end{cases}$$

This is the idele α (namespace `Parameters`) of `New/Data/WeilDataOfE.lean`; the notation $c_0 \in w$ means that the integer belongs to the finite prime ideal of E. All rational primes dividing N_1 divide $n_0 = (d+1)/4 = c_0(1-c_0)$. At a prime w above such a rational prime exactly one of c_0 and $\bar{c}_0 = 1 - c_0$ vanishes modulo w. The relative different generator divides N_1, and its valuation is invariant under Galois automorphisms over $\mathbb{Q}$ because the relative different ideal is preserved. Hence at every finite place w, $\alpha_w \bar{\alpha}_w D$ is a unit. Also α^{-1} and $N_1\alpha$ are integral finite adeles. These are `v_alpha_mul_relDifferentGen`, `inv_alpha_mem` and `N1_mul_alpha_mem` in their typographically analogous source declarations. As a consequence the fractional lattice

$$\mathcal{I} = E \cap \alpha\widehat{\mathcal{O}}_E$$

contains $\mathcal{O}_E$ and is contained in $N_1^{-1}\mathcal{O}_E$. Moreover $x\bar{y}D$ is integral whenever $x, y \in \mathcal{I}$; locally this follows by multiplying the two bounds and using $\alpha\bar{\alpha}D$ a unit. By the trace-dual characterization of the relative different, $\mathrm{Tr}_{E/K}(x\bar{y}) \in \mathcal{O}_K$. In particular the Hermitian trace pairing on $\mathcal{I}$ takes integral values. This proves the lattice assertions in `herm_mem_quadInt` in `New/Curve/TightFrameExists.lean`.

The finite splitting character in the Weil datum is denoted by χ_{fin}. Its necessary property can be stated as follows without any archimedean topology on the field of complex values. It is trivial on integral finite units; on a principal finite idele y one has

$$\chi_{\mathrm{fin}}(y) = \prod_{\tau \in \Phi} \tau(\bar{y}). \tag{8.15}$$

This is the type of the normalized finite character, not the unitary Hecke character prior to the absolute-value factor. The compatible splitting character is constructed in Proposition 6.4 and its finite normalized action in Theorem 7.6. The remaining comparison with the archimedean arithmetic theta model is a specific input; see `Parameters.instWeilDataE` in `New/Data/WeilDataOfE.lean` and `nonempty_splittingChar` in `Lib/ClassField/HeckeCharacter/SplittingCharExists.lean`. For the following normalization computation it suffices to assume exactly (8.15) and triviality on the finite unit subgroup. Thus this input is separated from the estimates about theta coefficients.

If z is any finite idele, some positive power z^h has the same valuations as a principal idele y, by finiteness of the ideal class group. Equation (8.15) implies

$$\chi_{\mathrm{fin}}(z)^h = \prod_{\tau} \tau(\bar{y}).$$

The latter is in $K^\times$. Taking absolute values of all its images under a p-adic embedding ι and extracting the positive h-th root proves

$$|\iota(\chi_{\mathrm{fin}}(z))|_p = \prod_{\tau\in\Phi} |z|_{\iota\tau\circ\mathrm{conj}}. \tag{8.16}$$

In the factor on the right, the component of z at the place induced by the indicated embedding is read in the local field. It depends only on the valuation of that component. This establishes `exists_chiFin_pow_eq_prod` and `norm_map_chiFin` of `Lib/ClassField/HeckeCharacter/SplittingCharOnU1.lean` directly from the ideal-class argument and character type.

Let s_0 be a nonzero f-th root of $\chi_{\mathrm{fin}}(\alpha)^{-1}$. The ideal of $\alpha\bar\alpha$ is the inverse relative different. The absolute norm of that different is $(\mathrm{rd}_E/\sqrt{d})^{2f}$. Taking norms of both sides shows that the ideal norm of α is $(\mathrm{rd}_E/\sqrt{d})^{-f}$. If α^h has the valuations of the principal ideal (y), it follows that $|N_{E/\mathbb{Q}}(y)| = (\mathrm{rd}_E/\sqrt{d})^{-hf}$. By (8.15), the squared ordinary absolute value of $\chi_{\mathrm{fin}}(\alpha)^h$ is this rational ideal norm. Since the value on the right of (8.15) is in K, any field automorphism γ over K fixes it. Thus, taking positive roots,

$$|\gamma(s_0)|^2 = \mathrm{rd}_E/\sqrt{d}. \tag{8.17}$$

The compact derivation is `Parameters.norm_gal_vacuumRoot` in `New/Curve/PeriodDefs.lean`. At a fixed p and ι, the norms of the components of α read through $\iota\tau\circ\mathrm{conj}$ are the same for all $\tau\in\Phi$. They are either all one or all equal to $|\iota(D)|_p^{-1}$: the embeddings τ extend the fixed embedding of K, and the valuations of the conjugate generators of the relative different coincide. Dividing (8.16) for $z=\alpha$ by f in logarithms gives $|\iota(s_0)|_p = |\alpha|_{\iota\tau\circ\mathrm{conj}}^{-1}$.

Now let u_0 be as in Proposition 8.11 and $s_1 = s_0/u_0$. Its finite valuations are the same as those of s_0 because u_0 is a unit. Therefore for every $y\in\mathcal{I}$,

$$|\iota(s_1\overline{\tau(y)})|_p \le 1, \qquad |\iota(s_1)|_p = 1 \quad (p\nmid N_1). \tag{8.18}$$

This is exactly the finite normalization required in Proposition 9.2, with the stronger equality of local norms of s_1 and the inverse component of α. It proves `norm_iota_periodFactor`, `norm_periodFactor_mul_conj_le_one` and `norm_iota_periodFactor_eq_one` in `New/Curve/TightFrameExists.lean`, from the splitting character property and the classical CM results. In particular the primitivity condition of the tight frame cannot be obtained by adjusting s_1 at each prime; it has to be built into the frame itself.

Finally, if a conjugate normalized curve has area $|\Omega_\gamma|^2\sqrt{d}/2$ and $t_\gamma = \gamma(s_1)(2\pi/\sqrt{d})^{1/2}/\Omega_\gamma$, then (8.17) and (8.14) give

$$|t_\gamma|^2 = \frac{\mathrm{rd}_E}{2\pi e^{\kappa_K}} = \Omega^{-1}. \tag{8.19}$$

This is the numerical relation used for the conjugate archimedean estimates. It shows exactly how the discriminant and logarithmic derivative in the main height cost enter: the normalized complex Hermitian coordinates have scale $(2\pi/\sqrt{d})^{1/2}$, whereas the CM multiplier is normalized by a unit using the mean area. Their ratio, after the finite different adjustment, is a fixed positive real number independent of the conjugate curve.

8.8 Constructing arithmetic tight frames of prescribed length

We now prove the arithmetic frame assertion of Proposition 9.10. This step deserves a separate proof: a family of lattice vectors may be a perfectly good tight frame without the length being a p-adic unit, and may have every individual vector too small in the coefficient normalization s_1. Both problems arise in the theta estimates. Write $N = N_1$, not the auxiliary level prime,

and $\zeta = \zeta_N$ as above. Throughout the construction $n_0 = (d+1)/4$. We have $p \mid n_0$ for every rational prime $p \mid N$. Recall that a tight frame is a family $w_i \in \mathcal{I}$ with positive integer multiplicities m_i such that

$$\sum_i m_i \sigma(w_i)\overline{\tau(w_i)} = M\mathbf{1}_{\sigma=\tau} \quad (\sigma, \tau \in \Phi). \tag{8.20}$$

The parameter M is called the frame length; the cardinality of the indexing family need not equal M. For example the N roots ζ^j, each with multiplicity one, give length N by the geometric sum $\sum_{j=0}^{N-1}(\sigma\zeta/\tau\zeta)^j$. A direct sum of two frames with lengths M_1, M_2 has length $M_1 + M_2$, and multiplying all multiplicities by an integer multiplies the length. We seek another frame of length coprime to N, with additional local unit vectors in this family of frames.

Lemma 8.12 (A cyclotomic factor of the different)**.** *For every rational prime $\ell \mid N$, put $r_\ell = N/\ell$, $A_\ell = \zeta^{r_\ell}$, and let $M_\ell = N/\ell^{v_\ell(N)}$. Then*

$$V_\ell = \frac{D(A_\ell - 1)M_\ell}{N} \in \mathcal{O}_E, \qquad V_\ell \mid M_\ell.$$

In particular V_ℓ is an ℓ-adic unit at every embedding above ℓ.

Proof. The standard cyclotomic derivative formula follows from $\Phi_N(X) = \prod_{a|N}(X^a - 1)^{\mu(N/a)}$ by differentiating and evaluating at ζ. Up to the unit ζ^{-1} it gives

$$D = N\zeta^{-1} \prod_{a|N,\ a<N} (\zeta^a - 1)^{\mu(N/a)}.$$

For any prime above a rational prime p, $\zeta^a - 1$ is a unit unless the order N/a of ζ^a is a power of p. Indeed if that order has nontrivial prime-to-p part, its reduction has that same prime-to-p part; otherwise the valuations of $1 - \zeta_{p^k}$ are the familiar $1/(p^{k-1}(p-1))$, which follow from $\Phi_{p^k}(1) = p$ and conjugacy of the factors. Since the Möbius exponent vanishes for a nonsquarefree integer, the only nonunit factor of the product at p occurs when $N/a = p$. Hence $v_p(D) = v_p(N) - 1/(p-1)$ when $p \mid N$, and zero elsewhere, uniformly in the embedding. The factor $A_\ell - 1$ has valuation $1/(\ell-1)$ above ℓ and zero elsewhere. Thus $v_p(V_\ell)$ is zero above ℓ and outside the primes dividing N, and equals $v_p(M_\ell) - 1/(p-1)$ above each $p \mid M_\ell$. These are nonnegative numbers at most $v_p(M_\ell)$, proving both assertions by the characterization of algebraic integers by their finite valuations. □

The source `Parameters.exists_cyclotomic_different_factor` in `New/Curve/TightFrameExists.lean` proves this by a polynomial factorization with a bound $M_\ell \le N$. The explicit prime-to-ℓ part above is one convenient choice satisfying its conclusions.

Since n_0 is divisible by all prime divisors of N, it is nilpotent modulo N. The map $z \mapsto n_0 z^2 + z$ on $\mathbb{Z}/N\mathbb{Z}$ is injective: the difference of its values at z, z' is $(z - z')(1 + n_0(z + z'))$, and the second factor is a unit modulo N. It is therefore surjective. Choose z_0 with $0 \le z_0 < N$ such that

$$n_0 z_0^2 + z_0 + 1 = eN, \qquad e = (n_0 z_0^2 + z_0 + 1)/N \in \mathbb{N}.$$

The integer e is positive and $e \le n_0 N + 1$. Define the integer of K

$$b_0 = 1 - c_0 + n_0 z_0.$$

Then $b_0 + \bar{b}_0 = 1 + 2n_0 z_0$, and calculating the norm with $c_0\bar{c}_0 = n_0$ gives $b_0\bar{b}_0 = n_0 eN$. In particular b_0 is a unit at the places above $p \mid N$ where c_0 vanishes, while at the other places

above those primes it is divisible by N: its conjugate is a unit there and its norm is a multiple of N.

For a rational prime $\ell \mid N$ take the following family, indexed by $0 \le j < r_\ell$ and $0 \le k, i < \ell$:

$$w_{\ell,j,k,i} = \zeta^j \left(\mathbf{1}_{k=0} A_\ell^i + \frac{b_0 M_\ell}{N} (A_\ell^i - A_\ell^k) \right). \tag{8.21}$$

This is an arithmetic modification of the usual cyclotomic frame. Both types of lattice condition are satisfied. For a finite prime above a rational prime outside N, all denominator factors are units. At a prime above N where c_0 does not vanish, b_0 is divisible by N, so the vector is integral. At a prime where c_0 vanishes, $A_\ell^i - A_\ell^k$ is an integral multiple of $A_\ell - 1$; multiplying the second term by D gives the integral factor $b_0 V_\ell$ times that multiple. Thus $D w_{\ell,j,k,i}$ is integral there. This is precisely membership in $\mathcal{I}$, whose local allowed denominator is D in the latter case.

Lemma 8.13 (Length of the modified cyclotomic frame)**.** *The family* (8.21)*, with multiplicities one, is a tight frame of length*

$$C_\ell = N + M_\ell(1 + 2n_0 z_0 + 2\ell n_0 M_\ell e). \tag{8.22}$$

Moreover for each $p \mid N$ the family with $\ell = p$ contains a vector $w \in \mathcal{I}$ with $|\iota(s_1 \overline{\tau(w)})|_p = 1$ at all τ, for each ι for which $|\iota(s_1)|_p < 1$.

Proof. Set $T = b_0 M_\ell / N$. For any two complex embeddings $\sigma, \tau \in \Phi$, let $A = \sigma(A_\ell)$ and $B = \tau(A_\ell)$. Both are primitive ℓ-th roots. The sums over i, k in the inner product of (8.21) can be expanded into four terms. The products of the two terms with $k = 0$ sum to $\ell \mathbf{1}_{A=B}$. Each mixed product is $\ell T \mathbf{1}_{A=B}$ or $\ell \bar{T} \mathbf{1}_{A=B}$: all sums with just one nontrivial power of A or B vanish. Finally

$$\sum_{i,k} (A^i - A^k)(\overline{B^i} - \overline{B^k}) = 2\ell^2 \mathbf{1}_{A=B}.$$

Multiplying by the remaining j-sum gives $\sum_j (\sigma\zeta/\tau\zeta)^j$, which is r_ℓ if $\sigma\zeta = \tau\zeta$ and zero if $A = B$ but the embeddings differ (in the latter case its ratio is a nontrivial r_ℓ-th root of unity). Thus the full matrix of products equals $N(1 + T + \bar{T} + 2\ell T\bar{T})$ times the identity. Substituting the trace and norm of b_0 gives (8.22).

For the second assertion take $\ell = p$, $j = 0$, $k = 1$, $i = 0$. The chosen vector is $w = (b_0 M_p / N)(1 - A_p)$, and $Dw = -b_0 V_p$. If $|\iota(s_1)|_p < 1$, the relevant components of α above p read through every $\iota\tau \circ \mathrm{conj}$ are D^{-1}. Thus all those places contain c_0. The product $-b_0 V_p$ is a unit at every such place, by the definition of b_0 and Lemma 8.12. Equation (8.18) with its equality of finite scale factors proves the claimed unit norm. □

Let $C = \sum_{\ell \mid N} C_\ell$, the length of the union of the modified frames. It is positive and coprime to N. Indeed for every $p \mid N$, the terms with $\ell \neq p$ in the sum are zero modulo p (since $p \mid M_\ell$), while the term with $\ell = p$ is $M_p \not\equiv 0 \pmod{p}$ (since $p \mid n_0$). This proves the coprimality at each prime divisor.

Proposition 8.14 (A prescribed-length frame with a simultaneous unit vector)**.** *If $M \ge (dN)^6(R+1)$ with $R \ge 0$, there is a tight frame of length M with multiplicities at least R and containing, for every rational prime p and every ι, some vector w such that $|\iota(s_1 \bar{\tau}(w))|_p = 1$ for all $\tau \in \Phi$.*

Proof. We first bound the lengths sufficiently crudely. The number of distinct prime factors of N is at most N, and each M_ℓ, ℓ, z_0 is at most N. Together with $e \le n_0 N + 1$ this gives

$C \le 8(n_0+1)^2N^5$, hence $NC \le (dN)^6$ (recall $n_0 = (d+1)/4$, $d > 3$). Also $C \ge N$ and $C + N \le NC$ as $N \ge 2$. Consequently

$$M \ge R(C+N) + (C-1)(N-1).$$

Since C, N are coprime, every integer $L \ge (C-1)(N-1)$ is a nonnegative integer combination $AC + BN$. To verify the numerical assertion, choose $B \in \{0, \dots, C-1\}$ with $L \equiv BN \pmod{C}$; then $L - BN \ge -(C-1)$ and is divisible by C, hence nonnegative. Apply this to $L = M - R(C+N)$. Take the union of all the modified frames with multiplicities $R + A$ and of the elementary cyclotomic frame with multiplicities $R + B$. Its length is M and all multiplicities are at least R. The elementary frame contains 1; this gives the required assertion whenever $|\iota(s_1)|_p = 1$. Otherwise $p \mid N$ by (8.18), and the vector in Lemma 8.13 gives it. If positive multiplicities on both families rather than nonnegative ones are desired in the case $R = 0$, observe that the same crude bound improves strictly to $NC \le (dN)^6/4$ for the present d, N; we may first subtract $C + N$ and still satisfy the numerical representability condition. Thus one can take both family multiplicities positive for all these large lengths, without changing the stated bound. □

This proves `Parameters.exists_isTightFrame` in `New/Curve/TightFrameExists.lean`, and explains the two roles of the modified cyclotomic family. Its length, not the roots of unity individually, removes a divisibility obstruction; its special difference vector can compensate exactly the finite different scale at every embedding above a given prime. The proof gives a deterministic construction using algebraic integers and ideal norms; no approximation of an arbitrary archimedean isometry by rational matrices is necessary.

8.9 Frobenius on the frame matrix

The reduction in the Frobenius step is not always to a field. It is important to establish the equality in a quotient of a valuation ring, rather than infer a stronger assertion from an equality of residue-field values. Let p be nonsplit, let v_p have $v_p(p) = 1$, and set

$$\omega_p = \begin{cases} 1 & p \text{ inert}, \\ 1/2 & p = d, \end{cases} \qquad J_p = \{x \in \mathfrak{o} : v_p(x) \ge \omega_p\}.$$

The quotient $\mathfrak{o}/J_p$ is a ring of characteristic p. Since all primes dividing N split in K, N is a unit in this quotient, and $\mathcal{I}$ is integral at p.

Lemma 8.15 (Reduction of conjugate embeddings)**.** *Let $\sigma \in \Phi$, and let e be odd if p is inert, arbitrary nonnegative if $p = d$. For every $w \in \mathcal{I}$ one has*

$$(\iota\overline{\sigma(w)})^{p^e} = \iota\sigma'(w) \pmod{J_p}, \qquad \sigma' = [-p^e]\sigma.$$

Here the bracket indicates the action on the cyclotomic embeddings: $\sigma'(\zeta) = \sigma(\zeta)^{-p^e}$, while $\sigma'|_K$ is the fixed CM embedding.

Proof. The cyclotomic field and K have coprime discriminants, so $\mathcal{O}_E = \mathcal{O}_K[\zeta]$. It suffices to check the identity on $\mathcal{O}_K$, ζ, and N^{-1}, since $\mathcal{I} \subseteq N^{-1}\mathcal{O}_E$ and both sides are ring homomorphisms to a ring of characteristic p. At an inert prime, the subring from $\mathcal{O}_K$ maps through $\mathcal{O}_K/(p) = \mathbb{F}_{p^2}$. There the p-power map is conjugation. Thus applying an odd number of Frobenius maps after conjugation returns the original CM embedding. At the ramified prime, $\sqrt{-d}$ belongs to J_p, and the two embeddings of K reduce to the same embedding of $\mathcal{O}_K/(\sqrt{-d}) = \mathbb{F}_p$. Consequently every power of Frobenius acts identically on that subring,

including the zeroth power. In either case the identity on ζ follows from $\overline{\sigma(\zeta)} = \sigma(\zeta)^{-1}$, and the identity on N^{-1} from the fact that the prime field is fixed. This verifies the claimed equality in the specified nonreduced quotient as well as in the residue field. □

Proposition 8.16. *Repeat the vectors of a frame according to their multiplicities, so that* $\sum_i \sigma(w_i)\overline{\tau(w_i)} = M\mathbf{1}_{\sigma=\tau}$. *For an integral scalar* β, *define* $r_{\sigma i} = \beta\iota(\overline{\sigma(w_i)})$. *In every ring of characteristic* p *obtained from* $\mathfrak{o}$ *by a homomorphism annihilating* J_p *one has*

$$\sum_i r_{\sigma i}^{p^e} r_{\tau i} = \beta^{p^e+1} M \mathbf{1}_{\tau=[-p^e]\sigma}$$

for precisely the exponents in Lemma 8.15.

Proof. Factor out β^{p^e+1} and apply Lemma 8.15 to the first factor of each product. The remaining sum is the image of $\sum_i \sigma'(w_i)\overline{\tau(w_i)}$, which is the original frame identity with its first row replaced by σ'. □

Take $\beta = N\iota(b')$ in the normalization of (9.13). This proves Lemma 9.12, the statement `Parameters.exists_frame_rows_frob` of `New/Curve/ThetaProducts.lean`. In particular the minus sign in the permutation is not an arbitrary transpose or sign convention. The complex conjugation in the columns is essential, and for inert primes it cancels the odd Frobenius on the quadratic subfield while it reverses the cyclotomic exponent. The restriction on the live set of exponents is therefore intrinsic to the use of the half-set of CM embeddings.

8.10 Theta-group orthogonality and the finite realization

The following general linear-algebra argument isolates the substantial, but classical, mechanism in the finite expansion of a vacuum. Let L be a full lattice in a complex vector space of dimension g, and let B be a positive Hermitian form whose associated Riemann form $\operatorname{Im} B/\pi$ is integral on L. Use the normalization of the theta shift

$$U_a S(z) = \exp\bigl(-B(a,a)/2 - B(a,z)\bigr) S(z+a)$$

with B conjugate-linear in its first slot (equivalently one can use the conjugate convention written in (8.3)). Assume an Appell–Humbert multiplier has been fixed, and let $L^\vee$ be the integral dual for the resulting alternating form. Then U_a for $a \in L^\vee$ preserves the theta space H, and the shifts for $a \in L$ act by scalars. The alternating form induces a perfect alternating pairing on the finite abelian group $D = L^\vee/L$. The dimension of H is $m = \sqrt{|D|}$, by the usual Riemann–Roch dimension theorem for positive theta spaces; this classical theorem is also obtained by writing theta series in a symplectic basis of L. We use only its dimension formula here.

Lemma 8.17 (Theta-group averaging)**.** *Let* R_x *be a representative operator of the finite theta group on* H *for each* $x \in D$, *and denote its inverse by* R_{-x} *in the inverse-pair convention for shifts. Then for* $F, v \in H$ *and every linear functional* ℓ *on* H,

$$m \sum_{x \in D} \ell(R_x v) R_{-x} F = |D| \ell(F) v. \tag{8.23}$$

Proof. For clarity we establish the averaging formula directly. The commutator pairing on D is nondegenerate. Therefore for $x \neq 0$ some R_y conjugates R_x to a nontrivial scalar multiple of itself, so $\operatorname{tr} R_x = 0$. Also $\operatorname{tr}(R_x^{-1} R_{x'}) = 0$ for $x \neq x'$ and equals m for $x = x'$. These identities make the $|D| = m^2$ operators R_x a basis of $\operatorname{End}(H)$. The operator $\sum_x R_x^{-1} A R_x$ commutes

with every R_y because the central factors cancel in conjugation. Hence it is scalar. Its trace is $|D|\operatorname{tr}(A)$, so it equals $|D|\operatorname{tr}(A)I_H/m$. Substitute the rank-one operator A and apply the equality to v. Since $\operatorname{tr} A = \ell(F)$ this proves (8.23). In particular the span of $R_x v$ for $v \neq 0$ is H: the displayed matrix basis shows that any vector can be obtained by applying a linear combination of these operators to v. □

In the intended application the vector space H is an archimedean theta space, and we require its realization in the finite Weil model. We state explicitly which facts about the comparison of these two models we use without reconstructing their entire representation theory. For the standard space of the manuscript, put $g = fh$, where $f = |\Phi|$ and h is the rank. Write Θ for the linear map from the finite part to its entire Fock theta functions and v for the vacuum. We use the following formal-source input from `Lib/Weil/Construction/FockThetaFunctions.lean`: Θ is injective, $\Theta v \neq 0$, and

$$\Theta\big(\varrho(\operatorname{finImage}\gamma)F\big) = U_{-\operatorname{realImage}\gamma}(\Theta F). \tag{8.24}$$

The last identity is `entireTheta_translate_finiteImage`; the comparison of multipliers is `semichar_finiteImage`. The archimedean multiplier is the canonical one defined by the quadratic integer used in the construction of the standard lattice. The local lattice is self-dual for the finite additive character, and rational finite points are dense. These assertions explain why the finite and real shifts give inverse central characters. We retain these exact model-comparison identities as formal-source inputs (they can be verified in the integral model by the Gaussian integral and the additive-character product formula). The remainder of the realization is elementary theta-group theory.

To see it, let $\delta \in \mathcal{O}_K$ with $D_\delta = \delta\bar{\delta} > 0$. Let L be the archimedean image of $\bar{\delta}\mathcal{I}^h$ with the sign and real scale used in (8.24). The function Θv belongs to its theta space: for $\gamma \in \mathcal{I}^h$, invariance of the vacuum under the finite lattice gives exactly the canonical multiplier of the shift by $-\operatorname{realImage}(\bar{\delta}\gamma)$. Every dual point $a \in L^\vee$ comes from a rational point γ_a because some positive integer multiple of a belongs to L. In fact $\delta \operatorname{finImage}(\gamma_a)$ is in the finite lattice. One can see this by commuting its finite shift with all finite lattice shifts from $\bar{\delta}\mathcal{I}^h$: the scalar commutator is the inverse real pairing, which is trivial by duality. Density extends this statement from rational lattice points to the finite lattice, and self-duality gives the asserted lattice condition. Conversely the image of these rational points is in $L^\vee$ by the same argument. Thus every dual translation of Θv belongs to the image of Θ in (8.24). By Lemma 8.17 these translations span the theta space. In particular any entire S satisfying these multiplier conditions lies in the image of a unique finite vector.

In particular $D_\delta L^\vee \subseteq L$: the dual rational representative satisfies $\delta\gamma_a \in \mathcal{I}^h$, so $D_\delta\gamma_a = \bar{\delta}(\delta\gamma_a) \in \bar{\delta}\mathcal{I}^h$. We can therefore average over the larger finite group $L^\vee/D_\delta L^\vee$. An element of the smaller theta group is counted a constant number of times, and the central factors cancel in the two inverse shifts. The order of the larger group, since $L^\vee$ has rank $2g$, is D_δ^{2fh}. If I indexes its representatives, they give finite coordinates αx_j with δx_j integral, exactly as in Proposition 9.11. Transport (8.23) through the injective map Θ and put $M_\delta = m$. This yields

$$\sum_{j\in I} M_\delta \ell(\varrho(\alpha x_j)v)\varrho(-\alpha x_j)F = D_\delta^{2fh}\ell(F)v,$$

for any linear functional on the finite part, not only for the theta functional. This is the finite vacuum expansion in `Space.VacuumExpansion`, `Lib/Weil/Construction/ConjugateJetDefs.lean`. The argument recovers `Space.exists_entireTheta_eq` from the stated finite–archimedean comparison and the usual positive theta dimension formula. The source's `ThetaOrthogonality.lean` develops the theta-group orthogonality directly using entire functions as well as bilinear pairings.

8.11 From an arithmetic frame to a theta section

We apply the preceding realization to a product of one-dimensional CM sections and record the area normalization. Let C be a normalized CM curve with area A, $a \in \Lambda_C$, $\delta\bar{\delta} = D_\delta > 0$, and let (w_i, m_i) be a tight frame of length M. For each coordinate $k = 1, \dots, h$ choose a degree-m_i monomial section G_{ik} on C. Set $s_{\mathrm{ar}} = \sqrt{2\pi/\sqrt{d}}$ and

$$L_{ik}(z) = \frac{a}{\bar{\delta} s_{\mathrm{ar}}} \sum_{\tau \in \Phi} \overline{\tau(w_i)} z_{\tau k}, \qquad S(z) = \prod_{i,k} G_{ik}(L_{ik}(z)).$$

In the normalization of this chapter the polarization form of the degree-one curve section is $B_C(u, v) = \pi \bar{u} v / A$. The arithmetic condition is the identity of areas

$$A = \frac{|a|^2}{D_\delta} M \frac{\sqrt{d}}{2}. \tag{8.25}$$

Indeed the sum of pullback forms, with their degrees, is

$$\sum_{i,k} m_i B_C(L_{ik}u, L_{ik}v) = \frac{\pi |a|^2 M}{A D_\delta s_{\mathrm{ar}}^2} \sum_{\tau,k} \overline{u_{\tau k}} v_{\tau k} = \sum_{\tau,k} \overline{u_{\tau k}} v_{\tau k},$$

by the frame identity and (8.25). This identity is necessary for the product S to realize the standard theta function, rather than one in a different polarization.

For each translated period $-\operatorname{realImage}(\bar{\delta}\gamma)$ with $\gamma \in \mathcal{I}^h$, its image under the linear form is $-a \operatorname{Tr}_{E/K}(\gamma_k \bar{w}_i)$, an element of Λ_C by the trace pairing integrality proved above. The multiplier of a pullback of a theta section is the pullback of its multiplier. Its canonical factor is expressed using the imaginary part of the polarization evaluated on the quadratic multiplication by $\bar{c}_0$ (as in the earlier computation of the multiplier on the basis of the lattice). The L_{ik} intertwine this multiplication, since the embeddings in Φ restrict to the fixed embedding on K. Products of the canonical multipliers are therefore the canonical multiplier for the sum of the pullback forms. Consequently S has exactly the required multiplier, and the preceding finite realization gives a unique F with $\Theta F = S$ and the vacuum expansion.

For any rational vector γ put

$$y_{ik} = -\frac{a}{\bar{\delta}} \operatorname{Tr}_{E/K}(\gamma_k \bar{w}_i).$$

Under (8.24) the entire theta function of the translate of F by γ is $\prod_{i,k} (U_{y_{ik}} G_{ik})(L_{ik} z)$, since the exponential factors in the shifts multiply using the equality of the polarization forms. Taking homogeneous Taylor coefficients is legitimate: at each degree the product is a finite sum, and the Taylor series along all complex lines determines the multivariable Taylor series. This gives both the product Taylor identity and the theta evaluation asserted in Proposition 9.11.

For integrality choose a positive integer n clearing the finite denominators of γ: $n\gamma \in \mathcal{I}^h$. The trace integrality shows that $n D_\delta y_{ik} = -\delta a \operatorname{Tr}_{E/K}(n\gamma_k \bar{w}_i) \in \Lambda_C$. Thus y_{ik} is torsion. The normalized curve monomial sections have norm at most one at every finite valuation and the height of zero is one. The torsion translation estimate above therefore gives $|\iota(U_{y_{ik}} G_{ik})(0)|_p \le 1$ at every embedding, and hence the product theta evaluation has norm at most one. The finite model identifies its point evaluations with the theta evaluations of its finite translates, via `Space.θ_ϱ` (theta-rho evaluation) in the notation of the source. All translates of a fixed finite vector are locally constant in the finite adelic variable. By density each finite translate agrees with one by a rational point, so the same integral bound holds at every finite point. This proves the product and integrality parts of `Parameters.exists_finitePart_of_isTightFrame` in

`New/Curve/FrameSections.lean`, conditional on the stated representation comparison and the standard theta dimension theorem. It is insufficient for this argument to check integrality at the origin alone without controlling rational torsion translates.

As a useful special case, take w_i to be all the powers of ζ, each repeated N times, giving length N^2, and take $\delta = N$, $a = \Omega_0$ and $C = \mathbb{C}/\Omega_0\mathcal{O}_K$. These choices exactly satisfy (8.25). Choose every section equal to the normalized degree-one theta section, which has constant value one. The resulting F has $\theta(F) = 1$. By taking the theta functional in the vacuum expansion we express v as a finite linear combination of shifts of F. Since the finite part is by definition the span of translates of v, translates of F also span the entire finite part. This is the generator in the first part of the uniform conjugation assertion.

8.12 Integral vacuum and the remaining reciprocity law

It is useful to distinguish what is required from the complex multiplication reciprocity theorem from the elementary bound on the integral vacuum. The needed reciprocity is not merely the assertion that the j-invariant generates a ring class field. We require a comparison of normalized Taylor coefficients of theta functions and their torsion translates. This normalized reciprocity theorem is proved in Theorem 12.5. In the notation of the source this comparison is `CMLattice.complex_multiplication_theta`, used at the beginning of `Parameters.exists_isJetConjugate_generator` in `New/Curve/ConjugateGenerator.lean`. Concretely, starting with the principal normalized CM curve, it selects, for each $\gamma \in \mathrm{Aut}(\mathbb{C}/K)$, a conjugate normalized curve C_γ of area $|\Omega_\gamma|^2\sqrt{d}/2$, an integer $e_\gamma > 0$ with $e_\gamma\Omega_\gamma \in \Lambda_\gamma$, and a finite idele defining the transport of torsion coordinates. The key equality compares the coefficientwise Galois action on the twisted Taylor series with the series on C_γ after this transport and a differential rescaling. We use this equality as a formal-source input, including the necessary compatibility of the finite torsion transport for all vectors, not merely existence of some torsion point for each evaluation. The classical reciprocity theorem for algebraic elliptic curves explains the choice of C_γ; tracking normalized theta multipliers and the splitting character requires its more precise form here. We do not independently reconstruct that form of the theorem.

One can nevertheless prove from this input why the same finite generator works in every rank and why its adelic translates have the desired conjugate bounds. For C_γ use the same frame of length N^2 from the last paragraph, but put $a = e_\gamma\Omega_\gamma$ and $\delta = Ne_\gamma$. Then

$$\frac{|e_\gamma\Omega_\gamma|^2N^2}{(Ne_\gamma)^2}\frac{\sqrt{d}}{2} = |\Omega_\gamma|^2\frac{\sqrt{d}}{2},$$

so the area condition (8.25) holds exactly. The linear forms in this frame have common scalar $\Omega_\gamma/(Ns_{\mathrm{ar}})$, and after the conjugate period scale this scalar is $\gamma(s_1)/N$, with

$$t_\gamma = \frac{\gamma(s_1)s_{\mathrm{ar}}}{\Omega_\gamma}, \qquad |t_\gamma|^2 = \Omega^{-1}$$

by (8.19). The rational coefficients in each linear form are the conjugates of the original cyclotomic coefficients, with the permutation of places determined by γ. The trace formula for the y_{ik} has the same pattern on both sides; its scale is $-\Omega_\gamma/N$. Applying the normalized reciprocity comparison to these trace torsion coordinates yields the equality of the coefficients of each factor. Multiplication of the finite number of factors proves the equality for the jets of rational translates. For arbitrary finite translates, use the common finite idele furnished by the reciprocity theorem. Multiplication by this idele is continuous, and both finite-vector orbits are locally constant, so on any prescribed pair of translates the equality reduces

to a rational approximation in the corresponding common open set. This is the claimed compatibility for the whole finite orbit, rather than just the generator at zero. Together with the theta-group expansion already proved, the argument explains the second part of `Parameters.exists_isJetConjugate_generator` subject to the normalized reciprocity input.

The remaining estimate concerns the size of the scalar by which the vacuum is normalized, not the norm of a particular representative of the differential. Let F be our integral vector with $\theta(F) = 1$, let I be the finite expansion set with $|I| = D_\delta^{2fh} = N^{4fh}$, and use the standard inner product, normalized by $\langle v, v\rangle = 1$. In the standard finite Weil model, for all finite shifts $R_x = \varrho(\alpha x)$,

$$\langle v, R_x v\rangle \in \{0\} \cup \mu_\infty. \tag{8.26}$$

Here the zero is for a shift outside the lattice supporting the vacuum, and the root of unity is its prescribed finite multiplier inside that lattice. This is the vacuum overlap calculation `Space.dataInner_rho_of_vacElem` in the source. We retain it as part of the comparison with the standard finite Weil model. These are the only possible values: for a translation in the stabilizer the vacuum transforms by its finite character, and for a translation outside the lattice, commute with a lattice element on which the resulting character is nontrivial and use invariance of the inner product. Thus the identity follows from the self-duality of the stabilizing lattice and unitarity of the finite representation. It does not need a numerical integral of a complicated theta function.

Proposition 8.18 (Vacuum clearing factor)**.** *Under the stated comparison and reciprocity inputs there is a nonzero scalar c such that every finite evaluation of cv is integral at every finite place, and every K-conjugate of every such evaluation has absolute value at most* $\exp(C_0 N^2 h)$ *for a constant C_0 independent of N, h.*

Proof. First choose the theta functional in the vacuum expansion, and take its inner product with v. Since $\theta(F) = 1$ and v has norm one, some $j \in I$ has $\langle v, R_{-x_j} F\rangle \neq 0$. Set

$$\ell(W) = M_\delta^{-1}\langle v, R_{-x_j} W\rangle, \qquad c = |I|\ell(F) \neq 0,$$

where the inner product is linear in the second argument. Substitute this functional into the expansion. The result is

$$cv = \sum_{i\in I} \langle v, R_{-x_j} R_{x_i} v\rangle R_{-x_i} F.$$

By (8.26) and the Heisenberg multiplication law, the coefficients of the sum are zero or roots of unity. There is no surviving denominator M_δ: this is the purpose of the normalization of ℓ. The pointwise integrality of F proved above is preserved by finite translations (their phases are roots of unity). Hence all evaluations of cv are integral; this follows from a finite sum with integral coefficients, without having to decide in advance which evaluation of the vacuum is nonzero.

For the archimedean bound let A' be the area of any normalized conjugate curve and ϑ' its normalized theta section. The value of its torsion translate at zero has absolute value

$$|(U_y\vartheta')(0)| = \exp(-\pi|y|^2/(2A'))|\vartheta'(y)|.$$

This continuous function of y is periodic on its lattice and attains a finite maximum. The family of conjugate normalized curves has only finitely many such maxima: a_1 is algebraic, its conjugates form a finite set, and the cubic with its marked point and normalized differential determines the period lattice and normalized theta section (with the possible finite ambiguity from fourth roots of unity preserving the value of the maximum). Write $C_1 \geq 1$ for a bound

on all of them. The normalized reciprocity comparison and the previous product construction bound each conjugate of each finite translated evaluation of F by $C_1^{N^2h}$, since the product has exactly N^2h degree-one factors. The coefficients in the expression for cv remain of absolute value one or zero under every conjugation. Thus the sum has absolute value at most

$$N^{4fh}C_1^{N^2h} \le \exp\big((4+\log C_1)N^2h\big),$$

where $f = \varphi(N) \le N$ and $\log N \le N$. This proves the bound with $C_0 = 4 + \log C_1$, or any integer above it. □

The scalar chosen in the formal construction is that for rank one, and the standard vacuum in rank h is the product of the rank-one vacuum evaluations. This factorization, `Parameters.intVacuum_finVec_eq_prod` in `New/Curve/IntegralVacuum.lean`, follows from the product definition of the standard diagonal space and its vacuum. Therefore $s^h_{\mathrm{vac}}v$ is integral and its conjugates are bounded by the same expression $\exp(C_0N^2h)$, as asserted by `Parameters.norm_intVacuum_le_one` and `Parameters.norm_intVacuum_le`. The integral module of standard data is the integral span of translates of this scaled vacuum by the specified local lattice. It is stable under this lattice because the finite shift multiplication factors are roots of unity. Stability under the restriction of the finite Weil action, and the equivariance at nonintegral adelic matrix arguments, require the exact representation identities specified in Proposition 9.20 and Proposition 9.22. Neither is an automatic property of a space of power series with integral coefficients. They remain formal-source inputs: `Parameters.omega_mem_standardData`, `finHeisOp_mem_standardData` (the finite Weil vacuum sum), `Parameters.mem_raises_of_jet_bound`, and `mem_raises_of_standardData` in `New/Curve/ConjugatesOfJets.lean`.

8.13 What the arithmetic normalization proves

The preceding proofs establish the normalized good cubic, its exact local height and section algebra, the CM dilation and logarithmic identities, the formal height input from good reduction, the period comparison from the unit-adjusted Chowla–Selberg calculation and character normalization, and both the frame existence and Frobenius orthogonality. The finite realization and vacuum expansion reduce to positive theta space dimension and orthogonality, plus precisely the finite–archimedean comparison identities. Vacuum integrality and its uniform archimedean bound then follow from the normalized reciprocity law. This separates the explicit algebra in the jet theorem from the assumptions on the construction of the Weil model, its normalized CM reciprocity, and its adelic equivariance. Those assumptions are identifiable statements in the supplied source, rather than additional analytic estimates silently inserted into the proof.

Chapter 9

Local CM theta analysis and Schur jets

The preceding chapter (Chapter 8) supplies expanded proofs of the CM arithmetic, local section identities, frame and Frobenius calculations used here. Chapter 12 proves the soft normalized CM reciprocity statement and develops the compatible integral seesaw model. The interface assertions in this chapter are understood with the derivations and remaining hypotheses in both companion chapters.

The contribution of this section is the local lower bound for valuations of the matrix coefficients of a theta kernel. We isolate it from the later transfer of Hecke operators: it is a result in commutative algebra about differential operators, after a particular CM theta function has been put into normal form. The numerical gain which it produces is the weighted prime-power sum. The subsequent section explains how to use the local estimates for actual lifts and repair the obstruction at primes for which two long rows collide.

9.1 The arithmetic local scales

Write v_p for the valuation normalized by $v_p(p) = 1$, and $|z|_p = p^{-v_p(z)}$. At a specified embedding ι into the algebraically closed p-adic field, put $\mathfrak{o} = \{z : |z|_p \le 1\}$ and $\mathfrak{m} = \{z : |z|_p < 1\}$. Let

$$q_p = \begin{cases} p^2, & \chi_K(p) = -1, \\ p, & \chi_K(p) \neq -1, \end{cases} \qquad \rho_p = \frac{1 + \chi_K(p)^2}{2(q_p - 1)}, \qquad \omega_p = \frac{1 - \chi_K(p)}{2}.$$

The values of ρ_p at split, inert and ramified primes are $1/(p-1)$, $1/(p^2-1)$ and $1/(2(p-1))$, respectively. These are the definitions of `nonsplitWeight` in `Lib/General/NumberTheory/QuadFieldSplitting.lean`. The congruence conditions on d ensure that a nonsplit prime is at least 13 (`Parameters.thirteen_le_of_chiK_ne_one`, in `New/Data/Different.lean`). We call an exponent e live at a nonsplit prime if $e \ge 1$ and either e is odd or $p = d$. Put $\mathfrak{p}_e(\sigma) = [-p^e]\sigma$, using the convention of the arithmetic data for the action on places. The normal-form statement retains also $e = 0$ at $p = d$; that term does not appear in the prime sum. Define

$$A_p(y, \sigma) = \omega_p \sum_{e \text{ live}} \left(y/p^e - w_{\mathrm{nom}}(\mathfrak{p}_e(\sigma))\right)_+, \qquad 0 \le y \le a \tag{9.1}$$

and set $A_p = 0$ at split primes. This is `Data.localAsk` in `New/Curve/LiveSum.lean`. The factor ω_p accounts for the ramified prime as well as inert primes. Since $w_{\mathrm{nom}} \ge a/x_1$, the sum is finite.

Lemma 9.1. *For every place σ and $0 \le y \le a$,*

$$A(y,\sigma) = 2\sum_p A_p(y,\sigma)\log p, \qquad \omega_p \sum_{e \text{ live}} p^{-e} = \frac{1}{p-1} - \rho_p \quad (p \text{ nonsplit}).$$

Proof. Expand the von Mangoldt function by prime powers. The factor $1-\chi_K(p^e)$ is 2 at odd powers of an inert prime and vanishes at even powers; it is 1 at all positive powers of the ramified prime, and vanishes at split primes. For the second assertion sum the odd geometric series $p/(p^2-1)$ in the inert case and one half of the full geometric series $1/(p-1)$ in the ramified case. □

The latter identity is `hasSum_live_inv_pow` in `New/Curve/LiveSum.lean`. Introduce

$$F(p) = \sum_\sigma A_p(w_{\mathrm{nom}}(\sigma),\sigma), \qquad g_p = m\left(\sum_{\tau\in\mathcal{T}} A_p(a,\tau) + kF(p)\right). \tag{9.2}$$

Then

$$2\sum_p g_p \log p = m\sum_{\tau\in\mathcal{T}} A(a,\tau) + mk\sum_\sigma A(w_{\mathrm{nom}}(\sigma),\sigma).$$

The integer m appears because there are m kernels, each with a distinguished derivatives at the tops and approximately $kw_{\mathrm{nom}}(\sigma)$ other derivatives at place σ. We shall derive the positive parts in (9.1).

9.2 The CM period and the differential coordinate

The Weil representation on a standard positive definite hermitian space uses a finite lattice normalized by an idele α; in the source the rational vector w has finite image $\alpha^{-1}w$. Write $\Lambda = \Omega_0\mathcal{O}_K$ for the chosen elliptic curve lattice, $\Omega_0 \neq 0$, and ϑ for its normalized theta section, with $\vartheta(0) = 1$. Invariant derivatives are derivatives in the complex linear covering coordinate, not in an integral formal group coordinate. To read their Taylor series in the p-adic field we transfer the coefficients by ι; we use no analytic continuity of that field map.

Set $s_{\mathrm{ar}} = (2\pi/\sqrt{d})^{1/2}$. The source introduces the factors s_0 (`vacuumRoot`), $s_1 = s_0/u_0$ (`periodFactor`), and

$$\mathrm{per} = s_1 s_{\mathrm{ar}}/\Omega_0.$$

Here u_0 is the nonzero normalizing constant of the Chowla–Selberg identity and s_0 is a chosen f-th root of the reciprocal of the finite splitting character evaluated at α. The precise definitions are in `New/Curve/PeriodDefs.lean`. The positive constant in the period identity is

$$\Omega = \frac{2\pi}{\mathrm{rd}_E}\exp(\kappa_K), \qquad \kappa_K = \gamma - \Re\frac{L'}{L}(1,\chi_K).$$

Proposition 9.2 (Formal-source input: the CM normalizations)**.** *For every automorphism σ of $\mathbb{C}$ over K, $|\sigma(s_0)|^2 = \mathrm{rd}_E/\sqrt{d}$. The factor u_0 is the one prescribed by the normalized period formula `CMLattice.chowla_selberg`. For every $\tau \in \Phi$ and p-adic embedding, the norm of $\iota(s_1)$ is the reciprocal of the norm of the component of α read through the embedding $w \mapsto \iota(\tau(\overline{w}))$. Consequently, if $w \in E \cap \alpha\widehat{\mathcal{O}}_E$, then*

$$|\iota(s_1\overline{\tau(w)})|_p \le 1, \qquad |\iota(s_1)|_p = 1 \quad \text{if } p \nmid N_1.$$

The exact first identity is `Parameters.norm_gal_vacuumRoot` in `New/Curve/PeriodDefs.lean`; the Chowla–Selberg formula is `CMLattice.chowla_selberg` in `Lib/CM/FormalGroup/FormalHeight.lean`. The finite norm identity and its consequences are `norm_iota_periodFactor`, `norm_periodFactor_mul_conj_le_one` and `norm_iota_periodFactor_eq_one` in `New/Curve/TightFrameExists.lean`. The derivation is developed in Chapter 8, retaining the precise normalized Weil-character input.

9.3 Integral and logarithmic theta identities

Fix p and ι; all Taylor series of complex functions in the next paragraphs are read through ι. Write $\phi(z)$ for the integral parameter `translatedWpDeriv`, normalized to start with z, and $\xi(z)$ for `translatedWp`, which begins in degree two. Put $\mathfrak{R}_\phi = \mathfrak{o}[[\phi]] \subset \mathbb{C}_p[[z]]$. This denotes the image of the integral series ring by substitution, not the ring of integral coefficients in the differential coordinate.

Proposition 9.3 (Normalized CM identities (see Chapter 8))**.** *The ring $\mathfrak{R}_\phi$ contains ξ and is stable under d/dz. Moreover $\vartheta = \exp(-\mathscr{L}(z))$ with $\mathscr{L}(0) = \mathscr{L}'(0) = 0$ and*

$$\mathscr{L}''(z) = L_2 + \xi(z), \qquad |\iota(\sqrt{-d})L_2|_p \le 1. \tag{9.3}$$

If (c, D) is in the closure of the pairs $(\iota\delta, \iota(\delta\bar{\delta}))$ for $\delta \in \mathcal{O}_K$, then

$$\phi(cz), \quad \exp(D\mathscr{L}(z) - \mathscr{L}(cz)), \quad (D\mathscr{L}(z) - \mathscr{L}(cz))' \ \in \mathfrak{R}_\phi. \tag{9.4}$$

If $0 \neq \delta \in \mathcal{O}_K$, $\delta y \in \Lambda$ and $D = \delta\bar{\delta}$, the twisted theta translate satisfies

$$(U_y\vartheta)(\iota(\bar{\delta})z) = \vartheta(z)^D V(z), \qquad V \in \mathfrak{R}_\phi. \tag{9.5}$$

These statements are `CMLattice.exists_log_theta` in `Lib/CM/KroneckerLimit/LogTheta.lean`, `taylorVia_translatedWp_mem_phiSeriesRing` and `deriv_mem_phiSeriesRing` in `Lib/CM/LocalHeight/IntegralParameter.lean`, and `rescale_translate_eq_theta_pow_mul` in `Lib/CM/LocalHeight/IsogenyIntegrality.lean`. The logarithmic identity follows from differentiation of the theta addition formula; the normalized integrality assertions use the integral model and integrality of torsion evaluations. The arithmetic proofs and their precise classical premises are developed in Chapter 8.

Lemma 9.4 (Crude integrality of jets)**.** *Let $F = \prod_i (U_{y_i}\vartheta)(\sum_j c_{ij}z_j)$ for $\delta y_i \in \Lambda$, $0 \neq \delta \in \mathcal{O}_K$. If $|c_{ij}/(\iota(\bar{\delta})d)|_p \le 1$, then for every polynomial Q with integral coefficients $(Q(\partial)F)(0)$ is integral.*

Proof. Let $D = \delta\bar{\delta}$. By (9.5), each factor comes from $\vartheta^D V$ with $V \in \mathfrak{R}_\phi$ by substituting $z = \sum_j (c_{ij}/\iota(\bar{\delta}))z_j$. Let $\mathfrak{R}$ be the $\mathfrak{o}$-algebra generated by the images under these substitutions of $\mathfrak{R}_\phi$, $\mathscr{L}'$, and ϑ for each i. All constants are integral. By the chain rule, a partial derivative of a generator is an integral multiplier times the substituted series obtained by applying d times the derivative in z. For $\mathscr{L}'$ the latter is $d(L_2 + \xi)$, in which dL_2 is integral; for ϑ use $\vartheta' = -\mathscr{L}'\vartheta$. The derivative stability of $\mathfrak{R}_\phi$ treats the remaining generators. Thus $\mathfrak{R}$ is stable under all partial derivatives. It contains F and hence all its iterated derivatives. Evaluating at zero and taking integral linear combinations proves the assertion. □

This is `CMLattice.jet_theta_prod_integral` of `New/Curve/FormalParameter.lean`. It applies also at split primes.

9.4 The regularized CM parameter at nonsplit primes

Suppose p is nonsplit. Write $q_p = p^\epsilon$, with $\epsilon = 2$ inert, $\epsilon = 1$ ramified, and $Q_0 = q_p - 1$. Thus $\omega_p = \epsilon/2$, $\rho_p = \omega_p/Q_0$. Choose a primitive Q_0-th root ζ in the p-adic field. The assertion that $(\zeta^j, \zeta^{j(p+1)})$ belongs to the norm-pair closure in Proposition 9.3 is the formal-source input `mem_normPairClosure_of_pow_eq` in `Lib/CM/KroneckerLimit/LogTheta.lean`. Average the parameter:

$$T(z) = Q_0^{-1} \sum_{j=0}^{Q_0-1} \zeta^{-j} \phi(\zeta^j z).$$

By (9.4) it lies in $\mathfrak{R}_\phi$, is tangent to the identity, and has support in degrees $1 \pmod{Q_0}$. It is an integral generator of the ring, since as a series in ϕ it has linear coefficient one and its inverse in that coordinate is integral by recursion. Let Z denote its inverse in the differential coordinate and write $Z(u) = u + z_{q_p} u^{q_p} + \cdots$. Uniqueness of compositional inverses implies $Z(\zeta u) = \zeta Z(u)$. For every $F \in \mathfrak{R}_\phi$ the series $F(Z(u))$ is integral. Differentiation of $T(Z(u)) = u$ therefore yields

$$Z'(u) \in \mathfrak{o}[[u]], \tag{9.6}$$

since $T' \in \mathfrak{R}_\phi$ has constant term one.

Take $c = p$ in the inert case and $c = \sqrt{-d}$ in the ramified case. Its norm-square is q_p, its p-adic valuation is ω_p, and the isogeny has degree q_p. The integral series $[c](u) = T(\iota(c)Z(u))$ is equivariant for the above roots of unity and begins with the nonunit $\iota(c)u$.

Lemma 9.5 (Degree of the reduced isogeny (see Proposition 8.9))**.** *Under these hypotheses $[c](u)$ has a unit coefficient in degree at most q_p.*

The source is `CMLattice.exists_unit_coeff_degree_le` in `Lib/CM/FormalGroup/FormalHeight.lean`, the odd-degree CM isogeny height input for such an integral coordinate. Since all intermediate coefficients vanish, the coefficient in degree q_p is a unit. Comparing through that degree in the composition gives $(\iota(c) - \iota(c)^{q_p}) z_{q_p}$, and hence $|z_{q_p}|_p = p^{\omega_p}$. Choose ϖ with $\varpi^{Q_0} = -z_{q_p}^{-1}$ and put $\pi = \varpi^{Q_0}$; then

$$|\varpi|_p = p^{-\rho_p}, \qquad |\pi|_p = p^{-\omega_p}. \tag{9.7}$$

Lemma 9.6. *There is an integral series R for which*

$$\varpi^{-1} Z(\varpi u) = u - u^{q_p} + \pi R(u).$$

Its compositional inverse S is integral tangent to the identity. In particular for an integral series E_0,

$$T(\varpi z) = \varpi S(z), \qquad S - S^{q_p} = z + \pi E_0(z). \tag{9.8}$$

Proof. The first two nonzero coefficients are as stated by the choice of ϖ. In degree $n = 1 + Q_0 j$, $j \geq 2$, (9.6) gives $v_p(z_n) \geq -v_p(n)$. After scaling the valuation of this coefficient is at least $j\omega_p - v_p(n)$. If $p \nmid n$ this is at least ω_p; otherwise $j \equiv 1 \pmod{p}$ and $j \geq p+1$. Then $n \leq p^\epsilon j$, and $v_p(n) \leq \epsilon + \log_p j \leq \epsilon(j-1)/2$, the last inequality holding for $p \geq 13$, $j \geq p+1$ by comparison at $j = p+1$ and monotonicity of the difference. This proves integrality of the remainder after division by π. Invert the integral series recursively and compose with this inverse, taking $E_0 = -R \circ S$. □

Thus the invariant differential is regularized at scale $p^{-\rho_p}$. The gap between $1/(p-1)$ and ρ_p is the source of the local gain.

9.5 The exceptional exponential factor

Let $c_2 = [z^2]\mathscr{L}$ and put

$$M(u) = (\mathscr{L}(z) - c_2 z^2)\big|_{z=Z(u)}, \qquad n(u) = (\mathscr{L}'(z) - 2c_2 z)\big|_{z=Z(u)}.$$

Then $M = O(u^3)$, $n(0) = 0$, and

$$M' = nZ', \qquad n' = (\xi \circ Z)Z' \in \mathfrak{o}[[u]], \qquad Z' \in \mathfrak{o}[[u]].$$

For a series G, let $G_{[j]}$ denote the part in degrees $j \pmod{Q_0}$. Since Z is equivariant for the Q_0-th roots of unity, this projection commutes with substitution of Z. Define $\kappa = \varpi^{p+1}/p$ and $\mu(u) = \kappa^{-1} M_{[p+1]}(\varpi u)$.

Lemma 9.7. *All coefficients of μ are integral. Modulo $\pi\mathfrak{o}[[u]]$, only its terms of degrees $p+1$ and $p+q_p$ may be nonzero.*

Proof. Since M vanishes through degree two, the possible degrees are $N = p+1+Q_0 I$, $I \geq 0$ (when $\epsilon = 1$ the degree-2 term of the congruence class vanishes). The coefficient of M in degree N is

$$\frac{1}{N} \sum_j n_{p+Q_0(I-j)}(1+Q_0 j) z_{1+Q_0 j},$$

where $j \geq 0$ and the subscript $i = p + Q_0(I-j)$ of n_i is positive. Thus $0 \leq j \leq I+1$ when ramified, and $0 \leq j \leq I$ when inert. By integrality of n', Z' a summand has valuation at least $-v_p(i)$, with an additional ω_p when $j = 1$ because $v_p(q_p z_{q_p}) = \epsilon - \omega_p = \omega_p$. Scaling by $\kappa^{-1}\varpi^N = p\pi^I$ makes its valuation at least

$$1 + \omega_p I - v_p(N) - v_p(i) + \omega_p \mathbf{1}_{j=1}. \tag{9.9}$$

For $I = 0$, $v_p(N) = 0$ and $v_p(i) \leq 1$, so the expression is nonnegative. For $I = 1$ one has $N = p + q_p$ and $v_p(N) = 1$. At $j = 1$, $i = p$, and (9.9) is $2\omega_p - 1 \geq 0$; otherwise $p \nmid i$. If $2 \leq I < p$, $N \not\equiv 0 \pmod{p}$ and the only divisible i is $i = p$ at $j = I$, giving valuation at least ω_p. For $I \geq p$, both positive integers N, i are at most $p^\epsilon(I+2)$. It follows that

$$v_p(N) + v_p(i) \leq 2\epsilon + 2\log_p(I+2) \leq 1 + \omega_p(I-1).$$

The last inequality holds at $I = p \geq 13$ since $\log_p(p+2) < 2$, and persists as I grows because the right hand side increases faster than the left. Thus (9.9) is at least ω_p whenever $I \geq 2$. The non-Archimedean triangle inequality finishes the proof. □

This reconstructs the valuation calculation of `norm_coeff_le_of_deriv` in `New/Curve/FormalParameter.lean`. The estimate works because the coefficients which can cause two divisions by p simultaneously fall into the first two selected degrees; in the second selected degree the exceptional coefficient of Z' is small enough.

Define $\Psi_0(z) = \kappa^{-1}\mathscr{L}_{[p+1]}(\varpi z)$. Using (9.8), the contribution of the non-quadratic part is $\mu(S(z))$. The quadratic term is retained only when $p = d$, in which case its coefficient is $u_0 = \kappa^{-1}\varpi^2 c_2$. From (9.3), $v_p(c_2) \geq -1/2$ (here 2 is a unit); hence $v_p(u_0) \geq 1 - \omega_p - 1/2 = 0$. Let $u_0 = 0$ in the inert case. We obtain

$$\Psi_0(z) = u_0 z^2 + u_1 S(z)^{p+1} + u_2 S(z)^{p+q_p} + \pi R_0(z), \tag{9.10}$$

with integral coefficients and $\Psi_0(0) = 0$. Moreover $|\kappa|_p = p^{\omega_p - 1/(p-1)}$, as is checked in the two cases by substituting ρ_p in $v_p(\kappa) = (p+1)\rho_p - 1$.

The formal source additionally proves that u_1 is a unit at inert primes. It uses the following formal-group input: if every $F \in \mathfrak{R}_\phi$ has $|[z^{p-1}]F'|_p < 1$, then $|[z^{p-1}]\xi|_p \geq 1$. This is `CMLattice.one_le_norm_coeff_translatedWp_nonsplit` of `Lib/CM/FormalGroup/FormalHeight.lean`. The required hypothesis follows from $T(z) = z$ through degree $p < q_p$: an integral series in T has integral coefficient in degree p, so its derivative coefficient in degree $p-1$ is divisible by p. Equation (9.3) now gives $|[z^{p+1}]\mathscr{L}|_p \geq p$; multiplication by $\kappa^{-1}\varpi^{p+1} = p$ makes the corresponding coefficient of Ψ_0 of norm at least one. The coefficient estimate gives the reverse inequality. This proves the unit assertion conditional on that formal-height input.

Proposition 9.8 (Local normal form for torsion translates)**.** *If $p \nmid n$ and $ny \in \Lambda$, there are a unit N and a constant κ_n in the p-adic field satisfying $|\kappa_n|_p = p^{\omega_p - 1/(p-1)}$ and an integral series V_y such that*

$$(U_y\vartheta)(\varpi z) = \big(\exp(\kappa_n \Psi_0(z))V_y(z)\big)\big|_{z\mapsto Nz}.$$

The same Ψ_0, N, κ_n work for all y with the given n.

Proof. By (9.4) for the norm pairs of the powers of ζ, the series $\exp(\zeta^{j(p+1)}\mathscr{L}(z) - \mathscr{L}(\zeta^j z))$ is integral in T. Replace z by $\zeta^{-j}z$ (which preserves the ring) and multiply these series over all j. Character orthogonality makes their product $\exp(Q_0(\mathscr{L}_{[p+1]} - \mathscr{L}))$. Apply (9.5) with $\delta = nQ_0$ to get

$$(U_y\vartheta)(nQ_0 z) = \exp(-(nQ_0)^2\mathscr{L}(z))V(z), \qquad V \in \mathfrak{R}_\phi.$$

Replace $\mathscr{L}$ by $\mathscr{L}_{[p+1]}$ in the exponential: the residual product just obtained, raised to the integer power n^2Q_0, is integral in T. Scaling by ϖ then gives integral coefficients for this product and V, since $T(\varpi z) = \varpi S(z)$. Finally take $N = (nQ_0)^{-1}$, $\kappa_n = -(nQ_0)^2\kappa$. This proves the required normal form. □

Equation (9.8), (9.10) and the proposition above give the nonsplit part of `CMLattice.exists_formal_parameter` in `New/Curve/FormalParameter.lean`.

9.6 Tight frames and theta vectors

The theta kernel on a space of matrices is not an unrelated product of one-variable theta functions. The product is built from a tight frame in the CM lattice.

Definition 9.9. A tight frame of length M is a finite collection of vectors $w_i \in E \cap \alpha\widehat{\mathcal{O}}_E$ with nonnegative integer multiplicities m_i such that

$$\sum_i m_i \sigma(w_i)\overline{\tau(w_i)} = M\delta_{\sigma\tau} \quad (\sigma, \tau \in \Phi).$$

After repetition of elements we may assume each multiplicity is one.

This is `IsTightFrame` in `Lib/Weil/Construction/ConjugateJetDefs.lean`. For instance the N_1 powers of ζ_{N_1} form a tight frame of length N_1: their integrality follows from the different normalization of α, and the frame equality is the geometric sum of $(\sigma(\zeta_{N_1})/\tau(\zeta_{N_1}))^j$. This is `Parameters.isTightFrame_zeta_pow`. To ensure good p-adic integrality of all theta vectors one needs lengths that are prime to p, and for the lower bounds also frames having some vector attaining all the finite norm bounds of Proposition 9.2.

Proposition 9.10 (Large frames with a unit vector (see Proposition 8.14))**.** *For every $R \geq 0$, every integer length $M \geq (dN_1)^6(R+1)$, and each p-adic embedding ι, there is a tight frame with multiplicities $m_i \geq R$ containing a vector w_j with $|\iota(s_1\overline{\tau(w_j)})|_p = 1$ for all τ.*

The precise arithmetic existence result is `Parameters.exists_isTightFrame`, in `New/Curve/TightFrameExists.lean`. It constructs additional frames using the cyclotomic different generator and solves for the desired length using the coprimality of two frame lengths. The general arithmetic construction is proved in Proposition 8.14 of the preceding chapter. For upper bounds take $M = p(dN_1)^6 + 1$, large enough and prime to p.

For a standard positive definite hermitian space V of rank h, denote by $\mathscr{F}_V$ the finite part of its Weil representation with Heisenberg operator ϱ and theta functional ϑ_V. The Taylor model of $F \in \mathscr{F}_V$, at differential scale c, is the series $\mathcal{T}_c F$ with coefficients

$$[z^\beta]\mathcal{T}_c F = \frac{c^{|\beta|}}{\beta!}\vartheta_V(v_\beta \otimes F) \qquad (\beta \in \mathbb{N}^{\Phi\times\{1,\dots,h\}}), \tag{9.11}$$

where v_β is the Fock vector of the monomial and $\beta! = \prod_x \beta_x!$. This is `Space.jetSeries` in `Lib/Weil/Construction/DataProduct.lean`. Evaluating a monomial differential operator introduces $\beta!$ once again. No factorial may be lost by confusing the Taylor coefficient with the invariant derivative.

Let $\delta \in \mathcal{O}_K$ satisfy $\delta\bar{\delta} = D_\delta > 0$, and assume $v \in \Lambda$, with the curve's area $|v|^2 M\sqrt{d}/(2D_\delta)$. For a frame define linear maps on the complex coordinates of V by

$$A_{ic}(z) = \frac{v}{\bar{\delta}s_{\mathrm{ar}}}\sum_\tau \overline{\tau(w_i)}z_{\tau,c}.$$

The sum of pullbacks along A_{ic} of the normalized curve's hermitian form, counted with multiplicity m_i, equals the standard form: expand the scalar products of $A_{ic}(z)$ and $A_{ic}(z')$ and use Definition 9.9. The factor $|v/\bar{\delta}|^2 M$ is exactly absorbed by the area normalization. In particular the product of the pulled-back theta sections has the appropriate transformation law.

Proposition 9.11 (Formal-source input: theta product and finite vector)**.** *Under these frame and area choices, a product of the pullbacks of degree-m_i theta monomial sections is the entire theta function of some vector $F \in \mathscr{F}_V$. For rational Heisenberg translates this theta function is the product of the translated curve sections with arguments prescribed by the same linear forms. All finite point evaluations of F are integral at all p-adic embeddings. There is a finite vacuum expansion*

$$\sum_{j\in I} M_\delta \ell(\varrho(\alpha x_j)\,\mathrm{vac})\varrho(-\alpha x_j)F = D_\delta^{2fh}\ell(F)\,\mathrm{vac}$$

for every complex linear functional ℓ, where M_δ is a natural number, $|I| = D_\delta^{2fh}$, and δx_j is integral at the finite places.

This is `Parameters.exists_finitePart_of_isTightFrame`, in `New/Curve/FrameSections.lean`; its vacuum expansion is the structure `Space.VacuumExpansion` of `Lib/Weil/Construction/ConjugateJetDefs.lean`. The finite vector reconstruction and its integrality are formal-source inputs here. The identity of the corresponding Taylor series follows by evaluating along every line and differentiating at zero (homogeneous polynomials are determined by their values on all lines). The corresponding lemmas are `Space.taylor_line_eq_subst_jetSeries` and `Space.jetSeries_eq_prod`. The source also proves `Space.eq_zero_of_jetSeries_eq_zero` by the identity theorem for entire functions.

9.7 Frobenius pairings in tight-frame products

Fix a nonsplit p. Consider a frame product of type `Parameters.thetaProducts` with all multiplicities repeated, frame length M prime to p, torsion denominators n prime to p, and

multiplier b' of p-adic norm at most one (not the scale b). Its factors for pairs (i,c) are

$$(U_{y_{ic}}\vartheta)\left(\sum_{\tau}\iota(b'\overline{\tau(w_i)})z_{\tau,c}\right), \qquad ny_{ic}\in\Lambda.$$

At first there is no extra denominator from K. By Proposition 9.8, after substituting ϖz the product has the form

$$F(\varpi z)=\exp(\kappa_n\Psi(z))V(z), \quad V\in\mathfrak{o}[[z]], \quad \Psi(0)=0, \tag{9.12}$$

where $\Psi=\sum_{i,c}\Psi_0(L_{ic})$, $L_{ic}=\sum_\tau r_{\tau i}z_{\tau,c}$, and $r_{\tau i}=N\iota(b'\overline{\tau(w_i)})$. At a nonsplit prime, $p\nmid N_1$, so Proposition 9.2 gives $r_{\tau i}\in\mathfrak{o}$.

Lemma 9.12 (Frobenius orthogonality (see Proposition 8.16))**.** *In characteristic p, modulo $\pi\mathfrak{o}$, the frame matrix satisfies, for any live exponent e and for $e=0$ when $p=d$,*

$$\sum_i r_{\sigma i}^{p^e}r_{\tau i}=(N\iota(b'))^{p^e+1}M\,\mathbf{1}_{\tau=\mathfrak{p}_e(\sigma)}. \tag{9.13}$$

Here the multiplier is read in the same reduction, with N regarded as the p-adic element from the normal form.

The formula is `Parameters.exists_frame_rows_frob` in `New/Curve/ThetaProducts.lean`; more precisely it holds under every map to characteristic p killing all elements of valuation at least ω_p. The proof uses the Frobenius permutation of the cyclotomic embeddings and the frame identity. At an inert prime the odd Frobenius interchanges the CM embeddings, explaining the minus sign in $[-p^e]$; at the ramified prime complex conjugation acts the same way modulo the local uniformizer. These arithmetic identifications of embeddings modulo the indicated ideals are used as input.

For a finite truncation degree $D>0$ let $I_D=\{0\le e\le\lfloor\log_p D\rfloor : e \text{ odd or } p=d\}$. For integral series $Z_{\tau,c}$ with zero constant term in any finite variable set, suppose the pairings

$$\sum_{\tau,c}Z_{\tau,c}^{p^e}Z_{\mathfrak{p}_e(\tau),c}=0 \quad (e\in I_D) \tag{9.14}$$

hold after reduction modulo π and terms of degree $>D$. We shall prove they kill the exceptional exponent in (9.12) and render the exponential integral through degree D.

Lemma 9.13 (Artin–Schreier cancellation)**.** *Let R be a commutative ring of characteristic p, and let finitely many elements t_i, L_i satisfy $t_i-t_i^{p^\epsilon}=L_i$. Assume each t_i is nilpotent, and $\sum_i L_i^{p^j+1}=0$ for all odd j if $\epsilon=2$, or all $j\ge 0$ if $\epsilon=1$. Then $\sum_i t_i^{p+1}=\sum_i t_i^{p+p^\epsilon}=0$.*

Proof. Iterate the given equation:

$$t_i=\sum_{a=0}^{A-1}L_i^{p^{\epsilon a}}+t_i^{p^{\epsilon A}}.$$

Nilpotence kills the last term for large A. Expanding $\sum_i t_i^p t_i$ by Frobenius gives terms $\sum_i L_i^{p^u}L_i^{p^v}$ for $u=\epsilon a+1$, $v=\epsilon b$. Each equals $(\sum_i L_i^{p^{|u-v|}+1})^{p^{\min(u,v)}}=0$; indeed the exponent difference is odd if $\epsilon=2$. Expand $\sum_i t_i^p t_i^{p^\epsilon}$ in the same way, taking $v=\epsilon(b+1)$. This proves both identities without requiring that R be reduced or perfect. □

This is `sum_pow_eq_zero_of_null` in `New/Curve/ThetaProducts.lean`.

Proposition 9.14 (Null-cone integrality)**.** *Whenever the integral zero-constant series $Z_{\tau,c}$ satisfy (9.14), the coefficients of $F(\varpi Z)$ through degree D are integral.*

Proof. Work in the quotient of the ring of integral series by the ideal of coefficients divisible by π together with terms of degree $> D$. For $t_{ic} = S(L_{ic}(Z))$, (9.8) gives $t_{ic} - t_{ic}^{p^e} = L_{ic}(Z)$, and t_{ic} is nilpotent by degree. Using (9.13) we have

$$\sum_{i,c} L_{ic}(Z)^{p^e+1} = (N\iota(b'))^{p^e+1} M \sum_{\tau,c} Z_{\tau,c}^{p^e} Z_{\mathfrak{p}_e(\tau),c} = 0$$

for $e \in I_D$. Larger e give zero automatically by degree. Lemma 9.13 shows that the two S-powers in (9.10) vanish after summing; at $p = d$ the quadratic terms vanish by the same computation with $e = 0$. It follows that the coefficients of $\Psi(Z)$ through degree D are divisible by π. The exponential in (9.12) is integral through this degree because

$$v_p((\kappa_n \pi)^j / j!) = j/(p-1) - v_p(j!) \geq 0.$$

There is no formal convergence problem, since the substituted series have zero constant term and we work in a finite truncation. Multiplication by $V(Z)$ preserves integrality. □

This is the first conclusion of `thetaProducts_normal_form` in `New/Curve/ThetaProducts.lean`, whose hypothesis is the predicate `IsIntegralOnNullCone`. In the source the couplings are required to vanish under every characteristic-p map annihilating π; it suffices in the proof to apply that assumption to the indicated quotient.

In the inert case, when b' has norm one, one also knows which coefficients of the exponent are units. Reducing in the residue-field series ring and differentiating (9.10), $u_0 = 0$ and the term with u_2 has zero derivative. Since $S - S^{p^2} = z$, $S' = 1$ in this reduction, and

$$\frac{\partial \overline{\Psi}}{\partial z_{\tau,c}} = \bar{u}_1 \sum_i r_{\tau i}\, t_{ic}^{p}, \qquad t_{ic} = \sum_{a \geq 0} L_{ic}^{p^{2a}}.$$

By the frame identity a coefficient is nonzero exactly at $z_{\sigma,c}^{p^e}$ with odd e and $\mathfrak{p}_e(\sigma) = \tau$. It is a unit because M, Nb', u_1 are units. No coefficient supported on several variables occurs. This proves `coeff_sum_frob_ne_zero_iff` and explains the sharpness assertion of `thetaProducts_normal_form`. The same pairings which reduce the guaranteed upper divisibility are genuinely present in the series. This fact is also used in the lower volume estimates.

9.8 Schur matrices and determinant losses

Let R_1, R_2 be the source and target ranks, $N = R_2$, and let $\kappa_\sigma(h)$ count columns of height $h < N$ in the prescribed tableau. A filling assigns h row indices to each column of height h; its polynomial $b_i(Z)$ is the product of minors in the indicated rows and first h columns. Put $D = \sum_\sigma \sum_{h<N} h\kappa_\sigma(h)$. The jet estimate must hold for every Schur matrix coefficient, not just for the highest weight vector.

Proposition 9.15 (Formal-source input: integral Schur expansion)**.** *For a weight satisfying the rank condition `FitsRanks` the Schur matrix $S_\kappa(X)$ satisfies*

$$b_{i'}(XY) = \sum_i S_{i'i}(X) b_i(Y) \prod_\sigma \det(Y_\sigma)^{d_\sigma}, \tag{9.15}$$

where the determinants use the first R_1 columns of the variable matrix Y_σ, and $d_\sigma = \kappa_\sigma(R_1)$ if $R_1 < N$, zero otherwise. This holds over every ring and commutes with base change; the coefficient coordinate maps in the standard tableau basis are integral on the associated Borel semi-invariants.

This is the standard integral tableau theorem as used by `PolyRep.sum_schur_smul_nablaBasis` and `schur_map` in `Lib/GLn/SchurModule/SchurMatrix.lean`. Its basis and weight are `nablaBasis` and `diagramWeight` in `Lib/General/Algebra/ZerosCostDegree.lean`; the integral basis theorem is in `Lib/GLn/SchurModule/MinorBasis.lean`. Write $J_Q(F) = Q(\partial)F(0)$ and $\mathscr{J}_\kappa(F)_{i'i} = J_{S_{i'i}}(F)$ (`schurJet` in `Lib/Weil/Construction/DataProduct.lean`).

Let ev be a finite set of places whose height-one couplings will be exempted. Define

$$a_\sigma^{ev} = \kappa_\sigma(1)\mathbf{1}_{\sigma\notin ev}, \quad c_\sigma = \sum_{\substack{h<N\\ h\neq 1}} \kappa_\sigma(h), \quad D_\sigma = \sum_{\substack{h<N\\ h\neq 1}} h\kappa_\sigma(h).$$

For the ordered coupling $\sigma, \tau = \mathfrak{p}_e(\sigma)$ put

$$\mathfrak{C}_{p,e}(\sigma,\tau) = \min(a_\sigma^{ev}/p^e, c_\tau + a_\tau^{ev}) + c_\sigma/p^e + \min(D_\sigma/p^e, D_\tau). \tag{9.16}$$

This is `couplingCount` in `New/Curve/LiveSum.lean`. Its three terms distinguish the three types of pairs of columns.

Lemma 9.16 (Simultaneous tilting)**.** *For any integral matrix arrays $Y_{\sigma,c,j}$ with source rows c and column indices $j \geq 0$, there are integral $h \times h$ matrices $M_{\sigma,h}$, $1 \leq h < N$, as follows. Write $Y_\sigma^{(h)}$ for the transposed array of the first h columns. For each $e \in I_D$, $\tau = \mathfrak{p}_e(\sigma)$, and rows u, u' from $M_{\sigma,h}Y_\sigma^{(h)}$, $M_{\tau,h'}Y_\tau^{(h')}$,*

$$\sum_c u_c^{p^e} u'_c \equiv 0 \pmod{\pi}$$

unless $h = 1$ and either $\sigma \in ev$ or $(h' = 1, \tau \in ev)$. Moreover

$$\prod_{\sigma,h} |\det M_{\sigma,h}|_p^{\kappa_\sigma(h)} \geq p^{-\omega_p \sum_{e\in I_D}\sum_\sigma \mathfrak{C}_{p,e}(\sigma,\mathfrak{p}_e(\sigma))}. \tag{9.17}$$

Proof. An elementary observation about valuation rings: if coefficient vectors in $\mathfrak{o}^h$ must satisfy finitely many conditions $L_i(v) \in u_i\mathfrak{o}$ with integral linear functionals L_i and nonzero $u_i \in \mathfrak{o}$, there is a matrix of h admissible rows of determinant valuation at most $\sum_i v_p(u_i)$. Inductively a unimodular change on the current coefficient basis makes the next linear functional depend on just one coordinate: choose among its coefficients one of largest norm and use divisibility to eliminate the others. Multiplying that coordinate by u_i suffices; earlier conditions are maintained by linearity. Requiring all h coordinates to be divisible by u_i instead costs $hv_p(u_i)$.

Fix e, σ, τ as in the statement. For source height one, the untransformed row y is fixed. If $\sigma \notin ev$, either scale it by an integral p^e-th root of π, of valuation ω_p/p^e, or impose the single congruence $\sum_c y_c^{p^e} u'_c \equiv 0 \pmod{\pi}$ on each target height except the exempt height-one row. The weighted costs are $\omega_p a_\sigma^{ev}/p^e$ and $\omega_p(c_\tau + a_\tau^{ev})$; choose the cheaper. If $\sigma \in ev$ there are no conditions of this type.

For source height greater than one and target height one, choose integral p^e-th roots of the entries of the untransformed target row y' and require

$$\sum_c (y'_c)^{1/p^e} u_c \in \pi^{1/p^e}\mathfrak{o}.$$

Taking p^e-th powers in characteristic p gives the required congruence modulo π. Further transformation of y' is scalar multiplication and preserves this relation. The weighted cost is $\omega_p c_\sigma/p^e$. Finally, for both heights greater than one, annihilate all pairings by scaling all coefficient rows on either the source or the target side, at the smaller cost $\omega_p \min(D_\sigma/p^e, D_\tau)$. All the conditions are on the original coefficient coordinates or the fixed linear forms specified above. Impose them simultaneously by the first observation; integral row operations preserve previously zero pairings, since the quotient has characteristic p. The costs add to the exponent of (9.17). □

This is the argument of `exists_null_coupling_frames` in `New/Curve/JetDivisibilit y.lean`. The source implements the valuation-ring observation with `IndexDvd.of_det_dvd`. In particular the matrices are not claimed unimodular; their determinant costs must be subtracted.

For each box with height h and row label a prescribe a nonnegative integer charge $g(\sigma, h, a)$. Require

$$g(\sigma, 1, a) > 0 \quad \text{or} \quad g(\mathfrak{p}_e(\sigma), h, a) > 0$$

whenever $\sigma \in ev$ or $(h = 1, \mathfrak{p}_e(\sigma) \in ev)$. For a tableau i' write $G_{i'}$ for the total charges of its boxes, with multiplicity. The charges will kill the exempt pairings.

Theorem 9.17 (Schur-jet bound)**.** *Suppose F satisfies the null-cone property of Proposition 9.14 for degree $D > 0$ and the scales (9.7). Then*

$$\begin{aligned} v_p(\mathscr{J}_\kappa(F)_{i'i}) \geq (1/(p-1) - \rho_p)D - \omega_p \sum_{e \in I_D} \sum_{\sigma} \mathfrak{C}_{p,e}(\sigma, \mathfrak{p}_e(\sigma)) \\ - fN^3(\lfloor \log_p D \rfloor + 1) - G_{i'}. \end{aligned} \tag{9.18}$$

Proof. Apply the jet functional in X to (9.15). The integral basis property converts any coefficient bound on the polynomial $J_{b_{i'}(XY)}(F)$ in Y to a bound on the entries of the Schur jet. The largest coefficient norm of a polynomial over an algebraically closed valued field is attained at some integral specialization: divide by a largest coefficient, reduce modulo the maximal ideal, and choose a point of the infinite residue field where the resulting nonzero polynomial does not vanish. Thus it suffices to bound the expression for all specialized integral Y.

Apply Lemma 9.16. Substitute in the theta series fresh variables $u_{\sigma,h,j,a}$, $0 \leq j < h < N$, $1 \leq a \leq N$, via

$$X_{\sigma,a,c} = \varpi \sum_{h,j} p^{g(\sigma,h,a)} (M_{\sigma,h} Y_\sigma^{(h)})_{j,c} u_{\sigma,h,j,a}.$$

Remove the scaling factor ϖ from the substitution. Expanding the Frobenius pairings of matrix coordinates (a, c) in characteristic p gives exactly the pairings of tilted coefficient rows. Unexempted terms vanish, and exempt terms vanish by their positive charges since $p \in \pi\mathfrak{o}$. The null-cone condition gives integral Taylor coefficients of the substituted theta series through degree D.

Now apply the product of specified minors in the fresh variables as a differential operator. It is homogeneous of degree D, with integral coefficients. There are at most $fN \sum_{h<N} h \leq fN^3$ fresh variables. A monomial u^β of degree D contributes a factor $\beta!$; the base-p digit formula gives

$$v_p(\beta!) = (D - \sum_i s_p(\beta_i))/(p-1) \geq D/(p-1) - fN^3(\lfloor \log_p D \rfloor + 1).$$

The same bound holds for the sum of the monomial pairings. The chain rule identifies that sum with

$$\varpi^{D}p^{G_{i'}}\left(\prod_{\sigma,h}\det M_{\sigma,h}^{\kappa_\sigma(h)}\right)J_{b_{i'}(XY)}(F).$$

Within each minor the substituted differential rows are transformed by the same $M_{\sigma,h}$, yielding its determinant, and each box contributes ϖ and the power of p prescribed by its charge. Removing these factors subtracts $\rho_p D$, $G_{i'}$, and at most the determinant loss (9.17). This proves (9.18). □

This is `norm_schurJet_le` in `New/Curve/JetDivisibility.lean`. The factorial gain is diminished by the coordinate and determinant costs, and the explicit logarithmic factorial correction. The minima in (9.16) are essential: replacing them by the more expensive alternatives would eliminate the positive-part terms in (9.1).

9.9 Application to the prescribed weights

For the arithmetic data $\mathcal{D}$, let $R_j = k + (j+1)(t+1)$, $0 \le j < m$, be the target ranks of the definite chain. Put $w_\sigma = w_{\mathrm{nom}}(\sigma)$,

$$a_\sigma = a\mathbf{1}_{\sigma\in\mathcal{T}},\quad c_{j,\sigma} = \sum_{\substack{h<R_j\\ h\neq 1}}\kappa_\sigma(h),\quad D_{j,\sigma} = \sum_{\substack{h<R_j\\ h\neq 1}}h\kappa_\sigma(h),\quad D_j = \sum_\sigma (a_\sigma + D_{j,\sigma}).$$

The weight multiplicities $\kappa_\sigma(h)$ are `Data.κ` in `New/Curve/LiveSum.lean`: height one occurs at tops with multiplicity a; height k has multiplicity $w_1(\sigma) - \lfloor (k+t+1)/2\rfloor$; heights above k congruent to k modulo $t+1$ have multiplicity $2h_2 - (t+1)$; all other multiplicities vanish. These expressions are nonnegative by admissibility.

Lemma 9.18 (Weight counts)**.** *For $0 \le j < m$ one has*

$$kw_\sigma \le D_{j,\sigma} \le kw_\sigma + x^{60},\qquad 0 \le c_{j,\sigma} \le w_\sigma + x^{60}.$$

Proof. The multiplicity of height k is at least w_σ because the rounded weight w_1 includes the gap $3h_2$. The difference from w_σ is at most the rounding step plus $3h_2$. There are at most $j+1$ higher heights, each bounded by R_j and of multiplicity at most $2h_2$. Thus the excess box count is at most $k(N_1 + 3h_2) + 2(j+1)R_j h_2$ (here N_1 in the rounding term denotes the cyclotomic order, not a target rank). With $k \ll x^5$, $m \ll x^{1/2}$, $R_j \ll x^{11/2}$, $h_2 = r^8 \ll x^{40}$ and the cyclotomic bound, this is smaller than x^{60} on the admissible asymptotic range. The column count is bounded by ignoring the heights. The formal source establishes these comparisons for every admissible datum as `Data.lowColumns_lowBoxes_bounds` in `New/Columns/JetExponentBound.lean`. □

For a nonsplit prime $p < x_1$ define the exponent

$$E_{p,j}(ev) = \left(\frac{1}{p-1} - \rho_p\right)D_j - \omega_p\sum_{e\in I_{D_j}}\sum_\sigma \mathfrak{C}_{p,e}(\sigma,\mathfrak{p}_e(\sigma)) - fR_j^3(\lfloor\log_p D_j\rfloor + 1). \tag{9.19}$$

This is `Data.jetExponent` in `New/Columns/JetExponentBound.lean`. Define

$$\mathrm{coll}_p(ev) = \sum_\sigma\sum_{e\ \mathrm{live}}(a/p^e)\mathbf{1}_{\sigma,\mathfrak{p}_e(\sigma)\in\mathcal{T}\setminus ev}.$$

The sum converges by geometric decay but need not have finite support. It records the loss when a long row has a long partner.

Proposition 9.19 (Prime-power gain of one jet)**.** *For nonsplit $p < x_1$ and $0 \le j < m$,*

$$E_{p,j}(ev) \ge \frac{g_p}{m} - \mathrm{coll}_p(ev) + \omega_p|\mathcal{T} \cap ev|\, a/x_1 - \frac{fa(\log b)^2}{pm}. \tag{9.20}$$

Proof. We give the comparison on the asymptotic admissible range, sufficient for the main theorem; the formal source proves it uniformly for all admissible data. Set $u = a/(pm)$. The separated scales imply $x^{60}, R_j^3 \le u$, $c_{j,\sigma}/p \le u$, and $5(\lfloor \log_p D_j \rfloor + 1) \le (\log b)^2$. For example $u \ge xb/(x_1 m)$ and $x_1 \le b^{10^{-6}}$ control all polynomial terms. The column bound and $m/x \to 0$ give $c_{j,\sigma}/p \le (a/x + x^{60})/p \le u$. The last inequality follows from $D_j \ll b^2$ for sufficiently large x using $p \ge 2$.

Fix a live exponent $e \in I_{D_j}$, write $P = p^e$, and let $\tau = \mathfrak{p}_e(\sigma)$. The elementary factorial contribution less the cost for that coupling is

$$\begin{aligned} &\frac{a_\sigma + D_{j,\sigma}}{P} - \mathfrak{C}_{p,e}(\sigma, \tau) \\ &= \frac{a_\sigma}{P} - \min(a_\sigma^{ev}/P, c_{j,\tau} + a_\tau^{ev}) \\ &\qquad + \left(\frac{D_{j,\sigma}}{P} - D_{j,\tau}\right)_+ - \frac{c_{j,\sigma}}{P}. \end{aligned}$$

The ordinary-box term is at least $k(w_\sigma/P - w_\tau)_+ - x^{60}$, by Lemma 9.18. The last term costs at most u. For the distinguished height-one term, if σ is not a top then the contribution and requested gain are zero. If both σ, τ are nonexempt tops, the loss relative to $(a/P - w_\tau)_+$ is at most a/P, the collision summand. If σ is a nonexempt top but τ is not a nonexempt top, then $a_\tau^{ev} = 0$ and $c_{j,\tau} \le w_\tau + x^{60}$, giving at least $(a/P - w_\tau)_+ - x^{60}$. If σ is an exempt top, the minimum is zero and a/P exceeds the requested positive part by $\min(a/P, w_\tau)$; at $e = 1$ this is at least a/x_1 since $p < x_1$ and $w_\tau \ge a/x_1$. These four cases give the pointwise desired bound with at most $3u$ of error per place and exponent, plus the exempt contribution for $e = 1$.

For live e above the truncation cutoff there is no coupling cost; the factorial contribution then bounds both requested positive parts. The extra contribution at $e = 1$ does occur within the truncation since $D_j \ge a > x_1$. In summing the factorial contribution over all live exponents use Lemma 9.1. The $e = 0$ term allowed at $p = d$ occurs only in the coupling cost, and costs at most $4u$ per place. Indeed $p = d$ is fixed and $c_{j,\sigma} \le u$ from $w_\sigma \le a/x$ and $m/x \to 0$. The place $[-1]\sigma$ cannot have the same class of tops as σ, nor can both be thick. Hence the first minimum is bounded by u, and the last by $ka/x_1 + x^{60} \le u$, since at least one term corresponds to the small weight. Summing, the comparison error is at most $4fu(\lfloor \log_p D_j \rfloor + 1)$, and the factorial error at most $fu(\lfloor \log_p D_j \rfloor + 1)$. Multiply the comparison by $\omega_p \le 1$ and use the last scale inequality above. The sum of the desired positive parts with factor ω_p is exactly g_p/m by (9.2). □

The exact uniform source declaration is `Data.le_jetExponent`, in `New/Columns/JetExponentBound.lean`. Formula (9.20) is the numerical interface to the chain. When $ev = \varnothing$ there are no charges, but long-row collisions can be expensive. The Hecke repair uses exempt sets which remove all such collisions; the extra positive term in (9.20) compensates part of the charge incurred in doing so.

9.10 From theta products to the standard integral kernel data

A bound for arbitrary formal products becomes a bound for the automorphic theta kernel only after identifying the normalized finite integral structures. We give this interface explicitly.

For a standard space V let vac_V be the standard finite vacuum. Fix a nonzero factor s_{vac} chosen for rank one and define the scaled vacuum of rank h to be $\mathrm{vac}_V^{\circ} = s_{\mathrm{vac}}^{h}\,\mathrm{vac}_V$. Its point evaluations are the products over the h coordinates of the rank-one evaluations. The integral module $\mathrm{Std}_\delta(V,\iota)$ is the span over $\iota^{-1}(\mathfrak{o})$ of $\varrho(\alpha x)\,\mathrm{vac}_V^{\circ}$ with δx integral at all finite places above p (x is a finite adelic vector). This definition is `Parameters.standardData` in `New/Curve/IntegralVacuum.lean`; it depends on $\delta \in \mathcal{O}_K$ and on the p-adic embedding.

Proposition 9.20 (Formal-source input: integral vacuum and conjugates)**.** *The scalar s_{vac} is nonzero, and for all p and all p-adic embeddings every finite point evaluation of vac_V° is integral. For $\sigma \in \mathrm{Aut}(\mathbb{C}/K)$ its conjugates are bounded by $\exp(CN_1^2 h)$ for a constant C depending only on the fixed K, N_0. Moreover there is a finite vector generating the finite part by Heisenberg translates and with theta functional equal to 1; after applying σ, its jets at scale* per *coincide, up to permutation of the places, with jets of translates of another such finite vector at a scale of squared absolute value $1/\Omega$.*

The precise uniform generator statement is `Parameters.exists_isJetConjugate_generator` in `New/Curve/ConjugateGenerator.lean`; integrality and the conjugate bound for the scaled vacuum are `norm_intVacuum_le_one` and `norm_intVacuum_le` in `New/Curve/IntegralVacuum.lean`. These inputs use the normalized CM reciprocity law. For higher rank the vacuum value bounds follow multiplicatively from rank one, as indicated by `intVacuum_finVec_eq_prod`.

The integral module is stable under the Heisenberg translations by the indicated lattice because the central factors in their multiplication law are roots of unity and hence units at p. It is also stable under the finite Weil action of an element integral above p. This follows from the finite vacuum-translation sum for the Weil action on a vacuum: it is a sum of translates by integral vectors with unit multipliers, and conjugation by the integral element preserves the translation lattice. This bounded-action assertion is used in its exact form `Parameters.omega_mem_standardData` and `finHeisOp_mem_standardData`, in `New/Curve/ConjugatesOfJets.lean`. We assume the finite Weil model's vacuum sum identity in this application.

Lemma 9.21 (Transfer of a jet bound)**.** *Let $\mathcal{L}$ be a complex linear functional on the Taylor series space (in particular any specified Schur jet entry), $c \in \mathbb{C}$ a scale multiplier, and Q a nonnegative real bound. Suppose the norm of $\iota(\mathcal{L}(F))$ is at most Q for every F whose image under ι is a frame theta product with the torsion condition $n\delta y_i \in \Lambda$ and with scalar norm $|\iota(cs_1)|_p$. Then for every $f \in \mathrm{Std}_\delta(V,\iota)$,*

$$|\iota(\mathcal{L}(\mathcal{T}_{c\ \mathrm{per}} f))|_p \le Q.$$

Proof. Take a frame with $M = p(dN_1)^6 + 1$, of length prime to p, using Proposition 9.10. Repeat its vectors so that all multiplicities are one. Use Proposition 9.11 with $\delta = M$ and the corresponding repeated frame of length M^2, choosing the monomial sections equal to the normalized theta section. Then $v = \Omega_0$ and $D_\delta = M^2$ satisfy the required area identity. The resulting finite vector F_0 has theta value one, since every normalized theta section has constant value one. In the vacuum expansion take ℓ to be the theta functional and solve for the vacuum. After scaling by s_{vac}^{h}, the coefficients are integral: they are the normalized vacuum evaluations times the natural multiplier divided by M^{4fh}, a p-adic unit. The phase factors from translating these expansions have absolute value one.

Consequently a lattice translate of the normalized vacuum is an integral linear combination of translates of F_0. Approximate each finite translate by a rational one, using density of rational points of the finite additive adeles and local constancy of the Heisenberg orbit on the finite model. The rational approximation can be made so that δ times the translation

vector is integral above p. Denominators at other primes are cleared by a rational integer prime to p. Applying the frame product identity gives the class of theta products in the hypothesis: the scalar factor is the product of cs_1 and factors whose norm is one from M; the torsion coordinates have denominators dividing $n\delta$ with $p \nmid n$. This exact rational approximation and denominator assertion is the one used in `Parameters.norm_jetSeries_le` in `New/Curve/ConjugatesOfJets.lean`. Linearity of $\mathcal{L}$ and the non-Archimedean inequality then bound every generator of the integral module, hence its entire span. □

This argument uses the stated finite-vector and period inputs rather than replacing the actual kernel data by a single formal theta product. For $\delta = 1$ at a nonsplit prime, Proposition 9.14, Theorem 9.17, and Lemma 9.21 bound the Schur jets on the entire module whenever the scalar multiplier has integral norm. For general $\delta \neq 0$, Lemma 9.4 applied using Proposition 9.2 gives the elementary integrality at scale $\overline{\delta}d$ per; this is precisely `Parameters.norm_jet_jetSeries_le_one` in `New/Curve/ConjugatesOfJets.lean`.

Finally, a finite theta kernel acts not just on one fixed Weyl coefficient vector but on vectors at all adelic arguments. The Weyl gauge at g of a vector ξ is the maximum absolute value of its coordinates after reading $g^{-1}\xi$ through the specified p-adic embeddings (where the action contains the determinant character as well as the Schur representation). This is `weylGauge` of `Lib/GLn/SchurModule/WeightRepBasic.lean`. An integral change of adelic frames above p has integral matrices in both directions under this representation, so the gauge is unchanged. For the standard data f the kernel is equivariant for these integral changes.

Proposition 9.22 (Formal-source input: gauge bound for a translated theta kernel)**.** *If the pair of adelic changes of the matrix variable factors into integral changes and a diagonal element with parameters $u_\sigma(a)$, then a uniform bound on every jet coefficient after weighting by*

$$\prod_{\sigma}\left(\prod_{a}|u_\sigma(a)|_p\right)^{e_\sigma}\prod_{\text{boxes }(\sigma,h,a)\text{ of }i'}|u_\sigma(a)|_p$$

is a bound on the factor by which the theta kernel raises the Weyl gauge. Here e_σ is the determinant exponent of the output weight. In particular the kernel with standard datum and scale c per, with no diagonal change, raises gauges by at most

$$\left(\frac{|\iota(c)|_p}{|\iota(\overline{\delta}d)|_p}\right)^{D},$$

where D is the number of boxes of the output tableau.

The exact interface is `Parameters.mem_raises_of_jet_bound` and `mem_raises_of_standardData` in `New/Curve/ConjugatesOfJets.lean`. Its finite-place factorization and equivariance are used as inputs. The matrix estimate itself is just the ultrametric operator bound, and the weights come from evaluating the diagonal matrix in the Schur representation and its determinant character. The final displayed bound follows by homogeneity and the crude integrality at scale $\overline{\delta}d$ per proved above. More importantly, at a non-split prime away from the conductor and level, the sharper Schur bound replaces that crude estimate by (9.18), with any extra weights charged box by box. This completes the local differential calculation needed for the Hecke and seesaw argument.

Chapter 10

The definite chain, collisions and Hecke repair

The purely algebraic local Hecke estimates, including the Frobenius root and coefficient arguments, are developed in Chapter 11. The actual rational orbit and finite integral support calculations for the geometric chain are developed further in Chapter 12. Reading them separately is useful: the former concern a finite coefficient lattice, whereas the seesaw identifies the lift with a construction on an actual adelic lattice.

We continue the construction of the preceding section. Throughout this section the prime d of K is fixed. Write $g = g_0 g_1$ for the product of the two level primes, and let $\pi_g \in \mathcal{O}_K$ have norm g. The letter π without a subscript denotes the element of norm N used to define the extra column level. The spaces W and C are, respectively, the good part of the row space and the entire initial column space, of dimensions n and c. We use $n_j = k + j(t+1)$ for the ranks of the definite Hermitian spaces X_j, $0 \le j \le m$; thus n_j is not a dimension of a space of automorphic forms. Set $w(\tau) = w_{\mathrm{nom}}(\tau)$. The local ask A_p, the prime sum A, the gain g_p, and the direction cost $\mathcal{U}$ are as in the preceding section; in particular

$$g_p = m \left(\sum_{\tau \in \mathcal{T}} A_p(a, \tau) + k \sum_{\tau \in \Phi} A_p(w(\tau), \tau) \right). \tag{10.1}$$

Here we prove the quantitative implications between the jet theorem and Theorem A. We also isolate the remaining nontrivial automorphic and local representation-theoretic assertions on which the argument depends. They are identified explicitly as formal-source inputs when we do not give an independent proof of their full strength.

10.1 Coefficient lattices, class spaces and index

It is important to formulate the lattice estimates without assuming that the coefficient field of the automorphic forms is a fixed number field. For a rational prime p, fix a field isomorphism $\iota : \mathbb{C} \simeq \overline{\mathbb{Q}}_p$ as in the coefficient construction, and transport the valuation ring and absolute value to $\mathbb{C}$. We denote this ring by O_ι and normalize $|p|_p = p^{-1}$. All assertions below are quantified over ι. For a Hermitian space X and adelic point g', let $\|v\|_{\iota,g'}$ denote the Weyl gauge of its coefficient vector in the integral structure at g'. Thus, if ρ_κ is the algebraic coefficient representation, the gauge is obtained by applying the inverse of the adelic coefficient matrix to v and then taking the coefficient supremum norm transported by ι. It is non-Archimedean; multiplication by a scalar z multiplies the gauge by $|\iota(z)|_p$. At the ordinary

complex absolute value the Weyl norm is instead the chosen Hermitian norm; the two kinds of estimate should not be confused.

For level K_j and weight κ of the given definite space X_j, define $M_j(\iota)$ to be the O_ι-module of forms of this type, every value of which has gauge at most 1. Write $M = M_0(\iota)$, a submodule of C. All finite-place estimates in the proof are made on this adelic integral module, not just on abstract polynomials in a single weight representation. The exact definitions are `Data.integralColumns` and `integralFormsAt` in `New/Exchange/CrudeModularData.lean` and the background files respectively. The notation M in this section always depends on the chosen embedding.

We recall the index convention used in the proof.

Definition 10.1. Let P, Q be submodules over a valuation domain O, and let $z \in O$. Write $I(P, Q; z)$ if for every finite family $(v_i)_{i=1}^h$ of elements of P there exists a square h by h matrix D over O such that $\det D \mid z$ in O and

$$\sum_j D_{ij} v_j \in Q \quad (1 \le i \le h).$$

For an O-linear endomorphism T preserving P, write $T \in \mathrm{Type}(p, P, s; e)$ if there is an integer $l \ge 0$ with $l \le e$ and

$$I(P, TP + p^s P; p^l).$$

The parameter s is a nonnegative integer and e is allowed to be a nonnegative real number.

These are `IndexDvd` and `HasType` in `New/Core/ModuleIndex.lean`. The divisibility of the determinant is in the stated direction: $\det D \mid p^l$ implies $|\det D|_p \ge p^{-l}$. It does not mean that $\det D$ is a multiple of p^l. Nor is it necessary to assign a finite cardinality to a lattice quotient. This definition is appropriate even when O is not discretely valued, and is sufficient for the later height argument.

Lemma 10.2 (Elementary index calculus)**.** *Let $I(P, Q; z)$ and $I(Q, R; a)$ hold. Then $I(P, R; za)$ holds. If u is an O-linear map with $u(P') \subset P$ and $u^{-1}(Q) \cap P' \subset Q'$, then $I(P', Q'; z)$ holds. If $u(Q) \subset Q'$ and $P' \subset u(P) + Q'$, then $I(P', Q'; z)$ also holds. Consequently for $T_1, \dots, T_h$ preserving P,*

$$T_i \in \mathrm{Type}(p, P, s; e_i)\ (1 \le i \le h) \Longrightarrow T_1 \cdots T_h \in \mathrm{Type}(p, P, s; \textstyle\sum_i e_i).$$

Finally, if $u_i : P \to V'$ are linear and $I(P, u_i^{-1}(Q'); z)$ holds for each i, then

$$I\left(Q' + \sum_{i=1}^{h} u_i(P), Q'; z^h\right).$$

Proof. The first assertion is obtained by applying the index witness matrix for $I(Q, R; a)$ to the family already obtained using the matrix for $I(P, Q; z)$, and taking their product. The determinant is multiplicative. For pullback, apply the index condition to $(u(v_i))_i$ and use linearity. For pushforward, write $v_i' = u(v_i) + q_i'$ with $q_i' \in Q'$ and use the same linear combinations of these expressions.

For the product assertion it suffices to consider two factors. Suppose we have index witnesses for $T_i P + p^s P$, with determinant bounds p^{l_i}. Apply the second operator to the index condition for the first, if necessary reversing the order in this statement. Since $T_i(p^s P) \subset p^s P$, the pushforward rule gives the required condition from $T_1 P + p^s P$ to $T_1 T_2 P + p^s P$, and the transitivity rule gives the sum $l_1 + l_2$. Now use induction. For the final assertion enlarge Q' successively by the images. In each step, push the index condition along u_i and use transitivity. □

For comparison, if $P = O^D$ is free over a discretely valued ring and T has elementary divisors $p^{\nu_1}, \dots, p^{\nu_D}$ (with $\nu_i = +\infty$ allowed), the determinant valuation of the canonical lattice $TP + p^s P$ is

$$\sum_i \min(\nu_i, s).$$

Thus a bound on the type measures a truncated sum of elementary divisors. The Hermite and Smith factorization arguments in `ModuleIndex.lean`, in particular `IndexDvd.of_det_dvd`, implement this relation for the valuation domains of the proof. We will use type bounds, not an assertion about each individual eigenvalue of a Hecke operator.

10.2 The groups and the chain of theta kernels

At stage j the source is the standard definite space X_j of rank n_j. The column level K_j is the intersection of the chosen principal level at the places dividing $\overline{\pi_g}$ with the subgroup whose last row is congruent to $(0, \dots, 0, 1)$ modulo π. The element π is chosen so that $\pi\bar{\pi} = N$. There are f places dividing π, one member of each conjugate pair over the completely split prime N. In particular K_j contains the principal level of $\overline{\pi_g}\pi$ and is compact open. The spaces of forms are finite sums of coefficient spaces indexed by rational adelic classes. The chosen principal level makes the automorphism groups of the classes trivial. These claims are `Data.columnLevel`, `Data.principalLevel_le_columnLevel`, `Data.card_piPlaces`, and `Data.classAut_columnLevel_eq_bot` in `New/Columns/ColumnDefs.lean`. The analogous triviality at the intermediate stages follows from `Data.classAut_eq_bot_of_le_chosenLevel`. We use this specific property as a formal-source input for the adelic lattice model; no estimate for a weighted stabilizer sum can replace it in the following argument.

Here is how the coefficient weight is structured. With $h = 2h_2$ as in the admissible data, the number of columns of height j' in the Young diagram at τ is

$$\kappa_\tau(j') = \begin{cases} a, & j' = 1,\ \tau \in \mathcal{T}, \\ 0, & j' = 1,\ \tau \notin \mathcal{T}, \\ w_1(\tau) - \lfloor (k+t+1)/2 \rfloor, & j' = k, \\ h - (t+1), & j' > k, \quad t+1 \mid j'-k, \\ 0, & \text{otherwise.} \end{cases}$$

Only columns of height below n_j contribute to the polynomial part at stage j. The determinant exponent of the representation is 0 at stage zero, and $e_j(\tau) = \lfloor n_{j-1}/2 \rfloor - w_j(\tau)$ for $j \geq 1$. The coefficient representation on X_0 is in particular

$$\bigotimes_{\tau \in \mathcal{T}} \mathrm{Sym}^a \mathbb{C}^k.$$

All factors at the other places have degree zero. The restriction on the places of height-one columns is central in the collision argument: once a p-adic Frobenius orbit meets both top places, the two factors of degree a can interact at precisely the prime-power levels that contribute to the required gain.

Let θ be an admissible twist of finite conductor of rational norm q_θ, and let μ_{j+1} be the prescribed splitting character of the $j+1$st lift, with twist θ^{j+1}. The theta kernel K_j is the one for $\mathrm{Hom}(X_j, X_{j+1})$ with weight κ, character determined by $\mu_{j+1}\theta^{j+1}$, and complex Taylor scale $\overline{\pi_g}$per. The identities between weights that make the lift available use the fact that there is no nonzero column height strictly between n_j and n_{j+1}. The declarations are `Data.i`

`sLiftWeight_kappa` and `Data.liftKernel` in `New/Exchange/CrudeModularData.lean`; the formulas for κ, e_j are in `New/Curve/LiveSum.lean`.

For a tuple of finite adelic matrices $\xi = (\xi_j)_{0\le j<m}$ in the corresponding Hom spaces, let $\phi_{\xi,j}$ be the finite Heisenberg translate by ξ_j of the normalized integral vacuum of the preceding section. Let

$$\mathsf{C}^{j}_{\theta,\xi} = \mathsf{K}_{j-1}(\phi_{\xi,j-1}) \circ \cdots \circ \mathsf{K}_0(\phi_{\xi,0}), \qquad \mathsf{C}^{0}_{\theta,\xi} = 1$$

be the composite definite lift to stage j. Individual terms need not be invariant under the intermediate level. A useful datum is instead a finite linear combination of tuples $\Psi = \sum_\xi c_\xi[\xi]$, which is required to be jointly invariant under the levels in the sense that the corresponding tensor of kernel data is fixed by their simultaneous actions.

At the final stage X_m has rank one greater than Y_m. The map to the last row space is not simply evaluation on $U(Y_m)$: it uses branching of the algebraic coefficient representation. Let D_{br} be the integral branching matrix and put

$$e_{\mathrm{br}}(\tau) = \sum_{j'<n_m} (\kappa_\tau(j') - \kappa'_{\mathrm{kept}}(j')) + e_m(\tau).$$

Here the summands count the columns at that height whose height is reduced in branching; the formal definition of κ'_{kept} is $h_2 - \lfloor (t+1)/2 \rfloor$ if $j' \ge k$ and $t+1 \mid j'-k$, and zero otherwise. Thus the linear restriction is

$$(\mathsf{R}v)(u) = \left(\prod_{\tau\in\Phi} \tau(\bar\pi)^{-e_{\mathrm{br}}(\tau)} \right) D_{\mathrm{br}} v(\mathrm{diag}(u,1)). \tag{10.2}$$

The relevant branching identity is `Data.restrictToY`, and the estimate `Data.le_branchExp` is

$$e_{\mathrm{br}}(\tau) \ge \kappa_\tau(1) - 2b.$$

For a prime $p \ne N$, the scalar in (10.2) is a p-adic unit; using the lattice-compatible branching identity the restriction cannot increase the gauge. At a complex conjugate σ the estimate is stronger:

$$\|\sigma\mathsf{R}v(u)\|^2 \le \exp\left(x^{70} - \textstyle\sum_\tau e_{\mathrm{br}}(\tau)\log N\right) \|\sigma v(\mathrm{diag}(u,1))\|^2. \tag{10.3}$$

Indeed under each conjugate $|\tau(\bar\pi)|^2 = N$, so the scalar gives the displayed negative exponent; the branching matrix contributes at most $e^{x^{70}}$ by the coefficient estimate. The precise lattice and norm compatibility statements `Data.GaugeLe.restrictToY` and `Data.weylNormSq_restrictToY_le`, in `New/Seesaw/LiftCharacters.lean`, are formal-source inputs. Their significance is that the archimedean gain of level N is large, while the prime N itself will be treated separately.

Write

$$\mathsf{B}^{\Psi}_{\theta} = \mathsf{R} \sum_{\xi} c_\xi \mathsf{C}^{m}_{\theta,\xi}.$$

This is `Data.Bμ` in the source. For the local calculations the datum must have the following exact properties (the record `Data.IsGoodDatum`, in `New/Columns/ChainValuation.lean`):

1. If ξ lies in the support, its components multiplied by π_g are integral away from the places of $q_\theta N$.

2. The coefficients c_ξ are integral at every ι with $p \nmid q_\theta$. For $q_\theta = 1$ the sum of the absolute values of their conjugates is at most $e^{x^{60}}$.

3. The total tensor is jointly invariant under the intermediate column levels and the final action specified in the source.

4. At N the chain has an additional paired description. Project each kernel datum to the places dividing π; sum the pairwise translations of adjacent data by last-column $-y/N$ and last-row $-\bar{y}/N$ for y modulo N. The resulting combination of paired lifts, with the same coefficients c_ξ, equals the original chain of m lifts. In the source $m = 2(\varrho + 1)$.

The fourth property, `chain_at_N`, will make it possible to move large apparent denominators at N onto a single column operator, instead of losing them at every Taylor coefficient.

Lemma 10.3 (Gauge bounds for a jointly invariant chain)**.** *Suppose $p \nmid q_\theta$, Ψ is good and each translated kernel after a chosen modification obeys*

$$\|\mathsf{K}_j(\phi_{\xi,j})v\|_{\iota,u} \le p^{r_j}\|v\|_{\iota,g'}$$

at the prescribed coefficient frames, uniformly for each pair of points and every ξ in the support. Assume also that the modifying operators commute with the required source and target level actions on the total invariant tensor. Starting with an integral column, the total modified chain has gauge at most $p^{\sum_{j<m} r_j}$ in the final frame.

Proof. For jointly invariant tensors all the integrals over definite groups reduce to finite class sums; an individual term is a composition, indexed by a tuple of representatives of the successive class spaces. The triviality of the stabilizers means that such a sum is a plain sum of the corresponding values, not a sum with rational denominators from the stabilizer orders. Each composite has gauge at most the product of the individual bounds. The coefficients c_ξ are integral since p is prime to q_θ. Thus the non-Archimedean triangle inequality leaves the largest gauge bound unchanged under summation. This proves the result. The source formalizes the reduction to the class sum and this argument as `Data.IsGoodDatum.gaugeLe_chain` via `Data.exists_chain_eq_sum` in `New/Columns/ChainValuation.lean`. □

The same argument at an archimedean place bounds the square root of the squared norm by the product of individual operator norm estimates, the sum of absolute values of the coefficients, and the product of the numbers of integral classes at each stage. This explains why the class-counting estimate and coefficient-size estimate must be checked even when no such factor is present in the non-Archimedean estimate.

10.3 The geometric seesaw and its denominator control

Let P be the (unprojected) Poincaré construction from columns to initial row forms. It comes from a Bessel period for the unitary space Y_0; the successive indefinite theta lifts run through the spaces $Y_0, \dots, Y_m$ of ranks $r + jt$. Our notation Θ_θ is reserved for the composite on W (or its extension to the unprojected principal-level row space before selecting W). Initially we do not identify it with any definite chain.

Theorem 10.4 (Seesaw identity, formal-source input)**.** *There is a nonzero $z_0 \in \mathbb{C}$, independent of the twist, and for every admissible twist θ a good datum Ψ_θ such that*

$$\Theta_\theta P(z_0 u) = \mathsf{B}_\theta^{\Psi_\theta} u \qquad (u \in C^{\mathrm{ram}}).$$

The declaration is `Data.seesaw` in `New/Seesaw/SeesawTheorem.lean`. The exact content includes the four properties of good data above, not just an equality of two complex-valued functions. In the following we explain the denominator mechanism and give the finite cut argument in a form that is independent of the deeper theta calculations.

In the adelic construction a stage j has a rational Bessel subgroup, with finite-adelic component denoted here by S_j, and a compact open subgroup $J_{j,\theta}$ of the row group. The function arising after the first j unfoldings is P_j, with P_0 the generic function from which the original Poincaré series is formed. The rule $P_{j+1} = \mathsf{U}_j(P_j)$ is a Bessel unfolding by a single step of theta lifting, at parameters (ξ_j, y_j) of the archimedean and finite row spaces.

Three deeper assertions of the source are needed in this construction; we do not reprove the full geometric or local theta identities.

1. The local constant-term identity preserves the generic character, and consecutive pairs of unfolding steps give matrix coefficients of a definite paired chain. At the even step $2i + 2$, the equality is

$$P_{2i+2}(\iota_X v \cdot (x, y_{2i+2})) = \mathrm{algCoeff}_{t_i}\,(\mathsf{C}^{\mathrm{pair}}_{i+1}(u))(\mathrm{emb}_X v)$$

   for all ramified columns u and all configurations y with initial point y_0 in the chosen level. The linear functional t_i depends on the archimedean parameters but not on the twist or on u. These are `Data.isGeneric_unfoldStage` and `Data.unfoldStage_even` in `New/Seesaw/UnfoldingStages.lean`; the local arrow-pair tensor and integral identity used in their proof are in `New/Seesaw/EdgePair.lean` and `New/Seesaw/DoubleLiftValue.lean`. In the formal proof one unfolds the two consecutive theta kernels, takes their constant terms with respect to the unipotent subgroup, and applies the identity for integrating a product of kernels. This gives $i + 1$ paired lifts from the assertion for i.

2. The support of each stage function is controlled by the lattice lines of its Hermitian space. If P_j is nonzero, the lines at places of the twist move at most j steps, and the integral lines outside those places do not move (`Data.hasNearLines_of_unfoldStage_ne_zero`, same file). In the induction, even steps are covered by the line estimate for the theta kernel. At odd steps the possible exceptional local lines at the level prime are eliminated by writing a ramified column as $u' - z_w u'$ for a central generator z_w at a place over N. The two terms of the matrix coefficient are then equal there, so cancel. This use of ramification is why the seesaw statement is on C^{ram}.

3. For each stage one can choose a compact open A_j independent of the twist and a finite cut $T_{j,\theta} \subset G_{j,f}/J_{j,\theta}$ stable under $S_j \cap A_j$ and contained in the level outside the twist conductor. Here $G_{j,f} = U(Y_j)(\mathbb{A}_f)$. The cut covers the support after translation by S_j and has cover multiplicities which divide powers of q_θ. At $j = 0$ it is contained in the chosen level. This assertion is `Data.exists_stage_cut` in `New/Seesaw/SeesawTheorem.lean`.

Lemma 10.5 (Finite-cut multiplicities)**.** *Let G be a locally compact totally disconnected group, S a closed subgroup, and $J \subset A$ compact open subgroups of G. Let $D \subset G$ be compact, and suppose that the set to be covered is in SD. Suppose for each $x \in D$ there is an e such that*

$$(xux^{-1})^{q^e} \in A \qquad (u \in J).$$

There is a finite $(S \cap A)$-invariant subset $T \subset G/J$ whose S-saturation covers the image of SD. For v in this saturation, write

$$m(v) = \#\{s(S \cap A) : s \in S,\ s^{-1}v \in T\}.$$

Then $m(v)$ divides a power of q.

Proof. The image of D in the discrete space G/J is finite. Choose one element $[x]$ in that image from each S-orbit, and take T to be the union of their $(S \cap A)$-orbits. These are finite

since $S \cap A$ is compact. If H_x is the stabilizer of $[x]$ in S, then the s with $s^{-1}[x] \in T$ constitute $H_x(S \cap A)$. Thus

$$m([x]) = [H_x : H_x \cap A].$$

This index is finite, and is invariant under translating $[x]$ by S. Pass from H_x to its quotient by the normal core of $H_x \cap A$; this is a finite group. As $H_x \subset xJx^{-1}$, every element in the finite quotient has order dividing a power of q. Indeed for each $h \in H_x$ and each conjugate by an element of H_x, that conjugate's q^e-power belongs to A, hence h^{q^e} belongs to the normal core. By Cauchy's theorem every prime divisor of the order of this finite group divides q. The index divides its order, proving the assertion. This is the finite-index argument in `exists_cut_coverMultiplicity_dvd`, `New/Seesaw/IntegralAsCosetSum.lean`. □

The cut hypothesis is justified in the source by the support of the unfolded stages. Outside the twist the local groups are integral. At the finitely many places of the twist, taking a sufficiently large power of the residue characteristic puts an element of the conjugated compact open level back into the fixed level. The admissible translations therefore lie in finitely many local compact sets modulo S_j; the stage-cut theorem applies the lemma to those sets. Thus the importance of the local geometric support bound is not just convergence: it is the prescribed prime support of denominators.

At the analytic level the unfolding reduces an integral over the rational Bessel quotient to a sum over a finite cut. In the notation of the lemma the formula has the form

$$\int_{B(F)\backslash G(\mathbb{A}_F)} P_j(g)z(g)\,dg = \lambda_j \sum_{v \in T_{j,\theta}} m(v)^{-1} b_j(v).$$

The scalar λ_j is determined by the fixed archimedean integral and Haar normalizations; it is independent of the twist. To see the finite sum mechanism, the weight $\mathbf{1}_T(v)/m(v)$ is a partition of unity after summing over $S_j/(S_j \cap A_j)$ on the support. Indeed exactly $m(v)$ terms contribute, with the same orbit multiplicity. It may therefore be inserted in the integral without changing it. Integration over a coset of the compact subgroup $J_{j,\theta}$ contributes its Haar mass; the standard level-index normalizations in the period over B cancel this factor. The remaining archimedean integral is λ_j. The analytic hypotheses are precisely the generic character identity, automorphy of z, integrability of $P_j z$ and the product form of its Bessel-group integral. The formal-source input for this quotient-integral statement, including its measure normalizations, is `Grading.exists_integral_eq_sum_over_cut` in `New/Seesaw/IntegralAsCosetSum.lean`.

The product of the resulting finite combinations along the chain is again a finite combination: its coefficients are sums of products of $1/m(v)$. For $p \nmid q_\theta$ these are integral. If the twist is trivial, $q_\theta = 1$, the multiplicities are all 1 wherever the cut is nonempty. For quantitative complex sizes one counts the cut at each stage by the index of the chosen principal level; this gives a power of `Data.cutBound` independent of the actual columns. The formal assertion `Data.exists_isAdmissibleComb_unfolding` in `New/Seesaw/SeesawTheorem.lean` carries out this induction, testing against every suitable form z and linear functional on the tensor space. The terminal stage identifies the coefficient with paired lifts, using the preceding even-step identity.

To turn the resulting tensors into the vacua of the definite chain one needs a further finite Weil representation calculation. For the evaluated kernel tensor at each configuration y which is integral outside q_θ, the source proves that it is an admissible linear combination of reduced vacua, up to a single constant C and with denominator $\mathrm{cutBound}^{2m}$, bounded size when the twist is trivial. The reduced vacua are the paired combinations of Heisenberg translates described above. The exact declarations are `Data.exists_isAdmissibleComb_reducedVacu`

`um`, `Data.isAdmissibleComb_reducedVacuum`, and `Data.exists_isGoodDatum_of_invariant` in `New/Seesaw/GoodDatumExists.lean`. We take the three assertions as formal-source inputs; they contain the explicit evaluation of finite Weil data and the local fiber geometry, rather than just the group-theoretic finite-cut calculation. The scalar C and the archimedean normalization from the integrals are absorbed into the one scalar z_0. The theorem then follows by pairing against the test vectors at stage m; equality of all those matrix coefficients gives equality of the lifts. These are the mechanisms underlying Theorem 10.4.

10.4 Ramification and short generators on the definite group

The set of ramified columns has an elementary description which is compatible with its adelic lattice. For each $v \mid \pi$ choose a central element of the local norm-one torus, integral everywhere, supported on $v, \bar{v}$, whose component at v reduces to a generator of $\mathbb{F}_N^\times$. Let z_v be its translation operator on columns. Define

$$R_N = \prod_{v|\pi}(1 - z_v), \qquad C^{\mathrm{ram}} = R_N C.$$

These operators commute. In particular every element of the image can be written $(1 - z_v)u'$ at each $v \mid \pi$. The definition and the lift of the generators are in `Data.ramifyingOp`, `Data.ramifiedColumns` and `Data.exists_centralGen` in `New/Columns/ChainValuation.lean`.

The following dimension and norm observations will be useful both here and for the spanning theorem.

Lemma 10.6. *For each $v \mid \pi$,*

$$\dim \ker(1 - z_v) \le \frac{c}{N-1}.$$

Consequently for any subspace $U \subset C$ and any subset T of the places dividing π, with $R_T = \prod_{v\in T}(1 - z_v)$,

$$\dim R_T^{-1}(U) \le \dim U + |T|\frac{c}{N-1}.$$

If u is integral at every embedding and every value of its conjugates has squared norm at most C_0, $R_T u$ is likewise integral, with squared norm at most $4^{|T|}C_0$.

Proof. The forms fixed by z_v have level enlarged by the cyclic group generated by the central translation. The index of the old level in this enlargement is at least $N - 1$: a positive power lying in the old level reduces to the identity at v. Since both levels are contained in the chosen level and their class automorphism groups are trivial, their spaces have dimension proportional to the numbers of adelic classes. The ratio of these class counts is the level index; thus the dimension of the fixed space is at most $c/(N - 1)$. This is the class-count formula specialized in `Data.finrank_ker_one_sub_centralGen_le`. For a linear operator L, $\dim L^{-1}(U) = \dim \ker L + \dim(U \cap \operatorname{im} L)$. Applying this successively proves the second assertion. The central element preserves the integral gauge and permutes the arguments of an algebraic form. Therefore $u - z_v u$ is still integral by the ultrametric inequality, and its squared complex norm is at most $4C_0$ by the ordinary triangle inequality. Induction gives the bound. This is `Data.ramifyingOp_finrank_comap_and_short`. □

To produce the initial short vectors the source uses a geometry-of-numbers estimate for definite Hermitian lattices. We record its role explicitly. Let I be a level whose class stabilizers are trivial, let the coefficient representation of rank u have a Young diagram with D_κ boxes

(counted over the infinite places), let $Q \geq 1$ be a rational integer in the different ideal of E, and suppose $M_0 \geq uf + D_\kappa$. Every class admits an algebraic basis of its associated lattice with the uniform bounds summarized in `exists_short_basis_stdSpace`. Put $R_0 = 6Q(fu)^2$. Then the coefficient forms which are everywhere integral with squared norms of all conjugates bounded by $R_0^{6M_0}$ span the whole class space (`span_isShort_eq_top`, in `New/Columns/ChainValuation.lean`). This is a formal-source input for the quantitative short basis bound; the coefficient argument is as follows. The class space is freely the direct sum of weight modules at the class representatives, so it suffices to span one such weight module. Apply the Schur representation to the bounded lattice basis. In the standard coefficient model the Weyl vectors obtained by evaluating integral flag polynomials at roots of unity span the weight module: if a linear functional annihilates all such vectors, successive application of the root bound for a one-variable polynomial shows that the corresponding polynomial vanishes identically. Multiplication by the determinant factor for the determinant weight also preserves integrality in the chosen lattice basis. Every conjugate of a root of unity has absolute value 1, so the norm bound follows from the power of R_0 for the bounded basis and the degree of the coefficient representation. This is why one obtains a spanning set rather than simply a single short vector. Applying this construction at level K_0 with the admissible scale estimates, and then applying R_N, gives the needed conclusion:

$$\operatorname{span}_{\mathbb{C}}\{u \in C^{\mathrm{ram}} : u \text{ is everywhere integral, } \|\sigma u(g')\|^2 \leq e^B \ \forall \sigma, g'\} = C^{\mathrm{ram}}. \tag{10.4}$$

Specifically $|\{v \mid \pi\}| = f$, and the bound before application of R_N is chosen at most $e^B/4^f$. The exact scale inequality and conclusion are `Data.span_short_ramifiedColumns`. We will use (10.4) with the ordinary complex norm squared, not its logarithm, bounded by e^B.

10.5 The collision loss and the marked set

For a rational prime p, put

$$\omega_p = \frac{1 - \chi_K(p)}{2}, \qquad s_p = \lceil \log_p m \rceil, \qquad \delta_p = s_p f(n_m + s_p).$$

A prime is called short if $p \equiv 1$ or $-1 \pmod m$. In this terminology shortness refers to a congruence and not to the size of p. Define

$$\ell_p = \frac{fa(\log b)^2}{p} \begin{cases} m, & p \text{ short}, \\ 1, & \text{otherwise}. \end{cases}$$

We will call p small nonsplit if $p < x_1$ and $\chi_K(p) \neq 1$. At such a prime, g_p in (10.1) is the desired total valuation. In the formulas below, all relevant place actions are the actions of the CM cyclotomic type, not of the single quadratic extension $K/\mathbb{Q}$ alone.

For a subset $E_0 \subset \Phi$ of places declared even, set

$$C_p(E_0) = \sum_{\tau \in \Phi} \sum_{\substack{e \geq 1,\ e \text{ odd or } p=d \\ \tau, -p^e\tau \in \mathcal{T} \setminus E_0}} \frac{a}{p^e}.$$

This collision count removes all pairs for which one of the two places belongs to E_0. Let o_0 be the order of the class of $-p$ modulo N_0 (using the place convention of the data). We call p served if it is inert in K, not short, and o_0 is odd with $o_0 \leq s_p$. These are the definitions `Data.IsShortPrime`, `Data.loss`, `Data.ceilLogM`, `Data.deltaP`, `Data.IsServed` and `Data.collisionSum` in `New/Curve/LiveSum.lean`.

Lemma 10.7 (Collisions without exceptional treatment)**.** *For p small nonsplit and not served,*

$$mC_p(\varnothing) \le \ell_p.$$

Moreover, for every small nonsplit p one has $C_p(\varnothing) \le 2fa/p$.

Proof. If all the exponents appearing in a collision sum are at least o, then

$$C_p(E_0) \le fa \sum_{e \ge o} p^{-e} \le 2fa/p^o,$$

since $p \ge 2$. The second assertion follows by taking $o = 1$. The first assertion for short p follows as soon as $(\log b)^2 \ge 2$. If p is not short, recall that a collision has $p \ne d$ and odd e divisible by o_0 (`Data.IsCollision.ne_and_odd_and_dvd`, `New/Columns/JetExponentBound.lean`). If there are any collisions for such a prime, o_0 is odd; the assumption that p is not served therefore gives $o_0 > s_p$. Consequently

$$mC_p(\varnothing) \le 2mfa/p^{s_p+1} \le 2fa/p \le \ell_p.$$

The source's assertion is `Data.mul_collisionSum_le_loss`. The arithmetic fact excluding collisions at d uses the congruence hypotheses on d; without those hypotheses the argument at a ramified prime would require modification. □

The issue at a served prime is different: exponents $e \le s_p$ can contribute, and multiplying a loss of size a/p by m cannot be absorbed in ℓ_p. The prime-by-prime repair uses a set of split places of E/E^+ even though p is inert in $K/\mathbb{Q}$. Fix ι and denote by $l(\tau)$ the prime of E defined by the embedding attached to τ and ι.

At a served prime, write $u = -p$ in the group of infinite places. Its order is divisible by m:

$$m \mid \mathrm{ord}(u). \tag{10.5}$$

For clarity, the basic 2-power estimate behind this assertion is the following. If an odd integer v is congruent to neither 1 nor -1 modulo 2^h, then the order of its image modulo 2^s has 2-part at least 2^{s-h+1} provided s is sufficiently large. For an even exponent L the identity

$$v_2(v^L - 1) = v_2(v-1) + v_2(v+1) + v_2(L) - 1$$

follows by repeated factorization of $v^{2j}-1$. Since one of $v-1, v+1$ has valuation 1 and the other less than h, the required bound follows from divisibility by 2^s. Here $m = 2^h$, the cyclotomic exponent is $s = 22h$, and the shortness exclusion allows the estimate with ample slack. The local embeddings of the cyclotomic tower, including primes dividing the fixed conductor, are handled in `Data.IsServed.m_dvd_orderOf` in `New/Columns/JetExponentBound.lean`. We use that exact assertion for the admissible data.

The other relevant arithmetic property is

$$l(u^e\tau) = \overline{l(\tau)} \quad (e \text{ odd}), \qquad l(\tau') \in \{l(\tau), \overline{l(\tau)}\} \iff \tau' = u^e\tau \text{ for some } e. \tag{10.6}$$

These are `Data.inducedPlace_neg_pow_mul_of_odd` and `Data.inducedPlace_eq_or_star_iff`, same file. In particular, a served prime of the data is unramified in the cyclotomic tower. One way to see the parity assertion is to express the decomposition group through Frobenius in $\mathrm{Gal}(E/\mathbb{Q})$: its restriction to K is conjugation and its restriction to the cyclotomic part is the pth power. The element $-p$ of the CM place group is obtained by composing this Frobenius with complex conjugation. Thus it interchanges the two induced primes above the real place. The order condition (10.5) ensures that these two primes are distinct. The full arithmetic assertions `smul_eq_iff_inert`, `smul_eq_star_iff_inert`, `q_eq_inert` and the unramified/splitting hypotheses used for these formulas are formal-source inputs for the specific extension E.

Lemma 10.8 (Marking)**.** *Fix a small served prime and a p-adic embedding. There is a finite set Λ of primes above p such that:*

1. *no $l \in \Lambda$ has $\bar{l} \in \Lambda$;*
2. *each marked prime l is induced by a top;*
3. *for every top τ, either $l(\tau)$ or $\overline{l(\tau)}$ belongs to Λ.*

Write $E_0(\Lambda) = \{\tau : l(\tau) \in \Lambda\}$ for the set of even places. Then

$$|\Lambda| \le f, \qquad C_p(E_0(\Lambda)) = 0.$$

Every τ with $l(\tau)$ or $\overline{l(\tau)}$ marked has $\tau \notin \mathcal{S}$.

Proof. For each top, multiplication by u^{o_0} stays within the tops and, because o_0 is odd, interchanges the two induced primes by (10.6). Hence the primes induced by tops occur in conjugate pairs. Choose one from each pair. This gives all three properties and the cardinality bound. Every collision has e odd, so the primes induced by its two top places are conjugate. At least one is marked, and consequently no collision remains outside $E_0(\Lambda)$.

For the last assertion, the orbit condition (10.6) gives a power of u taking the place τ to a top. Its class modulo N_0 is therefore in the cyclic subgroup generated by the class of u. That subgroup has odd order. A thick place has class -1, of order two in the residue group of the fixed modulus N_0; it cannot lie in an odd-order subgroup. The exact statements are `Data.exists_isMarked`, `Data.IsFullMark.collisionSum_evenPlaces_eq_zero` and `Data.IsMarked.notMem_thick`. □

We will also need the balancing property of marking. If $G(c, l)$ is any function of an integer weight and a prime, zero unless $l \in \Lambda$, then

$$\sum_{\tau} G(w_j(\tau), l(\tau)) = \sum_{\tau} G(w_j(\tau), \overline{l(\tau)}). \tag{10.7}$$

Indeed multiplication by u exchanges the two primes. All places in these paired orbits are non-thick by the previous lemma, so the functions w_j are constant there, with the common weight determined by b_2. Reindexing the first sum by $\tau \mapsto u\tau$ proves the equality. The source's exact assertion is `Data.IsMarked.sum_inducedPlace_sub_star_eq_zero`. The equality will ensure that the local factors from the splitting characters in the Hecke transfer are integral.

10.6 Moving the marked Hecke operators through the chain

For each marked l, consider the spherical Hecke operator T_{l,s_p} corresponding to

$$\mathrm{GL}_k(\mathcal{O}_{E_l}) \,\mathrm{diag}(p\mathbf{1}_{s_p}, \mathbf{1}_{k-s_p}) \,\mathrm{GL}_k(\mathcal{O}_{E_l})$$

on the column space. We use the unnormalized coset-sum convention of the source's `heckeSplit`. Put $A_l = a\#\{\tau \in \mathcal{T} : l(\tau) = l\}$ and set $Y_l = p^{A_l} T_{l,s_p}$ and $Y = \prod_{l \in \Lambda} Y_l$ in a fixed order. By the scale constraints p is outside the column levels g and N. The local group is split and p is a uniformizer at l. The use of s_p rather than just the rank-one Hecke operator is part of the local repair strategy.

Let $J_{p,j}(E_0)$ denote the jet exponent for the jth lift established in the preceding section, and let $L_j(\tau)$ denote the number of columns in the weight at τ, below n_{j+1} and of heights other than one. Define

$$D_{p,j}(\Lambda) = s_p \sum_{\tau \in E_0(\Lambda)} L_j(\tau) - J_{p,j}(E_0(\Lambda)).$$

The quantity $D_{p,j}(\Lambda)$ may be negative; it is a logarithmic p-adic gauge exponent, not a loss required to be nonnegative. The definition is `Data.liftLoss` in `New/Columns/LocalStatement.lean`. For the following discussion a jet datum multiplying gauges by at most p^u will mean exactly membership in the source's submonoid `Kernel.raises` with prescribed frames and bound p^u.

Proposition 10.9 (Two local compatibility assertions, formal-source inputs)**.** *Let p be small nonsplit, let Ψ be good for a twist with $p \nmid q_\theta$ and let $l \in \Lambda$. For every $j < m$:*

1. *There are integral coefficients $c_0, \dots, c_{s_p}$ such that, on the partial definite chain $\mathsf{C}^{j+1}_{\theta,\xi}$ with any $\xi \in \operatorname{supp} \Psi$,*
$$\sum_{i=0}^{s_p} c_i T^{(j+1)}_{l,i} \mathsf{C}^{j+1}_{\theta,\xi} v = p^{\delta_p} \mathsf{C}^{j+1}_{\theta,\xi} T^{(0)}_{l,s_p} v$$
for spherical columns v, with spherical target datum at l.

2. *If each u_l belongs to one of the basic double cosets at l of index $i_l \le s_p$, the modified jth datum, translated on the target side by the product of the u_l, multiplies gauges by at most*
$$p^{\sum_{l \in \Lambda} A_l + D_{p,j}(\Lambda)}.$$

The declarations are `Data.IsGoodDatum.exists_hecke_transfer` and `Data.IsGoodDatum.mem_raises` in `New/Columns/LocalStatement.lean`. We explain precisely how the stated estimates enter the chain, and the source of each term.

At a split unramified place, the Hecke polynomial of a theta lift differs from that of the preceding stage by a polynomial with explicit powers of the size of the residue field and values of splitting characters. In the proof the source polynomial is evaluated with a rank-dependent normalization, and the constant term of the extra polynomial is 1. One can therefore recover its coefficient of degree s_p by multiplying the target polynomial by the truncated formal inverse of the extra polynomial. This is simple polynomial algebra: for $U(X) = 1 + u_1 X + \cdots$ over a ring, $U(X)^{-1} \bmod X^{s_p+1}$ has coefficients in the same ring, recursively $v_j = -\sum_{i=1}^{j} u_i v_{j-i}$. The unramified transfer identity for the single kernel is applied successively to all earlier lifts before extracting the coefficient. The weight balance (10.7) controls the norms of the values of the splitting characters at l; the coefficient of extraction uses at most p^{δ_p} to clear the rank normalizations. The exact normalized coefficient identity is the deeper part of assertion (1); it is not a consequence of unnormalized Satake transfer without the denominator computation.

For assertion (2), write each basic double coset representative as an integral invertible matrix times a diagonal matrix times an integral invertible matrix. At the marked places its diagonal has p in a set of at most s_p coordinates, and at conjugate places its entries are inverted by the unitary action. At infinite places this changes the coefficient frame for the local Schur-jet estimate. The top height-one columns contribute a units of gauge cost for each top inducing l, yielding A_l. At the other columns one incurs at most s_p times the number of columns at even places. The tilted Schur-jet bound supplies divisibility $J_{p,j}(E_0(\Lambda))$. The factors from determinant characters cancel between conjugate marked places, again by (10.7). This gives the exponent in assertion (2). It is crucial that the theta-jet estimate applies to

the actual tilted data: the Hecke operator has modified the kernel vectors, rather than merely changed their scalar coefficient by p^{A_l}. The proof of `mem_raises` composes the diagonal-frame jet estimate of the preceding section with integral covariance.

Lemma 10.10 (Iteration of the repair)**.** *For integral columns v and $p \nmid q_\theta$, with Ψ good,*

$$\|\mathsf{B}_\theta^\Psi(Y^m v)\|_\iota \le p^{\sum_{j<m}(|\Lambda|\delta_p + D_{p,j}(\Lambda))}.$$

Here the bound is on every gauge of the restricted chain.

Proof. Move one marked operator through the partial chain ending at $j+1$, using Proposition 10.9(1). On the jth kernel datum this replaces it by the operator

$$p^{A_l-\delta_p}\sum_{i\le s_p} c_i T_{l,i}^{(\mathrm{datum})}$$

on the finite Weil representation on the target side of that kernel. On its gauge the scalar factor contributes $p^{\delta_p - A_l}$. Target-translation covariance and the spherical invariance of the datum guarantee this operator commutes with the required simultaneous level actions. When the same operation is applied at a second marked place, the first remains in the same finite kernel; the local groups at distinct places commute. Repeating for a list I of distinct marked primes gives a new datum at stage j and, for the as yet unprocessed local translates, gauge exponent

$$\sum_{l\in I}(\delta_p - A_l) + \sum_{l\in\Lambda} A_l + D_{p,j}(\Lambda).$$

Indeed the remaining Hecke actions are sums of basic coset translates as in assertion (2), with integral coefficients; the ultrametric inequality does not change the exponent when adding these terms. Taking $I=\Lambda$ leaves $|\Lambda|\delta_p + D_{p,j}(\Lambda)$. Modify every stage by this construction. Inductively the resulting first j kernels act on v as the original partial chain on $Y^j v$: the new last kernel has exactly absorbed one more Y through the partial chain. Thus the modified full chain equals the original chain on $Y^m v$. Lemma 10.3 adds the exponents. Finally the restriction to Y_m does not increase gauges, since $p \ne N$. This is the deduction `Data.IsGoodDatum.exists_move_markOperator` and `Data.IsGoodDatum.gaugeLe_markOperator_pow`. □

Here is the numerical consequence. The jet exponent of the preceding section obeys the exact inequality

$$J_{p,j}(E_0) \ge \frac{g_p}{m} - C_p(E_0) + \omega_p \#(\mathcal{T}\cap E_0)\frac{a}{x_1} - \frac{fa(\log b)^2}{pm}. \tag{10.8}$$

For marked primes the other bound required is

$$|\Lambda|\delta_p + s_p \sum_{\tau\in E_0(\Lambda)} L_j(\tau) \le \frac{\ell_p}{m} + \#(\mathcal{T}\cap E_0(\Lambda))\frac{a}{x_1}. \tag{10.9}$$

This is `Data.IsMarked.sum_lowColumns_le` in `New/Columns/JetExponentBound.lean`. To explain the special top contribution here, the diagram count gives $L_j(\tau) \le w(\tau) + x^{60}$ and the number of boxes in these columns is between $kw(\tau)$ and $kw(\tau)+x^{60}$. Every marked place is non-thick, hence $w(\tau) = a/x_1$. For marked nonempty Λ the prime is served and not short. If $p \le m$, the factor $s_p \le 1 + \log_p m$ times a/x_1 is negligible in ℓ_p/m because of the separation $x \ll x_1$. If $p > m$, then $s_p = 1$ and the served condition forces $o_0 = 1$. Any even place in a marked orbit of tops is therefore itself a top. The contribution a/x_1 per such place is accounted for exactly by the term on the right of (10.9). In both cases $|\Lambda| \le f$, and the

terms $|\Lambda|\delta_p$ and $s_p f x^{60}$ are absorbed by the remaining margin $fa(\log b)^2/(pm)$. The formal calculation makes the constants explicit uniformly up to $p < x_1$; its scale inequalities are used here as formal-source bounds. The large a/x_1 term in (10.9) could not simply be discarded: for p near x_1 it would be too large. It is exactly matched by the positive term in (10.8).

For a full marking, $C_p(E_0(\Lambda)) = 0$ and $\omega_p = 1$. Combining (10.8), (10.9) and Lemma 10.10 gives

$$\|\mathsf{B}_\theta^\Psi(Y^m v)\|_\iota \le p^{2\ell_p - g_p}. \tag{10.10}$$

There is also a useful estimate with no marking at all. One may take $\Lambda = \varnothing$ in Lemma 10.10, with no nontrivial operators. Then (10.8) gives

$$\|\mathsf{B}_\theta^\Psi(v)\|_\iota \le p^{mC_p(\varnothing) + \ell_p - g_p}. \tag{10.11}$$

Writing

$$j_p = \left\lceil \frac{2mfa}{p} \right\rceil,$$

the estimate $C_p(\varnothing) \le 2fa/p$ and scalar linearity imply

$$\|\mathsf{B}_\theta^\Psi(p^{j_p} v)\|_\iota \le p^{\ell_p - g_p}. \tag{10.12}$$

These are the declarations `Data.IsGoodDatum.gaugeLe_Bmu_of_isMarked`, `Data.IsGoodDatum.gaugeLe_collisionSum` and `Data.IsGoodDatum.gaugeLe_pow_jetCeil_smul`, all in `New/Columns/LocalStatement.lean`. Thus the submodule

$$Y^m M + p^{j_p} M$$

has good divisibility for every admissible twist prime to p. In particular, its definition and the gain bound do not involve the twist. What remains is to prove that it has small index; this is a question about algebraic modular forms and the special Frobenius-twisted Hecke operator.

10.7 The Frobenius-twisted elementary-divisor theorem

We give the local theorem in its representation-theoretic form, because its application to the automorphic lattice is not valid unless both the operator and its coordinate normalization agree.

Fix a small served prime p, and a marked place l of E. Put $u = -p$ as above, $P_0 = p^{o_0}$, and let ν be the order of u^{o_0} in the CM place group. The tops inducing l or $\bar{l}$ are exactly the orbit of a chosen one under this element, so they may be labelled by $s \in \mathbb{Z}/\nu$. Since o_0 is odd, the parity of $l, \bar{l}$ alternates along the orbit. The residue field of l has cardinality $Q_l = P_0^\nu$ and ν is even. This uses the equality $q_l = p^{\mathrm{ord}(u)}$ in the served case, as in `Data.IsMarked.hasType_markOperator`. Here ν is a local orbit length and is not the auxiliary level prime N.

Write R_l for the valuation ring of E at l. There are embeddings $\iota_s : R_l \to O_\iota$ whose reductions satisfy

$$\iota_{s+1}(x) \equiv \iota_s(x)^{P_0} \pmod{p}.$$

In the odd letters the embedding in the coefficient representation also includes conjugation, so that the action of $g \in \mathrm{GL}_k(R_l)$ in the sth factor is by

$$g_s = \begin{cases} \iota_s(g), & s \text{ even}, \\ \iota_s((g^{-1})^t), & s \text{ odd}. \end{cases}$$

The weight module is $\mathcal{P} = \bigotimes_{s\in\mathbb{Z}/\nu} \mathrm{Sym}^a O_\iota^k$, represented by polynomials in variables x_{bs}, $0 \le b < k$, of degree a in each letter s. Let $\rho(g)$ denote the above representation. The other coefficient factors, not in this orbit, are left as a separate factor. Write $r_p = s_p$ for the size of the first block of coordinates. A variable is deep if either s is even and $b \ge r_p$, or s is odd and $b < r_p$. Denote the indicator of this condition by $\Delta(b, s)$. For a monomial x^α its depth is $\sum_{b,s} \Delta(b, s)\alpha_{bs}$. These are the exact definitions `TwistingData.letterRep`, `TwistingData.deepIndicator` and `TwistingData.depth` in `New/TwistedHecke/TwistingData.lean`. The dimension of $\mathcal{P}$ over the coefficient field is $D_\mathcal{P} = \binom{a+k-1}{k-1}^\nu$.

The following Euler- and weight-like notation is only local. Put $\lambda = \nu/2$ and $c_* = 2(k + k^2)o_0\nu^2 + 2$. The numerical conditions of the local theorem, with parameters ϵ, y, z, are

$$\begin{aligned}
&0 < 2r_p \le k, \quad \epsilon \le 1 \le y,\\
&\nu ak(\lambda + 2) + a + k < P_0^\lambda, \quad 6c_* + 2 \le y,\\
&2(\nu + 2)k/y \le \epsilon,\\
&(y + \nu(k + 1))kP_0^{3r_p} + k + 2 \le a/y,\\
&\nu a\left(P_0^{-2r_p} + \frac{r_p^2}{kP_0}\right) \le z.
\end{aligned} \tag{10.13}$$

This is `HeckeSizes` in `New/TwistedHecke/Counting.lean`. They allow the Frobenius degree to be large enough for interpolation while retaining a small elementary-divisor budget.

Why the normalized spherical operator has this weight. Write $d_p = \mathrm{diag}(p\mathbf{1}_{r_p}, \mathbf{1}_{k-r_p})$. In a single even letter, the contragredient coefficient action of d_p in the class-sum formula scales monomials by negative powers of p supported on the first block. Multiplying the action by p^a changes the power to the degree in the second block. At an odd letter the inverse transpose interchanges the two prescriptions. Thus multiplying the total Hecke operator by $p^{A_l} = p^{a\nu/2}$ gives precisely the diagonal matrix

$$D_p = \mathrm{diag}(p^{\mathrm{depth}(\alpha)}).$$

There is one factor p^a for every even letter. This explains the normalization of Y_l at the automorphic lattice, not just its complex normalization.

The root inequality. For distinct coordinates b, b' let $y_s = x_{b's}$, $z_s = x_{bs}$ if s is even, and interchange the choices if s is odd. For a fixed letter i define the excess functional

$$\mathrm{exc}_i(\alpha; b, b') = \frac{1}{Q_l}\sum_s \left(P_0^{1+[s-1-i]_\nu}\alpha(y_s) - P_0^{[s-i]_\nu}\alpha(z_s)\right), \tag{10.14}$$

where $[u]_\nu$ denotes the residue in $[0, \nu - 1]$. This is exactly the weight in `TwistingData.excessWeight`. Let $w \in O_\iota$ be a square root of p. For a set of real variable weights c_v, use the Gauss filtration: f is of order at least A if every coefficient of x^α is divisible by $w^{\max(0,\lceil A - \sum_v c_v\alpha_v\rceil)}$. This filtration is multiplicative: for every product of monomials the desired divisibility follows by adding the inequalities, then summing proves it for polynomials.

Lemma 10.11 (The elementary multiplicity count)**.** *If $\Phi \in O[T]$ is a nonzero polynomial, O a ring, and the images u_i of q distinct residue classes of a local ring have pairwise invertible differences in O, and $(T - u_i)^E \mid \Phi$ for all i, then $qE \le \deg\Phi$.*

Proof. The monic linear factors are pairwise comaximal and their powers are pairwise comaximal. Thus their product divides Φ. Multiplication by a monic polynomial adds its degree to the degree of a nonzero polynomial over any nonzero ring. Compare degrees. □

The precise root lemma of the local proof is stronger than a multiplicity count, and the following statement is the root result proved in Theorem 11.1, `TwistingData.root_lemma` in `New/TwistedHecke/RootLemma.lean`: suppose b, b' are in the same block of a Levi subgroup and, for every letter s,

$$\frac{2}{5} \le c_{y_s} - c_{z_s} \le \frac{8}{5}.$$

Suppose every Levi translate of F is in the Gauss filtration of order $l + \sum_v c_v \alpha_v$, and $w^{l+1} \nmid F_\alpha$. Then $\mathrm{exc}_i(\alpha; b, b') \le 0$ for each i. One cannot obtain this directly by counting zeros of the individual coefficient polynomials of degree a, because the Levi action uses different Frobenius powers in different letters.

We describe the reduction in the formal proof. Choose the transvection from b to b' composed with interchange of the two coordinates, with parameter $t \in R_l$. At a letter s the parameter in the variable substitution reduces to a signed $P_0^{[s-i]_\nu}$-power of the parameter at i. The coefficient of a selected monomial is read after identifying the two variables in letter i, and substituting appropriate powers of an auxiliary variable T in the other letters. More precisely, the exponents assigned to z_s in the selected coefficient are $Q_l P_0^{[s-i]_\nu}$ and those recording the forced zeros from y_s are $P_0^{1+[s-1-i]_\nu}$. Modulo w^{l+1} the hypothesis on F_α gives a nonzero polynomial in T whose degree is at most

$$Q_l \sum_s P_0^{[s-i]_\nu} \alpha(z_s).$$

Its Taylor coefficients in degrees below $\sum_s P_0^{1+[s-1-i]_\nu} \alpha(y_s)$ vanish at every residue class t in R_l. The nonzero coefficient assertion in this calculation uses the minimum for the Gauss filtration: otherwise different monomials could contribute to the selected coefficient and invalidate the argument. The discarded substitutions differ by terms of higher Gauss order; this follows from the Frobenius congruence modulo $p = w^2$ and the two-sided bounds on the differences of variable weights. Thus the multiplicity count gives $Q_l \sum_s P_0^{1+[s-1-i]_\nu} \alpha(y_s) \le Q_l \sum_s P_0^{[s-i]_\nu} \alpha(z_s)$, which is the asserted inequality. This describes the coefficient calculation of the root lemma; the general-ring Gauss-filtration substitution estimates which justify its nonzero coefficient and all its divisibilities are part of that input. The essential gain is that each residue class contributes the same weighted multiplicity, and there are Q_l of them.

Two dimension estimates. Reduce coefficients modulo the maximal ideal and call the residue coefficient space $\overline{\mathcal{P}}$. Let $\mathcal{P}_{\mathrm{lev}}$ be the two-block parahoric of $\mathrm{GL}_k(R_l)$, namely the inverse image of the two-block parabolic over its residue field. Fix the adjacent coordinates $(r_p - 1, r_p)$. Define the subspace U_h of $\overline{\mathcal{P}}$ by requiring in every block-Levi translate that monomials of excess greater than h at this adjacent pair vanish, in every letter. Let V_j be the subspace of $\overline{\mathcal{P}}$ of forms such that every $\mathrm{GL}_k(R_l)$ translate has depth at least j. These are `TwistingData.excessCond` and `TwistingData.deepForms`.

The two relevant estimates are

$$[\mathrm{GL}_k(R_l) : \mathcal{P}_{\mathrm{lev}}] \dim U_y \le (1+\epsilon) D_{\mathcal{P}}, \tag{10.15}$$

$$\sum_{0 \le j < J} \dim V_j \le (1+z) D_{\mathcal{P}} \quad (J \ge 0). \tag{10.16}$$

These are `TwistingData.index_mul_finrank_excessCond_le` and `TwistingData.sum_finrank_deepForms_le`, in `New/TwistedHecke/Counting.lean`. We explain the combinatorics yielding these bounds. The uniform Frobenius minimum-weight calculations in their proof are proved in Chapter 11, with the numerical bounds specified below.

The first count uses the following cyclic interlacing cone. For a real array v_{bs} put $\beta(b,s) = \pm(P_0 v_{b,s+1} - v_{b,s})$, with sign positive in the even letters and negative in the odd letters. This is an invertible linear transformation: if $P_0 v_{b,s+1} = v_{b,s}$ for every letter, then cycling gives $(P_0^{\nu} - 1)v_{bs} = 0$. Substitution into (10.14), telescoping around the letters, gives

$$\mathrm{exc}_i(\beta; b, b') = (1 - Q_l^{-1})\big(\beta(y_i) - (v_{b,i} - v_{b',i})\big).$$

Thus the zero-excess inequalities are exactly the interlacing conditions

$$\begin{cases} P_0 v_{b+1,s+1} \le v_{b,s}, & s \text{ even}, \\ v_{b+1,s} \le P_0 v_{b,s+1}, & s \text{ odd}, \end{cases} \qquad 0 \le b < k-1.$$

Let $\Sigma(D)$ be the product of slabs $\beta_{bs} \ge 0$, $D_s - 1 \le \sum_b \beta_{bs} \le D_s$. The cone of zero excess, intersected with $\Sigma(D)$, has volume exactly

$$\frac{\mathrm{vol}(\Sigma(D))}{\displaystyle\prod_{j=1}^{k} \frac{Q_l^j - 1}{Q_l - 1}}. \tag{10.17}$$

Here is the signed-volume identity underpinning this formula. For decreasing real vectors x, y of length k, the matrix $\mathbf{1}_{y_c < x_b}$ has determinant 1 precisely when $x_1 > y_1 \ge x_2 > y_2 \ge \cdots \ge x_k > y_k$; otherwise it has determinant zero. This follows because the row of index b ends with a string of ones, preceded by zeros, and the nested rows are independent exactly when all their threshold positions differ. Take one such matrix of interlacing indicators for each pair of adjacent letters, using transpose in the alternating letters. Expanding the product of determinants yields a signed sum over tuples of permutations, whose product permutation will be denoted π'. The corresponding linear changes of variables $v \mapsto \beta$ have determinant, up to a common factor independent of π',

$$\det(1 - Q_l^{-1} A_{\pi'}),$$

where $A_{\pi'}$ is the permutation matrix. This follows by solving the equations around the cyclic chain after successively reordering the coordinates: the sole boundary condition is given by the product permutation. The walls where equalities occur have measure zero. Consequently integrating the signed indicators and accounting for sorting yields

$$\frac{\mathrm{vol}(\mathrm{cone} \cap \Sigma(D))}{\mathrm{vol}(\Sigma(D))} = \frac{(1 - Q_l^{-1})^k}{k!} \sum_{\pi' \in S_k} \frac{\mathrm{sgn}(\pi')}{|\det(1 - Q_l^{-1} A_{\pi'})|}.$$

The formula, including the cancellation of all intermediate permutation multiplicities, is `volume_inter_slab_eq_signed_sum` in `New/TwistedHecke/InterlacingCone.lean`. To evaluate the signed sum, it suffices to compare with the single-letter case with parameter Q_l instead of P_0. In that case the inequalities are $\beta_b \ge Q_l \beta_{b+1}$. The triangular change

$$\beta_b \longmapsto d_b(\beta_b - Q_l \beta_{b+1}), \qquad d_b = 1 + Q_l + \cdots + Q_l^b \quad (b = 0, \ldots, k-1)$$

with $\beta_k = 0$ carries the cone onto the nonnegative orthant. It preserves the sum of all coordinates because $d_b - Q_l d_{b-1} = 1$ and $d_0 = 1$; its determinant is $\prod_b d_b$. This gives (10.17). In the source this step is `prod_mul_volume_inter_slab`. The factor in the denominator is exactly the number of complete flags in a k-dimensional vector space over the residue field of size Q_l. This exact match explains why a leading constant $1 + \epsilon$ in (10.15) is possible.

The root lemma relates this volume to the dimension of U_y. To see the requisite amplification, one passes from a polynomial F to blocks of polynomials of the form $x^{\gamma(u)} F_u(x^D)$

with $D = Q_l^J$, and $u \in \{0, \dots, D-1\}^{k\nu}$. The exponents $\gamma(u)$ are congruent to u modulo D and have degree in each letter between kD and $(k+1)D$. Hence these summands are independent, and substitution $x \mapsto x^D$ commutes with the residual group action because D is a full Frobenius period. One chooses a least-weight monomial among its translates over the flags in the Levi subgroup. The root lemma at adjacent roots within either block forces the corresponding excess inequalities, with the chosen bound y retained at the single boundary root between blocks. A count of its available coefficients and the amplification limit $D \to \infty$ bounds the dimension by the number of Levi flags times the volume of the thickened cone of boundary excess $y + \nu(k+1)$. This is proved in Lemma 11.3, corresponding to `TwistingData.finrank_excessCond_le_volume`; it identifies the injection obtained by taking the relevant coefficients of the translates of the amplified polynomials.

To remove the thickening, the proof constructs a nonnegative vector σ with all coordinates at most $P_0^{3r_p}$, satisfying each zero-excess inequality and having boundary excess at most -1. One can see such a vector by taking $v_{b,s}$ proportional to $2P_0^{2(r_p-1-b)}$ for $b < r_p$ and zero for $b \geq r_p$, with the odd letters multiplied by P_0. Substituting into the linear map defining β and into the interlacing inequalities gives this bound. The assertion is `TwistingData.exists_mem_excessCone_mem_Icc`. Translation by $(y + \nu(k+2))\sigma$ embeds the thickened cone into the zero-excess cone; it changes the upper endpoint of each slab by at most $(y + \nu(k+2))kP_0^{3r_p}$, with allowance for the initial degree padding. The precise choice in the Lean proof uses $y + \nu(k+2)$ with each letter's shift and the size conditions (10.13). The volume of a slab of upper endpoint D_s is at most $D_s^{k-1}/(k-1)!$, as is seen by subtracting the volume of two simplices. With the size conditions the changed volume of all slabs is at most $(1+\epsilon)D_{\mathcal{P}}$ after multiplying by the flag count and canceling it against the denominator in (10.17). The exact inequalities including the roundings are those of `index_mul_finrank_excessCond_le`. They supply (10.15).

For the depth bound assign to each deep variable of coordinate $b \geq 2r_p$ cost $P_0^{-2r_p}$, and to each deep variable with $b < 2r_p$ cost $1/P_0$. All other variables have cost zero. The result proved in Proposition 11.6, `TwistingData.exists_le_weight_of_mem_deepForms` in `New/TwistedHecke/DeepFormCost.lean` says that any nonzero polynomial in V_j has a nonzero coefficient at a monomial of total cost at least j, provided $2r_p \leq k$ and $2r_p < \nu$. The source uses the upper unipotent matrix $1 + A$ with A supported in the block from the first r_p rows to the remaining rows and $A^2 = 0$. In the rational function field of the variables, the coefficients of A can be recovered by linear equations in Frobenius powers of the deep variables; evaluating each letter's variables in the specified order gives a triangular matrix with nonzero determinant. The equations express each matrix entry at residue points by polynomials of cost less than one. The condition of depth j in every translate then means vanishing to that weighted order on the whole family of residue points. Polynomial interpolation forces the existence of the claimed costly monomial. The recovery matrix and the weighted vanishing estimate are the substantive content of this input, rather than a generic vanishing-order assertion.

The rest of the depth argument is a direct count. Restriction to coefficients of cost at least j is injective on V_j by the stated input. Summing the resulting dimension bounds counts each monomial at most 1 plus its cost. There are $\nu/2$ deep letters for each fixed coordinate. Symmetry of the monomial basis gives average exponent a/k at each coordinate of each letter. Thus the mean cost over all monomials is at most

$$\frac{a}{k}\frac{\nu}{2}\left(2r_pP_0^{-1} + (k-2r_p)P_0^{-2r_p}\right) \leq \nu a\left(P_0^{-2r_p} + \frac{r_p^2}{kP_0}\right) \leq z.$$

This proves (10.16), using only the lower bound on cost in V_j and the last size condition in (10.13).

The approximate projector and the integral pairing. The dimension estimates must be converted into a determinant bound; they do not by themselves control Y_l. In particular, bounding the dimension of a conditioned representation modulo p does not give a bound on all Smith divisors of an arbitrary matrix with that representation in its target.

In the source, the approximation is constructed using an invariant diagonal pairing on degree-a polynomials. For a multi-index α of degree less than p^F, its multinomial coefficient has p-valuation at most (#variables)F. Indeed Legendre's formula for the factorial expresses the valuation as a sum of carry contributions at at most F levels, bounded by the number of input summands at each level. Consequently one may choose integral diagonal entries proportional to the inverse multinomial coefficients, preserving transpose covariance of the action. The exact formula is `ClassMean.exists_invariant_diagonal_form` in `New/TwistedHecke/ClassMeans.lean`.

A second input in that file, `ClassMean.exists_mean`, constructs an invariant linear form on polynomials of a generic Levi matrix. If x is a weight function on the matrix coordinates and P is a condition on the row weight, the linear form kills monomials whose row and column weights differ; on monomials satisfying P it agrees with evaluation at the identity multiplied by p^{k^2F}, up to the prescribed degree cutoff. It is in the span of Levi translates of linear forms killing monomials not satisfying P. The hypotheses are degree cutoff $D < p^F$ and $D + k < \#(R_l/\mathfrak{m}_l)$. For orientation, the finite-residue interpolation behind this fact is elementary: evaluation on the matrices with unit determinant spans the dual of the polynomial coefficient space in degrees D if $D + k$ is less than the residue cardinality. A polynomial vanishing on all such matrices, multiplied by their determinant, vanishes on the entire finite matrix space. The degree bound and successive one-variable root bounds force it to be zero. Lifting linear independence modulo the maximal ideal gives the integral statement on the finite free coefficient module. The invariant trace construction which incorporates the row-weight restriction is part of `exists_mean`.

Taking tensor products of these class means in the cyclic letters gives the approximate projector. In the Frobenius product of generic matrices the factor at the dth position is the entrywise P_0^d power of the generic matrix from letter $s - d$. At each letter one takes the first λ factors, plus one additional factor in each deep row; adjugate transpose takes the place of ordinary multiplication in odd letters. The construction tracks the Gauss weight and hence the excess of each output monomial. If the generic monomial has equal row and column weights outside a single letter, and the discrepancy between its weight at that letter and the output excess is greater than $2c_* + 1$, the coefficient of this monomial is divisible by w^{c_*}. This is proved in Proposition 11.9, corresponding to `TwistingData.pow_dvd_coeff_conj_frobOperator` in `New/TwistedHecke/FrobeniusFactors.lean`; it follows from the multiplicative Gauss weight bound `TwistingData.frobOperator_mem` and the Frobenius congruence for conjugated generic matrices in the same file. That bound is consistent with the local scaling $p = w^2$.

The resulting precise assertion, proved in Theorem 11.10, is `TwistingData.exists_projector`, in `New/TwistedHecke/ElementaryDivisors.lean`. There are a diagonal pairing matrix b_* and an operator Z_* such that Z_* maps every coefficient vector into the integral excess condition of bound y modulo w^{c_*}, and

$$y_0^t b_* Z_* F \equiv p^{(k+k^2)o_0\nu^2} y_0^t F \pmod{w^{c_*}}$$

for y_0 in the smaller excess condition with threshold $y - (4c_* + 2)$ modulo w^{c_*}. Here Z_* is not asserted to be an idempotent. The transpose compatibility $b_*\rho(\gamma^t) = \rho(\gamma)^t b_*$ is part of the statement. The precision $c_* = 2(k + k^2)o_0\nu^2 + 2$ leaves two extra powers of w beyond the scalar in the congruence. This pairing estimate is the nontrivial bridge between the counting result and the determinant calculation.

We describe the final abstract determinant theorem, proved in Section 11.8, so that its applicability to the coefficient lattice can be verified. In the following statement let X be a finite class set, I a coefficient index set identified with the monomial indices of $\mathcal{P}$ times a finite set H_0 of spectator indices (the remaining coefficient factors). Let $J = \mathrm{GL}_k(R_l)/\mathcal{P}_{\mathrm{lev}}$. Suppose for $e = (c, b) \in X \times J$ we have class $s(e) \in X$, matrix $g_e \in \mathrm{GL}_k(R_l)$ and integral invertible matrix u_e on the spectator indices, with

$$e \longmapsto (s(e), [(g_e^t)^{-1}]) \in X \times J \quad \text{injective.}$$

If γ_b are the chosen parahoric-coset representatives and $F = (F_c)_{c\in X}$ is a family of coefficients, consider

$$(TF)_c = \sum_{b\in J} \Big(u_e \otimes (\rho(g_e) D_p \rho(\gamma_b))^t\Big) F_{s(e)}.$$

With the orientation of representatives chosen in this formula, the theorem `TwistingData.hasType_twisted_hecke` states that on any linear image of the coordinate module to which T descends, the descended map preserves that module and has type bound

$$(z + \epsilon s + y)|X||I| \quad \text{modulo } p^s \tag{10.18}$$

when (10.13) holds. Its algebraic proof is in Section 11.8; application to the actual adelic coefficient lattice still requires the concrete compatibility arguments.

For its proof the map $F \mapsto (w^{\mathrm{depth}} \rho(\gamma_b) F_{s(e)})_e$ is split by a Hermite–Smith factorization into a primitive factor and a square factor. The determinant of the square factor divides a power of w with exponent at most the sum of dimensions of the depth spaces over the required truncation range. One way to read this bound is to reduce a column of the Smith factor modulo the maximal ideal: a column whose divisor exceeds w^j is forced into the jth depth condition in all frames. The independence of the reduced primitive columns bounds the count of such divisors by the dimension of that condition. This is `TwistingData.exists_depth_factorization`, which also ensures the primitive factor's coefficients lie in the integral excess condition with prescribed precision. Apply this factorization to each of the two half-depth factors of D_p (using $w^2 = p$). The half-factors are paired through Z_* so that the scalar $p^{(k+k^2)o_0\nu^2}$ controls the primitive part of the product, except for excess-condition vectors of total dimension bounded by (10.15). There is a second Smith factorization on the discrepancy matrix between the two primitive images. It gives the abstract pairing assertion `TwistingData.of_pairing`: if $RQ = 1$, $\beta M \equiv a_0 I \pmod{a_0 w}$, and the normalized difference $Qu - Mv = a_0 w F$ imposes the excess condition on F, then the image of βQ enlarged by p^s has index dividing $a_0^{|I'|} p^{sl}$, with $|I'| + l$ bounded by the dimension of the excess condition (including the class and coset multiplicities). Here I' is the domain index set of the primitive matrices. This is proved in the source by applying Smith form to $(QRM - M)R$, assessing each divisor relative to $a_0 w$, and using independence modulo the maximal ideal. The counts (10.15) and (10.16) now account for the three terms ϵs, z, and y in (10.18): the first is the exceptional share of the cokernel charged at depth s, the second is the aggregate half-depth cost, and the third pays for the controlled scalar in the projector. These contributions are additive by Lemma 10.2. This description specifies the remaining determinant theorem needed in the proof, including its hypotheses; proving the full matrix approximation independently would require the exact Smith-factor and projector identities, beyond the counting arguments already given here.

10.8 Application to the actual column lattice

The abstract coefficient operator above occurs on M itself, rather than on a putative subspace with convenient coefficients. This fact uses the trivial class stabilizers at K_0 and the availability

of class representatives integral above p. The latter assertion is `exists_reps_integralAbove`, used in `New/Columns/TwistedHeckeType.lean`. For each class representative γ_c, write the Hecke sum at the point γ_c using the local double-coset representatives indexed by $\mathrm{GL}_k(R_l)/\mathcal{P}_{\mathrm{lev}}$. Each product is rationally equivalent to one chosen representative $\gamma_{s(e)}$, times a compact-level element. Triviality of class automorphism groups gives injectivity of $e \mapsto (s(e), [(g_e^t)^{-1}])$: if the class and the parahoric coset agree, the two chosen rational identifications differ by an automorphism of the representative, which must be the identity, and the local cosets agree. The other primes above p give integral invertible matrices on the spectator factors. Thus in the coordinate identification of algebraic forms the normalized Hecke operator has exactly the shape of (10.18), including the depth diagonal. The exact comparison, with the class-representative and gauge conditions, is `Columns.hasType_markOperator` in `New/Columns/TwistedHeckeType.lean`. We use this coordinate comparison as a formal-source input. It is essential that the integral image addressed by that theorem is all of M, as opposed to the coefficient module with unrelated adelic transition functions.

The local congruence and numerical hypotheses are verified for our marked place by `Data.IsMarked.hasType_markOperator`, in `New/Columns/MarkedPrimeType.lean`. The precise result is the following.

Proposition 10.12 (Marked operator type)**.** *For each $l \in \Lambda$ and each integer $s \geq 0$, the operator Y_l preserves M and*

$$Y_l \in \mathrm{Type}\left(p, M, s; \left(\frac{128 A_l \log b}{mp} + \frac{s}{x^{14}} + x^{40}\right) c\right).$$

Indeed on the local orbit discussed above one has $A_l = a\nu/2$, the parity of the local embeddings is alternating, and the orbit size is even. The choices of the generic theorem's parameters are

$$\epsilon = x^{-14}, \qquad y = x^{40}, \qquad z = \frac{128 A_l \log b}{mp}.$$

The relation $P_0^{\lambda} \geq p^{m/2}$, from the divisibility of the order of u by m, supplies the interpolation inequality in (10.13). The served condition $o_0 \leq s_p$ gives $P_0 \leq p^{s_p} \leq pm$. The bounds $f \ll x^{11}(\log b)^{1/5}$ up to fixed factors, $k \asymp x^5$, $\nu \leq f$ and the hierarchy $x \ll x_1 \ll b$ leave room in the remaining inequalities; for example $P_0^{-2r_p} \leq m^{-2}$, and $r_p^2/(kP_0)$ is negligible against $\log b/(mp)$ on the admissible range. More precisely $P_0^{2r_p} \geq p^{2s_p} \geq mp$ since $P_0 \geq p$ and $m, p \leq p^{s_p}$. The large-prime range is covered by the separate estimate $s_p = 1$ when $p \geq m$. The stated Lean declaration checks these estimates with the fixed constants in (10.13), including the degree condition when ν and x_1 vary. Thus the proposition is a quantitative application of the coordinate comparison and the local theorem, not an assumption about Hecke eigenvalue slopes.

Corollary 10.13 (Type bound for the repair)**.** *With Y and j_p as above there exists an integer $e_p \geq 0$ such that*

$$I(M, Y^m M + p^{j_p} M; p^{e_p}), \qquad e_p \leq 3\ell_p c.$$

Proof. By Lemma 10.2 and the proposition, the exponent bound for the product of the marked operators, repeated m times, is at most

$$cm \sum_{l \in \Lambda} \left(\frac{128 A_l \log b}{mp} + \frac{j_p}{x^{14}} + x^{40}\right).$$

Now $\sum_l A_l \leq fa$, $|\Lambda| \leq f$, and $j_p \leq 2mfa/p + 1$. Thus the first contribution divided by $c\ell_p$ is at most $128/\log b$, the second at most

$$\frac{2m^2 f}{x^{14}(\log b)^2} + \frac{mp}{a x^{14}(\log b)^2},$$

and the third at most $mx^{40}p/(a(\log b)^2)$. All are at most 1 on the admissible scale, the second using the bound on f and $m^2 \le 4x$, and the third using $p < x_1$ and $x_1^{10^6} \le b$. The precise three-term comparison with the fixed constants is `Data.IsMarked.hasType_prod_markOperator_pow`. The type condition gives integrality and an integer witness e_p for the desired index. □

The crucial point is that the type estimate grows with the *truncated mean* of elementary divisors rather than with their largest possible value. Charging the whole space for one large elementary divisor would multiply ℓ_p by m and defeat the saving.

10.9 The prime N: moving denominators to a Fourier operator

The additional prime N splits completely in E and does not appear among the small nonsplit primes. Its effect in the complex bound is the large negative term in (10.3). At the finite places of N one must recover a usable lattice estimate despite the large negative powers of π in the restriction and the translations by y/N in the good datum.

Fix ι above N and let S be the set of conjugates of the primes induced by the given CM-type embeddings at ι (`Data.nonInducedPlaces`). At these split places use the coordinate embedding of the general linear group into the unitary group. Let d_j be the finite point whose component at the places in $S \cap \{v \mid \pi\}$ is $\mathrm{diag}(1, \dots, 1, N)$, with the corresponding conjugate component and identity elsewhere. These are `Data.diagPointN`. In the representation conjugated by d_j, the column level has mirabolic reduction, i.e. its last column is the final basis vector modulo N. The relevant level identity is `Data.atPlaces_conj_mem_mirabolic` in `New/Seesaw/LiftCharacters.lean`.

Let F_N be the Fourier operator on columns, given in this conjugated representation by multiplication by the inverse last-coordinate diagonal and a sum over integral matrices with prescribed final row and column, weighted by the additive character of their final residue coordinate. The formula in `IsSplitLevel.fourierOp` (`New/Seesaw/NamesAtSplitLevel.lean`) is

$$F_N = \rho(\mathrm{diag}(1, \dots, 1, N^{-1})) \sum_{y,\ y_{\mathrm{last}}\ \mathrm{unit}} \psi_N(y_{\mathrm{last}})\rho(\widetilde{g}_y).$$

All coordinates here are at the places of S, with the conjugated representation understood. Each $\widetilde{g}_y$ is an integral invertible lift of a matrix having final row and column determined by y. The additive character takes root-of-unity values. This operator preserves the column level (`Data.fourierOp_mem_column`) and is compatible with the ramification condition at N.

The key finite calculation is local and independent of the large weights. If L is a single theta-lift law, $\mathbf{b}$ is the indicator of the final-column condition and $\mathbf{c}$ the associated additive-character datum, then moving the diagonal off the character datum gives the Fourier operator applied to the indicator lift; it also equals a scalar times the indicator lift applied to the Fourier operator on the source. In the notation of `LiftLaw.u$_0$_charDatum_eq_fourierOp`, this is

$$\rho'(u_0)L[\mathbf{c}]v = F_N' L[\mathbf{b}]v = \gamma L[\mathbf{b}]F_N v.$$

The equality uses ramification: vectors in the image of $1 - z_v$ kill the terms in the finite character expansion which are invariant under that local scalar translation. One can consequently repeat the identity $m/2$ times to move the weighted Fourier operators from the paired chain onto the beginning of the column chain. This iteration is `compUpTo_pairs_eq_smul` in `New/Seesaw/NamesAtSplitLevel.lean`; its proof tracks the number of residue columns nonzero at every place and the scalar splitting-character factors at each stage. The finite iteration is proved in Proposition 11.14, conditional on the exact normalized lift law; that

actual geometric normalization is still a formal-source input. Applying it to the paired expression of the good datum yields `Data.chain$_2$_pairTranslates_eq_smul_chain`, `New/Columns/TwistedHeckeType.lean`: for every admissible twist, the paired chain of m kernels on a ramified column is a scalar of N-adic absolute value at most N^{2mfb} times a chain with data cut by $\mathbf{b}$, read in the frames d_j, acting on $F_N^{m/2}v$.

Each new cut kernel multiplies gauges by at most N^{fb}. This is `Data.name_lastColDatum_mem_raises`, same file; it is the crude CM integrality estimate applied to the cut datum rather than to arbitrary translates with denominator N. The branching restriction has the same gauge bound at N when its source gauge is computed at d_m: the negative powers of $\bar{\pi}$ are canceled by the corresponding final-coordinate powers in the branching formula. This is `Data.GaugeLe.restrictToY_of_diag` in `New/Seesaw/LiftCharacters.lean`. Both are used as precise formal-source inputs. Consequently, by Lemma 10.3,

$$\|\mathsf{B}_\theta^\Psi(v)\|_\iota \le N^{3mfb} \tag{10.19}$$

if $v \in C^{\mathrm{ram}}$ and $F_N^{m/2}v$ belongs to the column lattice measured at the frame d_0. This is `Data.gaugeLe_Bmu_at_N`.

The index cost of imposing this condition is much smaller than charging every coefficient its worst weight at the last coordinate. To see this, for each coefficient index I let $\alpha_I(\tau, b')$ be its tableau weight in coordinate b' at τ. For the initial column space these are just monomial exponents, and the sum over the coordinate b' in any factor is at most a: the factor is of degree a on the tops and degree zero elsewhere. With $\mathcal{X}$ the class set and $\mathcal{I}$ the Weyl indices put

$$T_{b'} = \sum_{c' \in \mathcal{X}} \sum_{I \in \mathcal{I}} \sum_{\tau \in \Phi} \alpha_I(\tau, b').$$

Coordinate symmetry of the initial weight implies $T_{b'}$ is independent of b', and the degree bound gives $\sum_{b'} T_{b'} \le fac$. In particular $T_{\mathrm{last}} \le fac/k$. The class-space action in the chosen coefficient frames is diagonal for the last-coordinate scaling: it multiplies a coefficient by $N^{\sum_\tau \alpha_I(\tau,\mathrm{last})}$, so its determinant is $N^{T_{\mathrm{last}}}$.

Precisely, if the column lattice at a frame $d_0 d(S')$ is changed to the lattice at d_0, where $d(S')$ is the inverse last-coordinate diagonal supported on a set of split places, the corresponding index divides $N^{T_{\mathrm{last}}}$. The adelic integrality here uses rational class representatives adjusted at N; the coefficient norm computation is a direct diagonal determinant comparison. The proof in the source pulls back the coordinate-diagonal index condition using the integral representative correspondence. Applying the Fourier operator consists of such a diagonal frame change and a sum of translations by integral invertible matrices; its character coefficients are N-adic units. Applying the pullback and transitivity rules of Lemma 10.2 once for each application of F_N, in addition to the initial frame change, gives an index bound dividing $N^{(m/2+1)T_{\mathrm{last}}}$. This is the assertion `Data.exists_indexDvd_fourierOp_pow` (`New/Seesaw/LiftCharacters.lean`), including the conditions on the changes of representatives. Since $a \le 2xb$ and $k \gg x^5$, the exponent is at most $4mfbc$ on the admissible data. Thus there is a submodule M'_N of M independent of the twist with

$$I(M, M'_N; N^e), \quad e \le 4mfbc,$$

on which (10.19) holds for all ramified columns. The presence of the average tableau weight rather than the largest coefficient degree is another place where symmetry of the integral weight representation is important.

10.10 The simultaneous local divisibility statement

Define a function on rational primes by

$$\epsilon_p = \begin{cases} 3\ell_p, & p < x_1,\ \chi_K(p) \neq 1, \\ 4mfb, & p = N, \\ 0, & \text{otherwise.} \end{cases}$$

Extend g_p by zero off the small nonsplit primes. For p outside those primes and $p \neq N$ the elementary integrality bound from the theta data suffices: every translated kernel is standard with integral gauge away from the twist conductor, and its scale $\bar{\pi}_g$per compensates for the denominators. There is no division by d, since in that case $p \neq d$. The resulting all-column estimate $\|\mathsf{B}_\theta^\Psi v\|_\iota \leq 1$ is `Data.IsGoodDatum.gaugeLe_one` in `New/Columns/LocalStatement.lean`. It is the application of the crude jet bound and Lemma 10.3 to the exact unmodified data.

Proposition 10.14 (Local divisibility and index)**.** *For every p, every ι above p, and every integer $\beta < g_p - \epsilon_p + 1$, there is a submodule M_p' of M and an integer $e \geq 0$ such that*

$$I(M, M_p'; p^e), \qquad e \leq \epsilon_p c.$$

For every admissible twist θ prime to p, every good datum Ψ and every $v \in M_p' \cap C^{\mathrm{ram}}$, the restricted lift $p^{-\beta}\mathsf{B}_\theta^\Psi v$ is integral at ι.

Proof. For a served small nonsplit prime take $M_p' = Y^m M + p^{j_p} M$ and the index supplied by Corollary 10.13. The two terms have gauges at most $p^{2\ell_p - g_p}$ and $p^{\ell_p - g_p}$ respectively by (10.10) and (10.12); the same bound $p^{2\ell_p - g_p}$ holds for their sum. Since $\ell_p \geq 1$, the condition on β gives $\beta + 2\ell_p - g_p \leq 0$, so multiplication by $p^{-\beta}$ makes the gauge at most one. For an unserved small nonsplit prime use $M_p' = M$, the collision bound Lemma 10.7, and (10.11). At N use the preceding frame-change index and (10.19); the condition $\beta < 1 - 4mfb$ ensures $\beta + 3mfb \leq 0$ since $mfb \geq 1$. At the other primes integrality is unconditional and the hypothesis $\beta < 1$ gives $\beta \leq 0$. This proves the proposition, including the common choice of submodule before the twist. It is exactly `Data.exists_indexDvd_Bmu_divisible`. □

The loss function is summable with the following quantitative bound:

$$\sum_p \epsilon_p \log p \leq B. \tag{10.20}$$

Here is a direct calculation, with generous constants for the scale comparison. If $p \equiv 1 \pmod m$ is prime, $p = um + 1$ with $u \geq 1$, so $m/p \leq 1/u$. If $p \equiv -1 \pmod m$, then $p = um - 1$ with $u \geq 1$, so $m/p \leq 2/u$ since $m \geq 2$. Dropping the primality conditions and bounding $\log p$ by $\log x_1$, one obtains

$$\sum_{p < x_1} \frac{\log p}{p} \begin{cases} m, & p \text{ short}, \\ 1, & \text{otherwise} \end{cases} \leq 4 \log x_1 (1 + \log x_1).$$

Thus the small-prime contribution is bounded by $12fa(\log b)^2 \log x_1(1 + \log x_1)$. Since $(\log x_1)^6 \leq \log b$ and $\log b$ is large, this is less than $B/2$. At N, one has $4mfb \log N \leq 120fm^2 b \log b \leq 480fa \log b$, using $m^2 \leq 4x$ and $a \geq xb$. This too is less than $B/2$ for the data in question. The statement with the precise admissibility thresholds is `Data.sum_allowedLoss_mul_log_le` in `New/Columns/SizeEstimates.lean`. It is the sparseness of the congruences for short primes which allows the factor m in ℓ_p and makes the global error budget independent of an overall factor m.

10.11 Archimedean growth and the level normalization

We now compare the finite-prime gains with the growth at every complex conjugate. Define

$$h_0 = m\left(\sum_{\tau\in\mathcal{T}}(J(a) - A(a,\tau)) + k\mathcal{U}\right) - \#\mathcal{T}\, a\log N.$$

Since $A(y,\tau) = 2\sum_p A_p(y,\tau)\log p$, it follows identically from the definition of $\mathcal{U}$ and g_p that

$$h_0 + 2\sum_p g_p \log p = m\left(\sum_{\tau\in\mathcal{T}} J(a) + k\sum_{\tau\in\Phi} J(w_0(\tau))\right) - \#\mathcal{T}\, a\log N. \tag{10.21}$$

This quantity is the leading exponent of the purely archimedean chain. In particular its expression in terms of prime asks is not a second analytic estimate; it is an exact algebraic identity, used to record the finite gain to be divided out.

The precise norm assertion of the definite chain is `Data.weylNormSq_Bmu_le`, `New/Columns/LocalStatement.lean`. For the trivial twist, good datum Ψ, and any conjugate σ, if all values of a column v have squared Weyl norm at most e^η, all values of its restricted lift obey

$$\|\sigma\mathsf{B}_1^\Psi v(g')\|^2 \le \exp\left(\eta + h_0 + 2\sum_p g_p\log p + B\right). \tag{10.22}$$

Here the norm assertion is conditional on the conjugate-jet bounds of the preceding section. We explain the size comparison to isolate the exact quantities going into its error budget.

For one lift the principal contribution from the Young diagram is the sum of $J(l)$ over its row lengths. The uniform diagram estimate `Data.diagramWeight_le_by_index` gives each row length at place τ the bound

$$\begin{cases} \kappa_\tau(1) + b_1 + b_2 + h_2^2, & \text{row zero,} \\ w_0(\tau) + h_2^2, & \text{rows } 1,\dots,k-1, \\ h_2^2, & \text{remaining rows.} \end{cases}$$

Positive error terms in the first row can be regarded as additional degree beyond the a-row (at a top) or the degree-zero row. For any $l \le r' + e$, one has the elementary bound

$$J(l) \le J(r') + e\big(\log(r'+e+1) - \log(\Omega/g)\big).$$

Indeed $l! \le (r'+e)!$ and $(r'+e)! \le r'!(r'+e+1)^e$; the exponential scale $\Omega/g \le 1$ on the data. This gives `Data.jetSize_le_add` in `New/Columns/SizeEstimates.lean`. Using it for the rows of the diagram gives an upper bound with leading contribution $\sum_{\tau\in\mathcal{T}} J(a) + k\sum_{\tau\in\Phi} J(w_0(\tau))$. All lower-order row contributions, conjugate-jet errors and the difference between actual and nominal placement are bounded by a polynomial in the ranks and f times $b\log b$ or by similar smaller expressions. The source's size calculation following `diagramWeight_le_by_index` checks the corresponding inequalities with the fixed scale constants. The other terms in the complex triangle inequality of Lemma 10.3 are tightly bounded: the number of classes at each stage is at most $e^{x^{30}}$ (a uniform bound `Data.card_stdClasses_le_exp`, same file), the sum of absolute values of coefficients of the good datum is at most $e^{x^{60}}$, and the final branching produces the large negative exponent in (10.3), with error at most $x^{70} + 2fb\log N$. Since $\log N \ll m\log b$, the latter error is bounded within the budget. Combining the per-lift factorial estimates with this branching calculation gives the right side of (10.21), plus at most B of error. The exact inequality (10.22), including the conjugation uniformity and class-count normalization, is a formal-source input at the quantitative scale being used.

To see how the level N is chosen, retain the notation ℓ_{CS} for the Chowla–Selberg constant and put

$$\overline{w} = \frac{1}{f}\sum_{\tau} \log w(\tau), \quad G_0 = \log(g\,\mathrm{rd}_E/(2\pi)), \quad M_0 = a(\log a - 1 - \ell_{\mathrm{CS}} - \overline{w}).$$

Thus $\Omega = e^{\ell_{\mathrm{CS}}} 2\pi/\mathrm{rd}_E$, $\mathcal{R} = \sum_{\tau\in\mathcal{T}}(A(a,\tau) - M_0)$ and $J(a) = \log(a!) - a\log(\Omega/g)$. The elementary Stirling upper bound for $a \geq 1$ is

$$\log(a!) \leq a\log a - a + 1 + \tfrac{1}{2}\log a.$$

It follows from the monotonicity of the Stirling remainder with the constant fixed at $a = 1$, or by the standard Stirling bound with positive explicit error. Substituting in $J(a) - M_0$ gives

$$J(a) - M_0 \leq a(\overline{w} + G_0) + 1 + \tfrac{1}{2}\log a.$$

By construction of the level window

$$\log N \geq m(\overline{w} + G_0) - G_0 - \tfrac{15}{2}\log b.$$

Consequently

$$\begin{aligned} h_0 &= H - m\mathcal{R} + \#\mathcal{T}\{m(J(a) - M_0) - a\log N\} \\ &\leq H - m\mathcal{R} + \#\mathcal{T}\left\{m(1 + \tfrac{1}{2}\log a) + a(G_0 + \tfrac{15}{2}\log b)\right\} \\ &\leq H - m\mathcal{R} + B. \end{aligned} \tag{10.23}$$

The last estimate uses $\#\mathcal{T} \leq f$, $G_0 = O(\log b)$ and $m \leq 2\sqrt{x} \ll a$, with the large $\log b$ admissibility margin. This is precisely the calculation in `Data.heightPerRow_le` (`New/Endgame/ShortAndEnough.lean`). The level is chosen to remove the large contribution $m\#\mathcal{T}a(\overline{w} + G_0)$, leaving the prime remainder itself.

10.12 Projection, product formula scaling and Theorem A

The good projection e_0 on the finite level quotient G is normalized by $|G|$. Write $P_s = |G|e_0 sP(z_0\cdot)$ with z_0 from Theorem 10.4. The equality at the level of lifts is translation by finite compact elements followed by the same group-ring projection: the value of the lift of $|G|e_0 sz$ at g' is

$$\sum_{u\in G} (|G|e_0(u))\,(\Theta_\theta z)(g'\widetilde{us}).$$

This is `Data.iota_goodProjTranslate_apply` (`New/Rows/MonomialCoefficientVector.lean`). The coefficient $|G|e_0(u)$ is an algebraic integer: the group-ring central idempotents are projections onto specified irreducible types and their coefficients are characters multiplied by dimensions, divided by $|G|$. Finite group characters take algebraic-integer values. Thus the normalized projection preserves integrality at every finite place, including twists prime to that place. This yields `Data.goodProjTranslate_mem_integralForms`. For the trivial twist its complex norm contributes at most $e^{x^{99}}$ times the previous squared-value norm bound when converted to the squared class-sum norm. Indeed the absolute values of the coefficients of $|G|e_0$ are bounded in terms of the finite group order, and the number of classes is bounded as above. The statement with exact constant is `Data.normSqAt_goodProjTranslate_le`, same file.

Theorem 10.15 (A: short forms)**.** *There is a positive rational q such that, with $Q_s = qP_s$, the everywhere-integral columns $u \in C^{\mathrm{ram}}$ for which*

$$\|\Theta_1 Q_s u\|_\sigma^2 \le \exp(H - m\mathcal{R} + L) \qquad (s \in G,\ \sigma \in \mathrm{Aut}(\mathbb{C}))$$

span C^{ram}. For every finite $J \subset G$ and every finite family y of elements of the images $Q_s(C^{\mathrm{ram}}_{\mathrm{int}})$, $s \in J$, the logarithmic index of y in the twisted integral structures is at most

$$|J|cB.$$

Proof. Take the spanning set of (10.4). Put $h_p = \lfloor \epsilon_p - g_p \rfloor$ and $q = \prod_p p^{h_p}$, with the product finite by support of the loss and gain functions. In particular

$$\log q = \sum_p h_p \log p, qquad |\iota(q)|_p = p^{-h_p}.$$

For each p the integer $-h_p = \lceil g_p - \epsilon_p \rceil$ satisfies the strict hypothesis of Proposition 10.14. Therefore on M'_p the factor q times the restricted lift is integral for every admissible twist prime to p. The upper bound (10.22) with $\eta = B$ now gives

$$\begin{aligned} \log \|\Theta_1 Q_s u\|_\sigma^2 &\le 2\log q + h_0 + 2\sum_p g_p \log p + 2B + x^{99} \\ &\le h_0 + 2\sum_p \epsilon_p \log p + 2B + x^{99} \\ &\le H - m\mathcal{R} + 5B + x^{99} \le H - m\mathcal{R} + L. \end{aligned}$$

The last inequality uses the numerical hypothesis of the data, $L = B \log b$ and x^{99} negligible compared to B. A zero norm already satisfies the desired bound; the logarithmic calculation is applied to nonzero norms.

For the index, fix p and ι. On $M \cap C^{\mathrm{ram}}$ pull back the condition $I(M, M'_p; p^e)$ of Proposition 10.14 by inclusion; thus its determinant bound is unchanged. Under every particular Q_s the restricted submodule has integral lifted rows by the seesaw and the normalized projection. If M'' is the inverse image of the integral lifted rows under Θ_θ, this gives $I(M \cap C^{\mathrm{ram}}, Q_s^{-1}(M''); p^e)$. The last part of Lemma 10.2 therefore shows that any family from the specified union of $|J|$ images has a square coefficient matrix over O_ι turning its lifts into integral forms, with determinant dividing $p^{|J|e}$. In particular, its p-adic absolute value is at least $\exp(-|J|c\epsilon_p \log p)$. The numerical costs are independent of ι and of the twist, though the matrices need not be. This exactly meets the definition of logarithmic index in the row space: there is one function of p bounding the costs for every embedding and every twist prime to p. Finally (10.20) gives an upper bound $|J|cB$ for the sum of the costs over every finite set of primes. This proves the second part and completes the proof from the stated inputs. The source declaration is `Data.short_forms`, in `New/Endgame/ShortAndEnough.lean`. □

We emphasize which part of the mathematical reconstruction is still an input. The theorem's last deduction uses only the explicitly stated local divisibility, conjugate-size and index bounds, and their proofs as far as given above. The deepest local subtheorems are the CM jet normal form and Schur bound of the preceding section, the genuine seesaw with uniformly integral finite data, and the Frobenius-twisted Hecke divisor theorem `TwistingData.hasType_twisted_hecke`, with its root inequality and approximate projector. The presentation above specifies the actual versions needed, including the level and gauge compatibilities. They are formal-source inputs, not consequences merely of the usual theta correspondence or Satake parameter formulas.

Chapter 11

The local Hecke estimate in full detail, and the integral seesaw

The exceptional-prime correction of Chapter 10 is not an estimate for generic eigenvalues of a Hecke algebra. It is an estimate for the elementary divisors of a particular integral operator on the full column lattice, followed by an application to its image in ramified columns. The distinction matters: a small characteristic polynomial, or a bound for a normalized spherical function, need not give the required estimate on an integral image. In this chapter we reconstruct the main finite-dimensional theorem of `New/TwistedHecke/ElementaryDivisors.lean`, explaining in particular the integral root, flag-counting and approximate-projector arguments. We finish with some detailed geometric and integral computations needed at the extra level prime. The notation in the Hecke part is local: k is the column rank, n_0 the polynomial degree in each letter (equal to a in the application), r_0 the size of the first block, and $\nu = 2\lambda$ the length of the local Frobenius orbit. These should not be confused with the dimension $n = \dim W$, the global parameter r, or the auxiliary prime N.

11.1 The cyclic coefficient representation and the two filtrations

Let O be a valuation domain of characteristic zero, p a rational prime which is not a unit in O, and fix $w \in O$ with $w^2 = p$. Let R be a local ring with residue cardinality $q = P^\nu$ for $P = p^{d_0} \geq 5$. For each letter $s \in \mathbb{Z}/\nu$ there are ring maps $\iota_s : R \to O$ with

$$\iota_{s+1}(x) \equiv \iota_s(x)^P \pmod{p}.$$

Let a choice of alternating parity on $\mathbb{Z}/\nu$ be fixed. For $g \in \mathrm{GL}_k(R)$ define $g_s = \iota_s(g)$ in an even letter and $g_s = \iota_s(g^{-t})$ in an odd letter. We use the action ρ on the polynomial module generated by the monomials of degree n_0 in each block $X_{b,s}$, given by $gX_{b,s} = \sum_c (g_s)_{cb} X_{c,s}$. Its dimension is

$$M_0 = \binom{n_0 + k - 1}{k - 1}^\nu.$$

For a boundary between coordinates $r_0 - 1$ and r_0, let $\delta(b, s)$ be one on the deep variables (even s and $b \geq r_0$, or odd s and $b < r_0$), and zero otherwise; put

$$d(\alpha) = \sum_{b,s} \delta(b, s)\alpha_{b,s}, \qquad D_w = \mathrm{diag}(w^{d(\alpha)}).$$

The two-block Levi L consists of invertible block-diagonal matrices; let $\mathcal{P}$ be the inverse image of the lower block-triangular parabolic over the residue field. The formula for the group is precisely the convention used in `TwistingData.p_dvd_letterMatrix_of_mem_parahoric`, where an entry with block-colour $i < j$ is zero mod the maximal ideal (even letter).

For distinct indices b, b' write $y_s = X_{b',s}$, $z_s = X_{b,s}$ in the even letters, and $y_s = X_{b,s}$, $z_s = X_{b',s}$ in the odd letters. For a letter i define

$$E_i(\alpha; b, b') = \frac{1}{q}\left(\sum_s P^{1+[s-1-i]_\nu}\alpha(y_s) - \sum_s P^{[s-i]_\nu}\alpha(z_s)\right), \tag{11.1}$$

where $[s-i]_\nu$ is the residue in $[0, \nu-1]$. Write $E_i(v)$ for the weight of a variable v, so that $E_i(\alpha) = \sum_v E_i(v)\alpha(v)$. The exceptional summand at $s = i$ in the first sum has coefficient q, not 1. All weights $E_i(v)$ lie in $[-1/P, 1]$. This follows because the second exponent is at most $\nu - 1$ and the first is at most ν. If O is replaced by its residue field K_0, put U_h equal to the subspace whose every Levi translate has zero coefficients at monomials of excess above h for *some* i, using the boundary $b = r_0 - 1, b' = r_0$. Likewise V_j consists of vectors whose every full-group translate has zero coefficients at depth less than j. These definitions correspond to `excessCond` and `deepForms` of `TwistingData.lean`. We also use U_h mod w^e for the integral congruence condition in which forbidden coefficients must belong to (w^e) instead of being zero.

We will repeatedly use a Gauss filtration. If a variable v has real weight u_v, say a coefficient polynomial F is in filtration at least A if its coefficient of exponent α belongs to $(w^{\max(0,\lceil A - u\cdot\alpha\rceil)})$. This definition makes sense over an arbitrary commutative ring. Its groups are nested, and the product of elements of levels A, A' has level $A + A'$ because the product of their coefficient divisors is sufficient for the sum inequality. Two successive Gauss filtrations (first for a new variable T and then for the old variables) amount to using the additive array of weights of all variables. If two substitutions on polynomial variables each give weight at least u_v, and their difference on v has weight at least $u_v + \eta$, with $\eta > 0$, they induce filtration-preserving maps and their difference raises filtration by η. Indeed the difference on any monomial is a sum obtained by changing one factor at a time, and each summand contains a changed factor of the extra weight. Linearity gives the general statement. This proves the substitution lemma `aeval_sub_mem_of_gaussFiltration` used repeatedly in `RootLemma.lean`.

11.2 The root estimate, including its coefficient extraction

The short argument using multiplicity of zeros in Chapter 10 can be made completely explicit. It is crucial to use a sufficiently large weight separation, rather than directly asking a low-degree polynomial over the residue field to detect all the twists. Here is the complete substitution calculation in `TwistingData.root_lemma`.

Theorem 11.1 (Frobenius root lemma)**.** *Fix $b \neq b'$ in the same block of any Levi, a real weight c_v on all the variables and a coefficient $F_\alpha \in O$. Suppose that for every s*

$$\frac{2}{5} \leq c(y_s) - c(z_s) \leq \frac{8}{5}.$$

Suppose $l \geq 0$, $F_\alpha \notin (w^{l+1})$, and every translate of F by this Levi is in Gauss filtration of level $A = l + c\cdot\alpha$. Then for every letter i, $E_i(\alpha; b, b') \leq 0$.

Proof. Fix i and renumber cyclically so that its relative index is zero; set $t_s = [s-i]_\nu$. Thus $t_i = 0$. Write

$$\omega = \frac{c(y_i) - c(z_i)}{q}, \quad m(z_s) = qP^{t_s}, \quad e(y_s) = P^{1+[s-1-i]_\nu},$$

with $m = e = 0$ off their indicated variables. In particular $e(y_i) = q$. Let π collapse the pair y_i, z_i to z_i and leave every other variable fixed. Define a substitution χ with new variable T by

$$\chi(v) = \pi(v) T^{m(v)}.$$

For $s \neq i$ put $\epsilon_s = -1$ if the parity of s equals that of i, and $\epsilon_s = 1$ otherwise. Put $\epsilon_i = -1$. Define a second substitution ψ by

$$\begin{aligned}
&\psi(z_s) = z_s T^{qP^{t_s}} - \epsilon_s y_s T^{P^{t_s}}, && s \neq i,\\
&\psi(y_s) = y_s - \epsilon_s z_s, && s \neq i,\\
&\psi(z_i) = z_i T^q, && \psi(y_i) = z_i,
\end{aligned}$$

and $\psi = \chi$ on all other variables.

First give T weight $-\omega$ and give the old variables weight $c(v) + \omega m(v)$. The weight of $\chi(v)$ is exactly $c(v)$, since the two weights at the collapsed pair coincide after the adjustment. Each correction term in $\psi(v) - \chi(v)$ has weight at least $c(v) + 1/20$: for $v = z_s$, $s \neq i$, the gain is

$$c(y_s) - c(z_s) - \omega P^{t_s} \geq \frac{2}{5} - \frac{8}{5P} \geq \frac{2}{25} > \frac{1}{20},$$

and for $v = y_s$, $s \neq i$, its gain is

$$c(z_s) + \omega q P^{t_s} - c(y_s) \geq \frac{2P}{5} - \frac{8}{5} \geq \frac{2}{5}.$$

Consequently $\chi(F)$ and $\psi(F)$ have filtration at least A in these weights, and their difference at least $A + 1/20$. Define a scalar-coefficient polynomial $\Phi(T)$ by extracting the coefficient of the old monomial $X^{\pi_* \alpha}$ from $\psi(F)$ and reducing it modulo (w^{l+1}). Put $M = m \cdot \alpha$. Both terms of filtration A at this old exponent have T^j coefficient divisible by w^{l+1} whenever $j > M$, since its required w-order is $l + \omega(j - M) > l$. At $j = M$ the extra filtration of the difference means that its coefficient is also zero mod w^{l+1}. In $\chi(F)$ the old exponents determine all exponents of F except the split between z_i, y_i; then the value of M determines that split uniquely, because changing the exponent of z_i by one changes M by q. The coefficient of T^M is therefore $F_\alpha \bmod w^{l+1} \neq 0$. We have proved

$$\Phi \neq 0, \qquad \deg \Phi \leq M. \tag{11.2}$$

Next give T weight $+\omega$ and the old variables weight $c(v) - \omega e(v)$. Again the adjusted weights of the two collapsed variables agree. For $x \in R$ use the Levi element which first makes the transvection on the root corresponding to (b, b'), with parameter chosen as $-x$ when i is even and x when i is odd, and then interchanges those two basis vectors. For this element g_x the substitution in the sth letter sends $z_s \mapsto y_s$ and $y_s \mapsto z_s + \epsilon_s \iota_s(x) y_s$. This follows by writing the two-by-two transvection matrix: on odd letters its inverse transpose reverses the sign. Compose its action with ψ and the translation $T \mapsto T + \iota_i(x)$.

We verify filtration preservation with respect to these second weights. At i the image of z_i is z_i, and that of y_i is

$$z_i((T + \iota_i(x))^q - \iota_i(x)).$$

Modulo p the expression in parentheses is T^q, by the Frobenius congruence iterated around the full cycle. The main term has weight $c(y_i)$, and the error (coefficients divisible by $p = w^2$)

has weight at least $c(z_i)+2 \geq c(y_i)$. For $s \neq i$, the image of z_s is $y_s - \epsilon_s z_s$ and each term has weight at least $c(z_s)$, by the first inequality in the preceding calculation. The image of y_s is congruent modulo p to

$$z_s T^{qP^{t_s}} - \epsilon_s y_s T^{P^{t_s}}.$$

Indeed $(T+\iota_i(x))^{P^{t_s}} \equiv T^{P^{t_s}} + \iota_s(x)$ and $(T+\iota_i(x))^{qP^{t_s}} \equiv T^{qP^{t_s}} + \iota_s(x)$ modulo p. The constant terms cancel against the terms from the transvection. The first displayed main term has weight $c(z_s)+\omega q P^{t_s} \geq c(y_s)$, the second has weight exactly $c(y_s)$. The error has weight at least the smaller of $c(z_s)+2$ and $c(y_s)-\omega P^{t_s}+2$, so it too has weight at least $c(y_s)$. All other variables are unchanged. This proves the required filtration preservation without any assumption that the valuations are discrete.

For the desired Taylor expansion of $\psi(F)$, apply the proved filtration-preserving composite $(T \mapsto T+\iota_i(x)) \circ \psi \circ g_x$ to $g_x^{-1}F$, which satisfies level A by hypothesis. Its resulting polynomial is exactly the translate $\psi(F)(T+\iota_i(x))$. At the selected old monomial, the coefficient of T^j now has w-order at least

$$l + \omega(e \cdot \alpha - j).$$

For $j < e \cdot \alpha$ this is strictly greater than l. Thus the polynomial Φ of (11.2) has Taylor coefficients zero in degrees below $E = e \cdot \alpha$ at $\iota_i(x)$ for every $x \in R$.

Choose lifts of all q residue classes of R. Their pairwise differences are units in R, so their images remain units in $O/(w^{l+1})$. The polynomials $(T-\iota_i(x))^E$ are pairwise comaximal, each divides the nonzero polynomial Φ, and their monic product has degree qE. Hence $qE \leq M$ (multiplication by a monic polynomial adds its degree even over a ring with zero divisors). Substitution of the definitions of m, e gives exactly (11.1) with nonpositive sign. This proves the theorem. □

The proof only uses ring identities, the specified congruences, and quotient rings by powers of w. In particular it also applies after reduction to a field of characteristic p, with $w = 0$, $l = 0$ and the convention $w^0 = 1$. The proof illustrates why the small adic precision $p = w^2$ is sufficient: its weight 2 is strictly larger than every possible loss in the variable weights. The root lemma is a simultaneous condition on all possible least monomials in all frames. It does not say that all coefficients of an arbitrary polynomial already obey the excess inequalities.

Example 11.2. Take two letters, one even and one odd, and the boundary between coordinates 0 and 1. The two inequalities at a least monomial include

$$q\alpha(y_0) + P\alpha(y_1) \leq \alpha(z_0) + P\alpha(z_1), \qquad q\alpha(y_1) + P\alpha(y_0) \leq \alpha(z_1) + P\alpha(z_0).$$

Even if $\alpha(y_0)$ is small relative to the total degree in the even letter, it may violate the first inequality. The cyclic coefficients are a reflection of the two different embeddings of the same residue field: treating the letters as independent finite fields would lose precisely these constraints.

11.3 The dimension bound from flags and convex volume

We now calculate the dimension of U_h over a field K_0 of characteristic p with the indicated reduction of the twisting data. Because the critical coefficient is asymptotic to the reciprocal of the degree of the local Hecke correspondence, it is worth keeping the leading volume constant.

For $h \geq 0$, let $\mathscr{C}_h$ be the cone in the exponent space $\mathbb{R}^{k\nu}$ defined by the excess inequalities for adjacent coordinates:

$$E_i(\beta; b, b+1) \leq 0 \ (b+1 \neq r_0), \qquad E_i(\beta; r_0-1, r_0) \leq h$$

for every letter i. Its intersection with the nonnegative orthant and with the sets where the coordinate sums are in bounded intervals is a compact convex polytope. Write $\Sigma(D)$ for the product of slabs $\beta_{bs} \geq 0$, $D_s - 1 \leq \sum_b \beta_{bs} \leq D_s$. Let I denote the Iwahori subgroup (lower triangular reduction) of $\mathrm{GL}_k(R)$. Write $F_L = [L : L \cap I]$ for the number of full flags in the two Levi blocks.

Lemma 11.3 (Amplification by a full Frobenius period)**.**

$$\dim U_h \leq F_L \operatorname{vol}(\mathscr{C}_{h+\nu(k+1)} \cap \Sigma(n_0 + k + 1)).$$

Here the endpoint in the slab is constant in every letter and the measure is ordinary coordinate volume.

Proof. Choose representatives g_j for $(L \cap I)\backslash L$ including the identity. Take $M = q^u$ with u any positive integer. For each array $v = (v_{bs})$ with $0 \leq v_{bs} < M$, define $\gamma(v)$ by adding

$$M\left(k - \left\lfloor \frac{\sum_b v_{bs}}{M} \right\rfloor\right)$$

to the coordinate $b = r_0 - 1$ in each letter s and leaving the other coordinates unchanged. The exponent $\gamma(v)$ is congruent to v modulo M coordinatewise, and its degree in each letter lies in $[kM, (k+1)M)$. Make the following map on $M^{k\nu}$ copies of U_h:

$$(F_v)_v \longmapsto F = \sum_v X^{\gamma(v)} F_v(X^M).$$

It is injective, since exponents for different v are distinct modulo M, while for a fixed v distinct monomials remain distinct. Raising coefficients in any letter action matrix to the Mth power leaves them unchanged: M is a power of the cardinality of the original residue field and the reduced actions have entries in its image. By the Frobenius identity, linear substitution in $F_v(X^M)$ therefore commutes with the letter action. This is the necessity of choosing a *full* Frobenius period, not just $M = p^u$ of arbitrary exponent.

Suppose $F \neq 0$. Among all nonzero terms of all $g_j F$, choose one exponent α with least weight relative to the array $c(X_{b,s}) = b$ in an even letter and $-b$ in an odd letter. Every translate by an element of L has all terms of weight at least $c \cdot \alpha$: such a translate is hg_j with $h \in L \cap I$. The action of h cannot decrease this weight. Indeed in an even letter all nonzero matrix entries in the residue field have row index at least column index, and in an odd letter the reverse inequality holds after inverse transpose. Thus the chosen coefficient meets the hypotheses of the root lemma over K_0, with $w = 0, l = 0$ and with the Levi L for any adjacent pair lying in the same block. For each such pair the weight separation is exactly one. All its excess inequalities are consequently nonpositive.

For the boundary pair, any nonzero coefficient of any $g_j F$ comes from an exponent of the form $M\beta + \eta$, where β is in the support of a Levi translate of some F_v, and η has degree less than $(k+1)M$ per letter. This follows by expanding each summand of the substituted polynomial; nonzero coefficient means at least one such term contributes. Thus $E_i(\beta) \leq h$ and $E_i(\eta) \leq \nu(k+1)M$ because all variable excess weights are at most one. Hence $E_i(\alpha/M) \leq h + \nu(k+1)$ at the boundary. The degree of α/M is in $[n_0 + k, n_0 + k + 1]$ in each letter. It follows that it belongs to the polytope in the lemma.

Therefore the coefficient restrictions at the lattice points α/M of this polytope of all the finitely many $g_j F$ give an injective linear map. If their values all vanished, the above least nonzero coefficient would yield a contradiction. Its target has dimension at most F_L times the number of lattice points in the polytope of mesh $1/M$. Comparing dimensions gives

$$M^{k\nu} \dim U_h \leq F_L \#((M^{-1}\mathbb{Z}^{k\nu}) \cap \mathscr{C}_{h+\nu(k+1)} \cap \Sigma(n_0 + k + 1)).$$

Divide by $M^{k\nu}$ and let $u \to \infty$. The lattice-point Riemann sums converge to volume: the set is bounded and its boundary is contained in finitely many affine hyperplanes, unless it lies in such a hyperplane, in which case both the limiting sums and the volume are zero. This proves the bound. □

This is the proof of `TwistingData.finrank_excessCond_le_volume`, `New/TwistedHecke/Counting.lean`. It is a pure dimension argument and does not require the irreducibility of the symmetric-power representations.

The exact volume in the basic cone can also be calculated without an asymptotic estimate. In the following calculation we use the sign $\epsilon(s) = 1$ on even letters, -1 otherwise. Introduce an invertible linear change

$$\beta_{b,s} = \epsilon(s)(Pv_{b,s+1} - v_{b,s}).$$

The determinant is nonzero because in the kernel the equations cycle to $(P^\nu - 1)v_{b,s} = 0$. A direct telescoping gives

$$E_i(\beta; b, b') = (1 - q^{-1})(\beta(y_i) - (v_{b,i} - v_{b',i})). \tag{11.3}$$

For clarity, set $i = 0$. The excess can be written as $(1-q^{-1})\beta(y_i)+q^{-1}\sum_s P^{[s-i]_\nu}(\beta(y_s)-\beta(z_s))$. Now $\beta(y_s) - \beta(z_s) = P(v_{b',s+1} - v_{b,s+1}) - (v_{b',s} - v_{b,s})$, in either parity. The cyclic sum is $(q-1)(v_{b',i} - v_{b,i})$. This proves the identity. Nonnegative β satisfying $\mathscr{C}_0$ correspond to

$$\begin{cases} v_{b,s} \le Pv_{b,s+1}, & Pv_{b+1,s+1} \le v_{b,s}, \quad s \text{ even}, \\ Pv_{b,s+1} \le v_{b,s}, & v_{b+1,s} \le Pv_{b,s+1}, \quad s \text{ odd}. \end{cases}$$

Away from the equality hyperplanes this forces all vectors v_s to be decreasing in b. In particular the signs of v_s do not enter the counting, since only the adjacent comparisons do.

Lemma 11.4 (Cyclic interlacing volume)**.** *For all endpoints D_s,*

$$\operatorname{vol}(\mathscr{C}_0 \cap \Sigma(D)) = \frac{\operatorname{vol}\Sigma(D)}{\prod_{j=1}^{k}(q^j - 1)/(q-1)}.$$

Proof. Ignore hyperplanes of equalities (measure zero). For decreasing vectors x, y let $C(x, y)$ be the matrix whose (b, c) entry is $\mathbf{1}(y_c < x_b)$. Its rows have nested supports in the final segments of the coordinate set. Its determinant is one for strict interlacing $x_1 > y_1 > x_2 > \cdots > x_k > y_k$, and zero otherwise. Indeed all rows must be distinct and nonzero, so their lengths must be exactly $k, k-1, \ldots, 1$. This gives the claimed determinant and condition.

Apply this identity to (Pv_{s+1}, v_s) in an even letter and (v_s, Pv_{s+1}) in an odd letter, transposing the former matrix if needed to give indices at s and $s+1$ consistently. The product of these determinants on sorted vectors is the indicator of the interlacing cone with nonnegative β. It is invariant under permutation of the coordinates at any single letter, because that permutation affects two consecutive determinants and their signs cancel. We may integrate over all vectors, subject only to the condition on the coordinate sums, and divide by $(k!)^\nu$. On expanding the determinants, each permutation tuple $(\sigma_s)_s$ gives the signed indicator of positivity in the linear transformation

$$(L_\sigma v)_{b,s} = \epsilon(s)(Pv_{b,s+1} - v_{\sigma_s(b),s}).$$

The sums of coordinates of $L_\sigma v$ in a letter are independent of σ, so the integral of the indicator in the sum equals $\operatorname{vol}\Sigma(D)/|\det L_\sigma|$. To calculate the determinant, put A_s for the

permutation matrix associated to σ_s, in the convention $A_s v_s = (v_{\sigma_s(b),s})_b$. Ignore the sign factors, which are units. The equations give $v_{s+1} = P^{-1}(A_s v_s + \beta_s)$; successive elimination of $v_1, \dots, v_{\nu-1}$ contributes powers of P whose product is $P^{k(\nu-1)}$. The remaining block for v_0 is $PI - P^{-(\nu-1)} A_{\nu-1} \cdots A_0$; multiplying its determinant by those factors gives precisely $|\det(qI - A_\pi)|$, where π is the product permutation. Replacing the v-coordinates by the original β-coordinates multiplies the result by $|\det L_1| = (q-1)^k$. For any fixed π there are $(k!)^{\nu-1}$ permutation tuples of product π and the product of their signs equals $\operatorname{sgn}\pi$. The volume ratio is thus

$$\frac{(q-1)^k}{k!} \sum_{\pi \in S_k} \frac{\operatorname{sgn}\pi}{|\det(qI - A_\pi)|}.$$

This expression depends only on q and k, not on the number of letters or their parities. Evaluate it by the same one-letter determinant formula with parameter q (the all-positive parity in that one letter is legitimate for this measure calculation). Then the cone is $\beta_b \geq q\beta_{b+1}$, $\beta_k = 0$ for numbering $b = 0, \dots, k-1$. Let $d_b = (q^{b+1} - 1)/(q-1)$. The triangular linear map $\beta \mapsto (d_b(\beta_b - q\beta_{b+1}))_b$ carries that cone onto the nonnegative orthant, preserves coordinate sum, and has determinant $\prod_b d_b$. Hence the volume ratio is its reciprocal, giving the formula. □

This proves the determinant-volume identities in `New/TwistedHecke/InterlacingCone.lean`. For example, if $k = 2$ the volume fraction $1/(q+1)$ is exactly the reciprocal of the number of lines over $\mathbb{F}_q$. For higher k the denominator is the number of complete flags, obtained recursively by choosing a line in the first quotient and then a complete flag of the quotient space.

We now need the thickened cone with nonzero allowance at the boundary. Define $u_b = 2P^{2(r_0-1-b)}$ if $b < r_0$, and $u_b = 0$ otherwise. Use $v_{b,s} = u_b$ on even letters and $v_{b,s} = Pu_b$ on odd letters. Then $(L_1 v)_{b,s}$ equals $(P^2 - 1)u_b$ in even letters and zero in odd letters. Call this vector σ. It has nonnegative coordinates bounded by P^{3r_0} (indeed $2P^{2r_0} \leq P^{3r_0}$). Since

$$u_b = P^2 u_{b+1} + 2\mathbf{1}(b+1 = r_0),$$

substitution in (11.3) shows $\sigma \in \mathscr{C}_{-1}$. At an even letter the inner expression is $P^2 u_{b+1} - u_b$, giving -2 on the boundary, zero elsewhere. At an odd letter it is $-P(u_b - u_{b+1})$, which is nonpositive and at most $-2P$ on the boundary. The factor $1 - q^{-1}$ leaves the boundary bound at most -1. This is the explicit vector of `exists_mem_excessCone_mem_Icc`.

Put $h' = h + \nu(k+1)$. Translating every exponent in the polytope of Lemma 11.3 by $h'\sigma$ embeds it in the zero-excess cone, while the endpoint of the slab in letter s increases by at most $h'kP^{3r_0}$. Suppose $\varepsilon \leq 1$ and the size assumptions (10.13) hold (with $a = n_0, y = h$). Then the new endpoints are at most

$$n_0 + k + 1 + h'kP^{3r_0} \leq n_0(1 + 1/h).$$

The volume of a nonnegative slab with endpoint D in dimension k is $(D^k - (\max(D - 1, 0))^k)/k! \leq D^{k-1}/(k-1)!$. Together the shifted slabs have volume at most

$$\left(\frac{n_0^{k-1}}{(k-1)!}\right)^{\nu} e^{\nu(k-1)/h} \leq (1+\varepsilon)M_0,$$

since $n_0^{k-1}/(k-1)! \leq \binom{n_0+k-1}{k-1}$, $\nu(k-1)/h \leq \varepsilon/2$, and $e^{v/2} \leq 1 + v$ for $0 \leq v \leq 1$. Finally $[\mathrm{GL}_k(R) : \mathcal{P}]F_L$ is the number of complete flags: reduction identifies the cosets of $\mathcal{P}$ with subspaces of dimension the appropriate block size, and a full flag refining the fixed block is a pair of flags inside the two blocks (the unipotent part lies in I). Lemmas 11.3 and 11.4 therefore give the needed bound

$$[\mathrm{GL}_k(R) : \mathcal{P}] \dim U_h \leq (1+\varepsilon)M_0. \tag{11.4}$$

This is the quantitative count `index_mul_finrank_excessCond_le` with exactly the needed leading constant.

11.4 Deep coefficients: a weighted Nullstellensatz and its application

The volume argument controls the dimension of U_h. The other loss in the Hecke theorem is the cumulative dimension of the depth spaces V_j. To obtain it one needs a genuinely different argument: the obstruction is high-order vanishing at a large family of twisted unipotent points. We prove the multiplicity-weighted version of the Nullstellensatz first.

Lemma 11.5 (Weighted vanishing)**.** *Let K be a field, $u_1, \dots, u_l$ independent variables, and $\kappa_1, \dots, \kappa_l$ positive integers. Give u_i cost $1/\kappa_i$. Let $Z \subset K^l$ be a finite set of at least $\prod_i \kappa_i$ distinct points. For every i suppose there is a polynomial g_i of cost strictly below one agreeing with $u_i^{\kappa_i}$ at all points of Z. If a nonzero polynomial F vanishes to order at least d at all points of Z, then some monomial of F has cost at least d.*

Proof. The assertion is clear for $d = 0$. Set $h_i = u_i^{\kappa_i} - g_i$, which vanish on Z. The polynomials

$$u^\alpha \prod_i h_i^{\beta_i}, \qquad 0 \le \alpha_i < \kappa_i,$$

have distinct initial monomials $u^{\alpha+\kappa\beta}$ of cost $\sum \alpha_i/\kappa_i + |\beta|$; all remaining terms have lower cost. They form a basis of the polynomial ring over K: independence follows by taking largest cost (and distinct initial terms at that cost), and spanning follows by subtracting the appropriate basis elements from the terms of maximal cost. This terminates because the cost is discrete, being a rational number with bounded denominator. In particular the polynomials of cost below d lie in the span of basis elements with $|\beta| < d$.

Consider the quotient of the polynomial ring by the intersection of the dth powers of the maximal ideals of the points in Z. These ideals are pairwise comaximal, so by the Chinese remainder theorem the quotient has dimension $|Z|\binom{d+l-1}{l}$. Since all basis elements with $|\beta| \ge d$ vanish to the indicated order, the quotient is spanned by the images of the $(\prod_i \kappa_i)\binom{d+l-1}{l}$ remaining basis elements. The cardinality hypothesis implies equality and makes their images independent. A nonzero polynomial of cost below d cannot belong to the intersection. This is the desired statement. □

The argument is the ordinary division proof of `exists_le_cost_of_vanishing`, in `Lib/General/Algebra/ZerosCostDegree.lean`. It also shows why merely using the degree of a polynomial on a finite grid would not suffice: certain variables have extremely different costs, and the grid is not a product in those coordinates.

Proposition 11.6 (Everywhere deep forms have a costly coefficient)**.** *Suppose $2r_0 \le k$ and $2r_0 < \nu$. Give each deep variable of index $b < 2r_0$ cost P^{-1}, each deep variable of index $b \ge 2r_0$ cost P^{-2r_0}, and all other variables cost zero. Every nonzero member of V_j has a nonzero monomial of cost at least j.*

Proof. Put $k = 2r_0 + h_0$. We use upper unipotent matrices $1 + A$ with entries A_{ab} only in the following positions: the r_0 paired positions $(a, r_0 + a)$, $0 \le a < r_0$, and all $r_0 h_0$ positions (a, b) with $b \ge 2r_0$. In particular $A^2 = 0$, and $1 + A$ is invertible. Index the positions by (a, c), $0 \le a < r_0$, $0 \le c \le h_0$, where the second coordinate of the matrix entry is $r_0 + a$ for $c = 0$ and $k - c$ otherwise. Denote the resulting pair of matrix indices by $\rho(t), \sigma(t)$ for a position t.

Take each entry $A_{\rho(t),\sigma(t)}$ to be a lift of a residue parameter $u_t \in \mathbb{F}_q$ and write $a_s(t) = \bar{\iota}_s(u_t)$ for its residue in K_0. The Frobenius congruences give $a_{s+e}(t) = a_s(t)^{P^e}$.

Let K be the field of rational functions in the shallow variables over K_0 (one may harmlessly adjoin independent indeterminates also for the original deep variables). In the new polynomial algebra over K, keep only the deep variables as variables and regard the others as scalar coefficients. Define a point $z(u)$ by

$$z(u)_{b,s} = \sum_{\sigma(t)=b} a_s(t) X_{\rho(t),s}, \qquad s \text{ even},\ b \geq r_0,$$

$$z(u)_{b,s} = -\sum_{\rho(t)=b} a_s(t) X_{\sigma(t),s}, \qquad s \text{ odd},\ b < r_0.$$

Thus $\rho(1-A)$ sends a deep variable to the corresponding variable minus its coordinate of $z(u)$, and fixes shallow variables. This follows directly from g_s being the ordinary matrix in even letters and the inverse transpose in odd ones. Since $\rho(1+A)F$ has depth at least j, writing $F = \rho(1-A)\rho(1+A)F$ shows that $F \in \mathfrak{m}^j_{z(u)}$ for every u.

We claim that the coordinates satisfy the hypotheses of Lemma 11.5 with $\kappa_{b,s} = P$ for the deep variables $b < 2r_0$ and $\kappa_{b,s} = P^{2r_0}$ for the other deep variables. To prove this, it suffices to recover every $a_i(t)$ as the value at $z(u)$ of a polynomial of cost less than one, uniformly in u. For fixed i and each paired position $(a, 0)$, take the coordinate of z in letter i corresponding to the deep root variable of $(a, r_0 + a)$, with exponent one. For the other position (a, c) choose $j_t \in \{2a, 2a+1\}$ such that the letter $i - j_t$ is even, and use the coordinate $z^{P^{j_t}}_{k-c,i-j_t}$. Its cost is $P^{j_t - 2r_0} < 1$, since $j_t \leq 2r_0 - 1 < \nu$. The paired coordinates also cost less than one.

At the points, these expressions are linear combinations of the $a_i(t')$ by the Frobenius identities. Their coefficient matrix over the rational function field is invertible. Here is a direct check: specialize each shallow variable $X_{b,s}$ to one if $b \geq \lfloor [i-s]_\nu/2 \rfloor$, and to zero otherwise. For a row indexed by (a, c) with $c > 0$ its only nonzero coefficients correspond to (a', c) with $a' \geq a$, since $\lfloor j_t/2 \rfloor = a$; its diagonal coefficient is one. For a row $(a, 0)$ in an even i only the corresponding paired matrix entry contributes, with coefficient one. For odd i only entries of row a contribute (the paired position and the far columns), with sign minus one, and its diagonal coefficient is minus one. In the lexicographic order on (a, c) the specialized matrix is upper triangular with unit diagonal (up to signs), and hence its determinant is nonzero. This proves the invertibility claim before specialization.

The claim now follows: after raising any coordinate of $z(u)$ to the indicated $\kappa_{b,s}$ and applying Frobenius, the result is a linear combination of residue parameters $a_{s+1}(t)$ or $a_{s+2r_0}(t)$ over the rational function field. Replace each by its polynomial expression of cost below one. The points are distinct, because recovering $a_i(t)$ from the coordinates recovers u_t (the map on residue fields is injective). Their cardinality is $q^{r_0(h_0+1)}$. The product of the costs' denominators is exactly this: at each letter there are r_0 deep variables among the first $2r_0$ coordinates, giving $P^{\nu r_0}$; the h_0 far coordinates are deep in half the letters, giving $P^{2r_0 h_0 \nu/2}$. This proves the required cardinality. Lemma 11.5 applies to the polynomial in the deep variables over the rational function field. Its high-cost coefficient must come from at least one monomial of the original F with the same deep exponents. The shallow variables have zero cost. This proves the proposition. □

This reconstructs `exists_le_weight_of_mem_deepForms` in `New/TwistedHecke/DeepFormCost.lean`, including the recovery matrix in the proof. The triangular specialization is essential: there cannot be a separate parameter for every deep variable because that would eliminate the gain P^{-2r_0} in the estimate.

The restriction to costly coefficients is injective on V_j: a nonzero vector in its kernel would contradict Proposition 11.6. Each monomial is counted at most 1 plus its cost when one sums

these dimensions over $j \ge 0$. The mean exponent of any specified variable among degree-n_0 monomials in a letter is n_0/k, by coordinate symmetry. Hence the mean cost is

$$\frac{n_0\nu}{2k}\left(\frac{2r_0}{P}+\frac{k-2r_0}{P^{2r_0}}\right) \le \nu n_0\left(P^{-2r_0}+\frac{r_0^2}{kP}\right) \le z$$

under (10.13). The same size assumptions imply $2r_0 < \nu$: indeed $P^{3r_0} \le n_0 < P^{\lambda}$, hence $3r_0 < \lambda$. We obtain, for every J,

$$\sum_{0\le j<J} \dim V_j \le (1+z)M_0. \tag{11.5}$$

Thus both dimension bounds needed in the Hecke matrix theorem are consequences of explicitly proved algebraic estimates. We have not appealed to a semisimplicity statement in characteristic p.

11.5 Integral diagonal pairings and class means

Although the dimension bounds were proved in characteristic p, our application concerns an integral matrix. It is not enough to assert that a desired projector exists after reduction mod p. We need an integral approximation to it, with an explicit scalar normalization. The source constructs this approximation from invariant linear functionals on polynomials in a generic matrix. We spell out their construction.

For a finite set of coordinates S and monomial exponent α of total degree $a_0 < p^F$, the multinomial $m_\alpha = a_0!/\prod_i \alpha_i!$ divides $p^{|S|F}$ in O. Indeed the part of the integer prime to p is a unit in the local valuation ring. Legendre's formula expresses the exponent of p in m_α as

$$\sum_{j\ge 1}\left(\left\lfloor\frac{a_0}{p^j}\right\rfloor-\sum_i\left\lfloor\frac{\alpha_i}{p^j}\right\rfloor\right).$$

Terms for $j \ge F$ vanish, and each other term is at most $|S|-1$, by the floor inequality. Therefore $q_\alpha = p^{|S|F}/m_\alpha$ belongs to O. These numbers give the reciprocal-multinomial diagonal inner product for the symmetric power action. More precisely for any $S \times S$ matrix G over O, let L_G act on polynomials by substituting $X_b \mapsto \sum_c G_{cb}X_c$. The Fischer pairing is diagonal with $\prod_i \alpha_i!$ on X^α. Expanding powers of linear forms (or checking the identity by the binomial formula) gives

$$\langle L_G X^\beta, X^\alpha\rangle = \langle X^\beta, L_{G^t}X^\alpha\rangle.$$

Multiplying by $p^{|S|F}/a_0!$ in the fraction field and using the integral values yields

$$q_\alpha[X^\alpha]L_G(X^\beta) = q_\beta[X^\beta]L_{G^t}(X^\alpha). \tag{11.6}$$

This is `ClassMean.exists_invariant_diagonal_form`. In particular it holds integrally for noninvertible G too.

We also need the elementary interpolation span over the residue field. Let d_1 be a homogeneous degree and Δ a nonzero homogeneous polynomial of degree d_2 over R with a unit value. Suppose $d_1 + d_2 < q$. Then the evaluations on points where Δ is a unit span the full dual of the homogeneous polynomial module of degree d_1 over O. Indeed a vector in the kernel of all those evaluations over the residue field gives a polynomial F for which $F\bar{\Delta}$ vanishes at every $\mathbb{F}_q$-point. Applying the one-variable root bound in each variable successively makes this polynomial zero, because its degree in each variable is below q. As $\bar{\Delta}$ is not

identically zero, $F = 0$. This argument also works over a residue coefficient extension, by the same root bound with points in the embedded residue subfield. Therefore a finite subset of the evaluation vectors has invertible determinant mod the maximal ideal and lifts to a basis over O. We will only apply this strict inequality case of `ClassMean.eq_top_of_eval_monomial`; the source includes equality in total degree for its homogeneous version.

For the next construction fix the generic $k \times k$ matrix X, a real weight array x_b on its indices, and a subset $T \subseteq \mathbb{R}$ (not necessarily an interval). For a monomial in matrix entries let its row weight be the sum of weights x_b of its row indices counted with multiplicities, and column weight defined similarly. Put $C_g F(X) = F(\iota(g)^{-1} X \iota(g))$ for g in the Levi. We use C_g for this algebra automorphism, not for the column space.

Proposition 11.7 (Weighted conjugation mean)**.** *Suppose $D < p^F$, $D + k < q$. There is an integral linear form Λ_T on the polynomial ring in X with properties:*

1. *It is invariant under C_g for all $g \in L$.*
2. *Its value on monomials of different row and column weights is zero, and its value on a monomial X^α with row weight in T is*
$$\Lambda_T(X^\alpha) = \begin{cases} p^{k^2 F} X^\alpha(1), & |\alpha| \le D, \\ 0, & |\alpha| > D. \end{cases}$$
3. *It belongs to the O-span of forms $\ell \circ C_g$, where ℓ vanishes on monomials whose row weight does not belong to T.*

Proof. The problem splits by homogeneous degree; construct the form in one degree $a_0 \le D$ and extend by zero on higher degrees. Let $U = (O^k)^{\otimes a_0}$ with its tensor basis $e_{t_1,\dots,t_{a_0}}$. Denote by $Y(X)$ the block-diagonal part of X^t (transpose included in this definition). For any homogeneous dual functional ℓ define $S(\ell) : U \to U$ by applying ℓ to each entry of $Y(X)^{\otimes a_0}$. This is an O-linear map in ℓ. The evaluation vectors at block-invertible matrices span the dual by the interpolation argument with polynomial $\Delta = \det(Y(X))$ of degree k, which has unit value at the identity. Therefore the image of S is the O-linear span of the tensor powers of the transposes of matrices in the Levi. Write this image as $\mathscr{A}$; it is an algebra, since the product of two such matrix evaluations is another such evaluation.

Let U_T be the $\mathscr{A}$-submodule generated by the tensor basis vectors e_t for which $\sum_j x_{t_j} \in T$. It is a finitely generated torsion-free O-module, hence finite free. (One proof of the standard valuation-ring fact is pivot elimination on finitely many generators of a submodule of a free module: an entry dividing the other finitely many entries in its row can be split off and induction applies.) The action of $\mathscr{A}$ on U_T consequently has a trace. For each exponent α of degree a_0 let $K(X^\alpha) = q_\alpha[X^\alpha]$ be a functional on the homogeneous polynomials, with q_α as in (11.6) for $|S| = k^2$. Define

$$\Lambda_T(F) = \mathrm{Tr}_{U_T}(S(K(F))|_{U_T}).$$

Equation (11.6), applied to the matrix substitution under conjugation, gives covariance of $S(K(F))$ under the tensor action of a Levi element. With $C_g(F)(X) = F(g^{-1} X g)$ the transpose of the substitution on entry variables is $X \mapsto g^{-t} X g^t$. On the transposed block matrix $Y(X)$ this induces $Y \mapsto g Y g^{-1}$. Thus $S(K(C_g F)) = S(\mathrm{ev}_{g^t}) S(K(F)) S(\mathrm{ev}_{g^{-t}})$, with ι implicit in the matrix entries. Since the Levi is transpose-stable, this makes Λ_T invariant by cyclicity of trace.

Suppose the variable monomial has unequal row and column weights. The matrix $S(K(X^\alpha))$ on U sends a basis vector of weight ω to one of weight $\omega + \delta$, where δ is the column minus row

weight of α. Here $\delta \neq 0$. Since the set of tensor weights is finite, sufficiently large positive powers of this matrix are zero. Its restriction to the stable U_T is nilpotent, hence has zero trace (check over the fraction field). This proves the first part of property 2. If instead the row weight lies in T, the matrix annihilates all e_t not in the indicated generating set and has image in U_T. Its trace on U_T equals its trace on all U; again this may be checked after passing to the fraction field and choosing a complement to U_T. The trace of the universal tensor matrix is $(\sum_b X_{bb})^{a_0}$, so extracting the coefficient of X^α and multiplying by q_α gives $p^{k^2F} X^\alpha(1)$. This gives the remaining equality in property 2.

To prove the span assertion, choose a basis (z_j) of U_T and let λ_j be its dual. Express each z_j as a finite sum of terms $cS(\mathrm{ev}_g)e_t$ with e_t of allowed weight, possible by the definition of U_T and the evaluation-span assertion. The trace is a sum of expressions

$$c\lambda_j(S(K(F))S(\mathrm{ev}_g)e_t).$$

Move $S(\mathrm{ev}_g)$ to the left by the covariance just proved. Each term is now of the form $\ell(C_{g'}F)$ for an integral functional ℓ, where $g' = g$ or g^{-1} as appropriate; ℓ vanishes on forbidden row weights because $S(K(X^\alpha))e_t = 0$ there. This proves property 3. Summing the finitely many homogeneous constructions completes the proof. □

This is the trace construction of `ClassMean.exists_mean`, `New/TwistedHecke/ClassMeans.lean`. The name "mean" need not indicate averaging with an invertible group order. The bounded power of p replaces division by the order of a large finite group, and the trace construction works in the integral setting. Replacing it by the ordinary group average would destroy the denominator estimate.

11.6 Frobenius products of matrices and the approximate projector

Continue with the same local representation; we now take the particular adjacent pair $b = r_0 - 1, b' = r_0$. Write

$$x_b = \frac{\mathbf{1}(b = r_0) - \mathbf{1}(b = r_0 - 1)}{P^\lambda}, \qquad \nu = 2\lambda.$$

Use a generic matrix X_j for each letter $j \in \mathbb{Z}/\nu$. Put $\Psi_{s,d} = X_{s-d}^{(P^d)}$, taking the entrywise power. The product

$$H_{s,m} = \Psi_{s,m-1} \cdots \Psi_{s,0}$$

for $m = \lambda + \delta(b, s)$ supplies row b of a generic coefficient action matrix A_s: row b of $H_{s,m}$ in an even letter, row b of its transposed adjugate in an odd letter. Let

$$\mathscr{F}(\Psi) = \bigotimes_s \mathrm{Sym}^{n_0} A_s.$$

The row-dependent product length is essential. In the formal source it is `TwistingData.frobOperator` in `New/TwistedHecke/FrobeniusFactors.lean`. If all generic matrices are conjugated by the corresponding letter embedding of $g \in L$, the induced operator changes by conjugation by $\rho(g)$ (with the order determined by the substitution on functions). Indeed at fixed s both $H_{s,\lambda}$ and $H_{s,\lambda+1}$ undergo ordinary conjugation. The selection of rows commutes with the block diagonal g_s, because δ is constant on each block; and the identity $\mathrm{adj}(GHG^{-1})^t = G^{-t}\,\mathrm{adj}(H)^t G^t$ gives the inverse-transpose rule in the odd letters. Finally the symmetric-power action respects products. This proves the exact covariance `frobOperator_map`. We distinguish this exact invariance from the approximate compatibility of conjugation with entrywise Frobenius powers.

Lemma 11.8 (Weights of products)**.** *Let $\mathcal{F}^u$ be any decreasing multiplicative filtration on a commutative coefficient ring. Let $\alpha_{s,d}(b)$ be real weights with $\sum_b \alpha_{s,d}(b) = 0$, and $|\alpha_{s,0}(b)| \le u$, $u \ge 0$. Suppose*

$$(\Psi_{s,d})_{bc} \in \mathcal{F}^{\alpha_{s,d+1}(b)-\alpha_{s,d}(c)-u} \qquad (0 \le d \le \lambda).$$

Put $\widetilde{x}(b,s) = \alpha_{s,\lambda+\delta(b,s)}(b)$ in an even letter and its negative in an odd letter. Then every entry in row β of $\mathscr{F}(\Psi)$ has filtration at least

$$\widetilde{x} \cdot \beta - \nu n_0 k(\lambda+2)u. \tag{11.7}$$

Proof. Expanding the product $H_{s,m}$ entry by entry and applying the multiplicative filtration gives $(H_{s,m})_{bc} \in \mathcal{F}^{\alpha_{s,m}(b)-\alpha_{s,0}(c)-mu}$, because the intermediate indices' weights telescope. This is at least $\alpha_{s,m}(b)-(m+1)u$. In an even letter this bounds the row of A_s. In an odd letter each cofactor of the transposed adjugate in row b is a signed sum of products using every row of $H_{s,m}$ except b. The sum-zero assumption on $\alpha_{s,m}$ gives filtration at least $-\alpha_{s,m}(b)-(k-1)(m+1)u$. Both estimates are at least $\widetilde{x}(b,s) - k(\lambda+2)u$. In the symmetric power, a contribution to the coefficient of the output monomial β_s uses exactly $\beta_s(b)$ entries from row b of A_s; the total loss per letter is at most $n_0k(\lambda+2)u$. Taking the product over all letters gives (11.7). □

This is the abstract filtration lemma `TwistingData.frobOperator_mem`. In particular the degrees in the variables of any one generic matrix which occur in $\mathscr{F}(\Psi)$ are at most $D = \nu n_0 k(\lambda+2)P^\lambda$: apply the same argument to the degree filtration of negative weights, or simply note that each symmetric-power term uses at most $n_0k(\lambda+1)$ factors of degrees at most P^λ in a letter, and then sum over the ν letters. This slightly coarse bound is sufficient for the interpolation cutoff in Proposition 11.7.

Now let $g \in L$ and replace the generic matrices used above with

$$\Psi_{s,d} = \iota_s(g)\big(\iota_{s-d}(g^{-1})X_{s-d}\iota_{s-d}(g)\big)^{(P^d)}\iota_s(g^{-1}). \tag{11.8}$$

Modulo p this is exactly $X_{s-d}^{(P^d)}$, by the coefficient Frobenius congruence and the fact that entrywise raising to a power of p is a ring homomorphism over characteristic p. In particular the error terms have coefficients divisible by $p = w^2$ and are homogeneous of degree P^d.

Proposition 11.9 (The weighted coefficient estimate)**.** *Let η be an exponent in the generic matrices, with restriction η_j to X_j. Assume its row weight equals its column weight at every $j \ne i$. Put t_i equal to its row weight at i, with weights x_b above. Assume $\nu n_0 k(\lambda+2) \le P^\lambda$. If the output exponent β satisfies*

$$|E_{i+\lambda}(\beta; r_0-1, r_0) - t_i| > 2c+1,$$

then every coefficient of X^η in row β of $\mathscr{F}(\Psi)$ for (11.8) is divisible by w^c.

Proof. Let $e \in \{-1,1\}$ be the sign of the indicated difference and put

$$A_a(b) = \frac{e}{2}\frac{P^a x_b}{q}, \qquad u = \frac{1}{2P^\lambda}.$$

Give the variable $(X_j)_{bc}$ weight $A_\nu(b)$ if $j = i$, and $A_{[j-i]_\nu}(b) - A_{[j-i]_\nu}(c)$ otherwise. These weights, denoted ω_{jbc}, satisfy $|\omega_{jbc}| \le P^{-\lambda}$. Set $\alpha_{s,d}(b) = A_{d+[s-d-i]_\nu}(b)$. Then $|\alpha_{s,0}(b)| \le u$ and their sum over b is zero. Consider the Gauss filtration on generic polynomials for the variable weights ω and coefficient parameter w.

The main monomial $(X_{s-d})_{bc}^{P^d}$ of the entry of (11.8) has weight at least $\alpha_{s,d+1}(b)-\alpha_{s,d}(c)-u$. For $j = s-d \ne i$, $d+1+[j-1-i]_\nu = d+[j-i]_\nu$, so its weight is exactly $\alpha_{s,d}(b) - \alpha_{s,d}(c)$.

For $j = i$ the weight is $A_{d+\nu}(b)$ and the desired expression may further subtract $A_d(c)$; since $d \le \lambda \le \nu$, its absolute value is at most u. This proves the main bound in both cases. An error monomial has coefficient divisible by w^2, and degree P^d, hence has Gauss order at least $2 - P^{d-\lambda} \ge 1$. Every $\alpha_{s,d}(b)$ and $\alpha_{s,d+1}(b)$ in this range has absolute value at most $1/2$ (the largest exponent is at most $d + \nu$), so its Gauss order is again sufficient. Lemma 11.8 applies.

For the special adjacent pair, the deep root variable y_s has $\delta = 1$ and z_s has $\delta = 0$, and the $\widetilde{x}$ of that lemma is precisely $(e/2)$ times the variable weight of $E_{i+\lambda}$. This is immediate by substituting $d = \lambda + 1$ for y_s and $d = \lambda$ for z_s, with the odd sign introduced by inverse transpose. At every other variable both sides vanish. On the generic exponent η the weight sum is $(e/2)t_i$: the other letters vanish by the row-column equality hypothesis. The filtration loss in (11.7) is at most $1/2$. Thus the required order of the coefficient is at least

$$\frac{e}{2}(E_{i+\lambda}(\beta) - t_i) - \frac{1}{2} > c.$$

In particular it is divisible by w^c, proving the proposition. □

This gives the coefficient calculation of `pow_dvd_coeff_conj_frobOperator` without an appeal to generic representation theory; the only input is the Frobenius congruence modulo p.

Let now $a_* = (k + k^2)d_0\nu^2$, $c_* = 2a_* + 2$, and assume $D + n_0 + k < P^\nu$ with the stronger size condition of (10.13), namely $\nu n_0 k(\lambda + 2) + n_0 + k < P^\lambda$. In particular $D + k < P^\nu$ and $n_0 < P^\nu = p^{d_0\nu}$. Take $F = d_0\nu$ in the diagonal pairing and the conjugation means, and take the homogeneous cutoff D. Let b_β be the product over letters of the integral reciprocal-multinomial diagonal entries q_{β_s}. Each divides $p^{kF\nu}$ in O (it is a quotient of that power by an integer multinomial product). Put $b'_\beta = p^{kF\nu}/b_\beta$, another integral diagonal entry. In tensor notation let B and B' denote these two diagonal matrices in the local coefficient space. They are distinct from the global budget B of Chapter 4; in this subsection the distinction is made by attaching a matrix argument, as in BF. They satisfy $BB' = p^{kF\nu}I$ and $B\rho(g^t) = \rho(g)^t B$ by (11.6).

Theorem 11.10 (Approximate projector)**.** *For every real h, there is an integral operator Z on the coefficient module such that*

$$ZF \in U_h \bmod w^{c_*} \quad \textit{for every } F,$$

and, for every $y_0 \in U_{h-(4c_*+2)} \bmod w^{c_*}$ and every F,

$$y_0^t BZF \equiv p^{a_*} y_0^t F \pmod{w^{c_*}}. \tag{11.9}$$

Proof. In every letter take the means of Proposition 11.7 for the weight x_b. Write Λ^{all} for the one with every row weight allowed, Λ^{hi} for that allowing row weight above $h - (2c_* + 1)$, and $\Lambda^{\mathrm{lo}} = \Lambda^{\mathrm{all}} - \Lambda^{\mathrm{hi}}$. These are invariant for conjugation in their letter, and vanish on unequal row-column weights. The last mean vanishes on row weights above $h - (2c_* + 1)$, since the first two agree there. Apply the tensor of the Λ^{lo} to the coefficients of the generic matrix $\mathscr{F}(\Psi)$ with $\Psi_{s,d} = X_{s-d}^{(P^d)}$, and multiply on the domain by B'. This defines Z.

To verify the first assertion, translate by any $g \in L$ and look at an output row β with excess above h in letter $i + \lambda$. Move the translate into the generic matrix using exact covariance and the invariant means; this changes the generic matrix to (11.8) and introduces only an integral change of the input. A monomial surviving the means has equal row and column weights at all letters, and row weight at i at most $h - (2c_* + 1)$. The difference in Proposition 11.9 is strictly greater than $2c_* + 1$. Therefore the output coefficient is divisible by w^{c_*}. We have obtained the first assertion for all Levi translates, not just for the chosen coordinates.

For the congruence, expand $\bigotimes_s(\Lambda^{\mathrm{all}} - \Lambda^{\mathrm{hi}})$. The all-term is $p^{k^2F\nu}$ times evaluation at the identity matrices, because the generic degrees are bounded by D. There the operator $\mathscr{F}$ is the identity. The contribution to the left of (11.9) is consequently $p^{k^2F\nu}y_0^tBB'F = p^{a_*}y_0^tF$.

Every other term contains at least one distinguished factor Λ^{hi}, say in letter i. By the span property of Property 3 in Proposition 11.7, express that factor as a linear combination of forms supported on row weights above $h-(2c_*+1)$ after a Levi conjugation. Invariance of the other factors permits the same conjugation in every letter. Move it off the generic matrices by covariance. The transpose of the conjugating coefficient action can be moved off B using $B\rho(g^t) = \rho(g)^tB$. Thus the new testing row vector is a Levi translate of y_0 with an additional diagonal factor B. Its coefficients at output rows of excess greater than $h-(4c_*+2)$ are zero modulo w^{c_*}; multiplication by B cannot spoil this vanishing. For all other output rows, the difference between the row weight supported in the distinguished generic factor and the output excess is strictly larger than $2c_*+1$. The nondistinguished factors still vanish on unequal row-column weights because they are invariant means. Proposition 11.9 applies again, so the entire contribution is zero modulo w^{c_*}. This proves (11.9). □

In `TwistingData.exists_projector` the cutoff constant is called c, here c_*. The estimate requires that the test vector lie in the smaller condition, not only that the output lie in the larger condition. This shift by $4c_*+2$ is precisely what makes the pairwise argument work after changing frames.

11.7 Primitive half-depth matrices

We pass from the preceding polynomial estimates to a determinant estimate for a sum over classes. First observe that the normalized half-depth operator is well-defined on all frames associated to the two-block parahoric. If $h \in \mathcal{P}$, any ordinary (even-letter) entry in the prohibited off-diagonal block is divisible by p under ι_s. Indeed an element x in the maximal ideal of R satisfies $1-x^{q-1}$ a unit; the full-cycle congruence $\iota_s(x)^q - \iota_s(x) \in (p)$ implies $\iota_s(x) \in (p)$. The same is true for the opposite prohibited entries in an odd letter by applying this to h^{-1} and transposing. On the coefficient variables, conjugating h_s by the single-letter diagonal with entries $w^{\delta(b,s)}$ therefore gives integral coefficients: a factor w^{-1} can occur only in the prohibited block, where it is canceled by a factor of p. Taking symmetric powers and then their tensor product gives

$$D_w\rho(h)D_w^{-1} \in \mathrm{GL}_{M_0}(O) \qquad (h \in \mathcal{P}). \tag{11.10}$$

Invertibility follows by the same assertion for h^{-1}. This is the half-depth version of the parahoric lemma of `TwistingData.lean`.

Let X be a finite class set and E_0 a finite set of frames with entries $x_e \in X$, $g_e \in \mathrm{GL}_k(R)$ such that $(x_e, g_e^{-1}\mathcal{P})$ covers X times the parahoric cosets. Allow a finite spectator index H_0, which need not undergo the nontrivial Frobenius action; equivalently replace X by $X \times H_0$ in the next construction. Consider the rectangular matrix

$$J : F = (F_x)_{x\in X} \longmapsto (D_w\rho(g_e)F_{x_e})_{e\in E_0}.$$

If all coordinates of JF are divisible by some scalar $a \in O$, the same is true of $D_w\rho(g)F_x$ for *any* x, g. This follows from surjectivity on the coset set and (11.10): if $g^{-1}\mathcal{P} = g_e^{-1}\mathcal{P}$, then $g = hg_e$ with $h \in \mathcal{P}$. In particular, J is injective over the fraction field of O, since all powers of w are nonzero and each x appears.

Proposition 11.11 (Primitive factorization)**.** *Write $J = Qd$ where d is square and Q is left invertible over O. Such a factorization exists and, with the maximal depth denoted by x_0, it can be chosen with*

$$\det d \mid w^{|X| \sum_{0 \le j < x_0} \dim V_j}.$$

Moreover every frame component of QF satisfies $U_h \bmod w^c$ whenever $h \ge 2c$ (at the boundary pair $r_0 - 1, r_0$).

Proof. Choose a Smith factorization of the full-column-rank matrix J over O. Such a factorization exists also for a nondiscrete valuation domain: select an entry among the finite matrix entries dividing all the others, clear its row and column by integral operations, and continue. Thus $J = Qd$ with Q primitive and d invertible over the fraction field of O. Denote its elementary divisors by t_i. There is an integral change of basis in the domain with basis vectors F_i for which the coordinates of JF_i are all divisible by t_i. By the all-frame divisibility just proved, t_i divides $w^{d(\alpha)}[X^\alpha]\rho(g)F_i(x)$ for every x, g, α. Since $d(\alpha) \le x_0$, taking $g = 1$ shows t_i divides every coordinate of $w^{x_0}F_i$. In an invertible basis each column has a unit coordinate, so $t_i \mid w^{x_0}$.

Let j_i be the smallest integer j with $t_i \mid w^j$. If $j < j_i$ then in every frame t_i divides the depth-weighted coefficient, yet t_i does not divide w^j. For any monomial of depth below j, therefore, the coefficient in that frame is a nonunit: indeed $w^{d(\alpha)} \mid w^j \mid t_i$ and the quotient $t_i/w^{d(\alpha)}$ cannot be a unit. Reducing these basis vectors modulo the maximal ideal gives independent vectors satisfying the depth condition in every frame, so

$$\#\{i : j_i > j\} \le |X| \dim V_j.$$

Multiplying divisors and summing this inequality over $j < x_0$ gives the asserted bound for $\det d$.

It remains to prove the excess condition on the primitive factor. The previous divisibility $t_i \mid w^{x_0}$ implies $w^{x_0}d^{-1}$ is integral; hence every column vector QF can be expressed as $w^{-x_0}JF'$ with F' in the integral domain of J. For each x, g and exponent α the inequality

$$w^{x_0} \mid w^{d(\alpha)}[X^\alpha]\rho(g)F'_x \tag{11.11}$$

holds by the all-frame property. Fix a letter i, and use the real variable weights $c(v) = \delta(v) - E_i(v)/2$ where E_i is for the boundary pair. For the two root variables the difference of weights is

$$c(y_s) - c(z_s) = 1 - (E_i(y_s) - E_i(z_s))/2 \in [2/5, 8/5],$$

using the explicit range $[-1/P, 1]$. For any fixed class x, all its full-group translates $\rho(g)F'_x$ have Gauss level at least x_0 with these weights. If not, choose a nonzero term with w-order l, a nonnegative *integer* defined by the largest power of w dividing that coefficient, and $l + c \cdot \alpha < x_0$. One can choose such a term with globally minimal integer order plus weight: for each exponent the set of attained finite integer orders below its violating threshold is finite, with a minimum, and there are finitely many monomials. Infinite valuation in powers of w cannot violate a finite bound. By the root lemma for the full group as Levi, that monomial must satisfy $E_i(\alpha) \le 0$. But (11.11) then forces

$$l + c \cdot \alpha = l + d(\alpha) - E_i(\alpha)/2 \ge x_0,$$

a contradiction. The Gauss inequality for all coefficients follows.

In a primitive frame the coefficient of QF is that of $w^{-x_0}D_w\rho(g_e)F'$. The Gauss inequality implies that a term of excess above $2c$ in letter i is divisible by w^c. Finally, D_w commutes with the block-Levi action: monomials are transformed within blocks, so their deep degree in each

letter is preserved. Thus the same conclusion applies in every further Levi translate of each frame. This gives membership in U_h mod w^c for $h \geq 2c$. The proof works with an arbitrary finite spectator index since every frame can be treated separately in those coordinates. □

This is the depth factorization `TwistingData.exists_depth_factorization`. Notice the use of the full-group root lemma in the second part, despite the final condition being only in block-Levi translates.

11.8 The Smith pairing lemma and the local Hecke theorem

We recall that the index condition $I(M, M'; z)$ of Chapter 10 means that every finite family in M can be made a family in M' by a square integral matrix whose determinant divides z. In a finite free module it suffices to exhibit such a matrix on a basis: more generally the condition on an arbitrary family then follows by pulling back under the map from its free coefficient module, as in Lemma 10.2. In particular a square matrix with determinant dividing z and image contained in M' gives the condition. We need the following consequence of Smith elimination.

Lemma 11.12 (Pairing rectangular factors)**.** *Let $Q, M : O^D \to O^{D'}$ be integral matrices, Q left invertible with inverse on its image R_Q, and $\beta : O^{D'} \to O^D$ integral. Suppose u_0, w are nonzero, w is a nonunit, and*

$$\beta M \equiv u_0 I \pmod{u_0 w}.$$

Suppose that for all v, v', F satisfying $Qv - Mv' = u_0 wF$, the image of F over the residue field belongs to a fixed vector subspace S of the reduced target. Then there is an integer $l \geq 0$ with $D + l \leq \dim S$ such that the submodule $\mathrm{im}(\beta Q) + zO^D$ has index dividing $u_0^D z^l$ for every $z \in O$.

Proof. Identify the target with the direct sum of the primitive image of Q and a complement; the latter is finite free over the valuation ring. Then Q is the first inclusion, and write $M = \binom{M_1}{M_2}$ relative to that decomposition. Make a Smith factorization of M_2 by changing the basis of the complement and the source of M, so it is diagonal rectangular with elementary divisors t_j, including zero divisors for missing pivots. Change the output basis of β compatibly with the new basis of the source of M; this preserves the congruence $\beta M \equiv u_0 I \bmod u_0 w$.

Let J be the set of positions with nonzero t_j dividing $u_0 w$ and $l = |J|$. The space S contains the reductions of all columns of Q (take $v = u_0 w e_j, v' = 0$). For $j \in J$ choose $v' = (u_0 w/t_j)e_j$, $v = M_1 v'$: the resulting vector F is (up to sign) the corresponding complement basis vector. These are independent modulo the image of Q, so $D + l \leq \dim S$.

For a position $j \notin J$, the divisor t_j is zero or is divisible by $u_0 w$ (valuation-domain comparability). The vector $\beta Q M_1 e_j = \beta M e_j - \beta\binom{0}{M_2 e_j}$ is congruent to $u_0 e_j$ mod $u_0 w$. For $j \in J$ instead use the vector $u_0 z e_j$, which lies in zO^D. These vectors are all in the desired submodule. Factor u_0 out of each of their columns and factor z from the columns $j \in J$. The determinant of the remaining matrix is congruent to one modulo w, and is therefore a unit. The index condition follows. This argument is independent of the complement coordinates, which were only changed by invertible integral matrices. □

This is `TwistingData.of_pairing`; the complement-block proof is an equivalent way to read the Smith factorization of the discrepancy matrix $(QR_Q M - M)R_Q$ used by the formal development.

We can now prove the precise local elementary-divisor estimate from Chapter 10. Let X be the finite class set, I_0 the spectator index set, and $V = O^{X \times I_0 \times \mathcal{M}}$, where $\mathcal{M}$ indexes the

monomials in the cyclic coefficient representation. Thus $D = \dim V = |X||I_0|M_0$. Write J_0 for the parahoric coset set. For each $e = (x, b) \in X \times J_0$ there are $s(e) \in X$, $g_e \in \mathrm{GL}_k(R)$ and an invertible spectator matrix u_e on O^{I_0}, with $e \mapsto (s(e), (g_e^t)^{-1}\mathcal{P})$ bijective (injectivity suffices since the finite sets have the same cardinality). Let T be the local Hecke matrix whose block from $s(e)$ to x is

$$u_e \otimes (\rho(g_e)D_w^2\rho(\gamma_b))^t,$$

where γ_b are the specified representatives with inverses representing the cosets J_0. Contributions with repeated class pairs are summed. This is the matrix of (10.18) in its chosen orientation. Define two half-depth maps into the frame module

$$(J_1F)_e = D_w\rho(g_e^t)F_{s(e)},$$
$$(J_2F)_e = D_w\rho(\gamma_b)F_x,$$

with trivial action on the spectator coordinates. Both sets of frames cover the requisite cosets. Write $J_1 = Qd$, $J_2 = Q'd'$ as in Proposition 11.11; both Q, Q' have integral left inverses. In the frame module let U be the block-diagonal operator acting as u_e on the spectator factor in frame e, and put $\mathbb{B}$ for the repeated diagonal matrix B of Theorem 11.10. Then by the transpose identity and $D_wB = BD_w$ one has an exact matrix identity, with B also repeated over the domain:

$$TB = J_2^tU\mathbb{B}J_1 = d'^t(Q'^tU\mathbb{B})Qd. \tag{11.12}$$

To check it, the contribution from frame e on the right, before the primitive factorizations, is $u_e \otimes \rho(\gamma_b)^tD_wBD_w\rho(g_e^t)$. Moving B to the right using its transpose covariance gives exactly the required block times B. This calculation is why the automorphic Hecke operator must be identified with the specific transpose action rather than merely an abstract endomorphism with the same eigenvalues.

Set $\beta = Q'^tU\mathbb{B}$ and choose a left inverse R' of Q'. Let Z of Theorem 11.10 act in each frame on each spectator factor, and define the rectangular matrix $M = U^{-1}ZR'^t$. The second primitive factor has every frame component in $U_{h-(4c_*+2)}$ mod w^{c_*}: indeed $h \geq 6c_* + 2$ by the size hypotheses, so this threshold is at least $2c_*$. Therefore the approximate projector gives, by summing over the finitely many frames and spectator coordinates,

$$\beta M = Q'^t\mathbb{B}ZR'^t \equiv p^{a_*}Q'^tR'^t = p^{a_*}I \pmod{w^{c_*}},$$

where $a_* = (k + k^2)d_0\nu^2$ and $c_* = 2a_* + 2$. Both Q and M have components in U_h mod w^{c_*}; for M this follows from the projector and the fact that U^{-1} acts only on spectator coordinates. If $Qv - Mv' = p^{a_*}wF = w^{c_*-1}F$, cancellation of w^{c_*-1} in the integral congruences means that F satisfies the excess condition modulo w, hence modulo the maximal ideal. We can take for S in Lemma 11.12 the direct sum of copies of U_h in every frame and spectator coordinate. The flag estimate (11.4) gives

$$\dim S = |X||J_0||I_0| \dim U_h \leq (1 + \varepsilon)D.$$

Lemma 11.12 yields an index condition for $\mathrm{im}(\beta Q) + p^sV$ with exponent at most $(a_* + s\varepsilon)D$ in p. Left and right multiplication by the square factors in (11.12) adds their index costs. Each determinant divides $w^{|X||I_0|\sum_j \dim V_j}$, whose sum of two exponents is at most $2(1 + z)D$ by (11.5); using $w^2 = p$ the corresponding extra exponent is $(1 + z)D$. The determinant pullback, pushforward and composition rules apply even after adding p^sV: an integral matrix sends p^sV into itself, and passing to the image can only reduce the required determinant. Consequently

$$I(V, TBV + p^sV; p^{(a_*+1+z+s\varepsilon)D})$$

in the sense that an integer exponent bounded by the displayed real exponent suffices. The $a_* + 1$ term is $c_*/2 \le h$, so the bound is at most $(z + \varepsilon s + h)D$. If $f : V \to M'$ is any integral linear map under which the Hecke action descends to Y on $f(V)$, then $f(TBV) \subseteq Yf(V)$ and the pullback rule transfers the same estimate to $Yf(V) + p^s f(V)$. The action preserves $f(V)$ by its very intertwining equation. This completes a mathematical proof of `TwistingData.hasType_twisted_hecke`, including the leading constants of the dimension and index estimates.

In the global application $n_0 = a$, $r_0 = s_p$ is the rank of the marked Hecke block, $d_0 = o_0$ and the local residue field has cardinality $q = P^\nu$. The spectator factor consists of the weight variables belonging to other places and letters, and all computations above act trivially on their monomial indices except for the integral matrices u_e. The injective correspondence of local frame labels, verified on the actual column class model, is the local-to-global formal-source identification `Data.IsMarked.hasType_markOperator` in `New/Columns/MarkedPrimeType.lean`. The present proof discharges the abstract coefficient-matrix theorem; it should not be read as supplying the local-to-global identification of g_e and u_e if the chosen level or weight representation is changed.

11.9 The finite Fourier operator at the level prime

The arithmetic prime p in the preceding argument is a prime of the base of the column coefficient lattice. We now work at the different prime N, the auxiliary split level prime of the chain. Unlike the twisted Hecke estimate, the Fourier operation at this prime is based on a small finite sum with no divisions by the size of a residue field. We give the argument with all the scalar factors; it supplies a useful test of the integrality of the denominator trade described in Chapter 10.

Let S be the finite chosen set of split places of E above N, $\Lambda = \prod_{w\in S} \mathcal{O}_{E_w}$, $F_S = \Lambda/N\Lambda = \prod_{w\in S} \mathbb{F}_N$. All definitions below refer to the places in S, not to their conjugates. Write V_j for a representation of $G_j = \mathrm{GL}_{k_j}(\prod_{w\in S} E_w)$, with action π_j, and let e_j denote the last standard basis vector of rank k_j. An F_S-valued coefficient vector is called everywhere nonzero if its image in $\mathbb{F}_N^{k_j}$ is nonzero at every $w \in S$; this is stronger than being a nonzero element of $F_S^{k_j}$. A scalar in F_S is a unit if each component is nonzero. The finite additive character used below is exactly

$$\psi_N(t) = \psi_0(-\,\mathrm{residueOverN}(t)),$$

where the w-component of $\mathrm{residueOverN}(t)$ is the chosen integer representative of t_w divided by N, and all components outside S vanish.

A residue function is extended to local matrices by zero outside the integral matrices. In rank k_{j+1} by k_j define three such functions:

$$b_y(X) = [\overline{X} e_j = y], \qquad e_z(X) = [e_{j+1}^t \overline{X} = z^t],$$
$$c_z(X) = \psi_N(e_{j+1}^t \overline{X} z).$$

Here y has length k_{j+1}, z length k_j, brackets denote an indicator, and $b = b_{e_{j+1}}$, $e = e_{e_j}$. Put $u_j = \mathrm{diag}(1, \dots, 1, N^{-1})$ at every place in S. These are the finite data `lastColDatum`, `lastRowDatum`, `charDatum`, and `diagNInv` in `New/Seesaw/NamesAtSplitLevel.lean`. For any rectangular local matrix X, one has the exact identity

$$e(u_{j+1}^{-1} X u_j) = b(X). \tag{11.13}$$

Indeed, in blocks the transformation sends $\left(\begin{smallmatrix} A & v \\ w & d \end{smallmatrix}\right)$ to $\left(\begin{smallmatrix} A & v/N \\ Nw & d \end{smallmatrix}\right)$. For its result to be integral with last row congruent to $(0, \dots, 0, 1)$ requires A, w, d integral, $v \in N\Lambda^{k_{j+1}-1}$ and $d \equiv 1$

(mod N); these say just that $b(X) = 1$. This proves the identity also when the matrix is not integral. Omitting that case would not suffice for the kernel covariance calculation.

A lift law consists of a linear map $f \mapsto L_j[f]$ from the space of local functions to linear maps $V_j \to V_{j+1}$, together with a scalar γ_j satisfying

$$\pi_{j+1}(B)L_j[f(\overline{B} \cdot \overline{A})]\pi_j(A) = L_j[f], \qquad A \in \mathrm{GL}_{k_j}(\Lambda), \tag{11.14}$$

$$\pi_{j+1}(u_{j+1})L_j[e] = \gamma_j L_j[b]\pi_j(u_j). \tag{11.15}$$

The notation in (11.14) also requires $B \in \mathrm{GL}_{k_{j+1}}(\Lambda)$ and f a residue function. This is precisely the structure `IsSplitLevel.LiftLaw`. The existence of the law for a concrete Weil kernel, with its correct normalization γ_j, is a separate assertion about kernel equivariance and Haar factors (see `New/Seesaw/LiftCharacters.lean`); we use that formal-source input when applying the algebra below. In particular (11.13) is the support calculation in the second law, not by itself a determination of its scalar.

Call a vector $v \in V_j$ ramified at S if, for each $w \in S$, there is an integral scalar unit a_w whose reduction is 1 at each other place, such that $v \in \mathrm{im}(1 - \pi_j(a_w I))$. This weaker formulation than a fixed character eigenspace is the exact one used in the formal proof. It immediately implies an annihilation rule. If a residue function is unchanged by multiplication of its matrix argument by every unit scalar supported at w, then $L_j[f]v = 0$: choose a_w witnessing ramification and apply (11.14) with $A = a_w I$, $B = I$. The operator $L_j[f](1 - \pi_j(a_w I))$ is zero. The same argument works for any function invariant under the specific selected a_w.

For $y \in F_S^{k_j}$ with y_{last} a unit, select an integral invertible K_y whose reduction has both last row and last column given by y (with the row transposed). Such a choice is possible: for $y = (v, t)$ the matrix

$$\begin{pmatrix} I + vv^t/t & v \\ v^t & t \end{pmatrix}$$

has determinant t, by the Schur complement formula, and satisfies both conditions. Lifting arbitrary entries to Λ preserves invertibility since its determinant is a residue unit. Define

$$\mathcal{F}_j = \pi_j(u_j) \sum_{\substack{y \in F_S^{k_j} \\ y_{\mathrm{last}} \in F_S^\times}} \psi_N(y_{\mathrm{last}})\pi_j(K_y). \tag{11.16}$$

This operator is defined using a chosen set-theoretic section; its action on the indicated invariants is what matters, not an assertion of independence of its choices on the whole representation. In particular, $\mathcal{F}_j$ commutes with the central scalar operators used to define ramification.

Proposition 11.13 (Sliding the finite Fourier operator)**.** *For a ramified vector v and a lift law as above,*

$$\pi_{j+1}(u_{j+1})L_j[c_{e_j}]v = \mathcal{F}_{j+1}L_j[b]v = \gamma_j L_j[b]\mathcal{F}_j v. \tag{11.17}$$

Proof. The function $c_{e_j}(X) = \psi_N(\overline{X}_{\mathrm{last,last}})$ is the sum over all bottom right entries $t \in F_S$ of $\psi_N(t)[\overline{X}_{\mathrm{last,last}} = t]$ on integral matrices. If t is a nonunit, there is a place w with $t_w = 0$. The displayed indicator is unchanged under central scalars supported at w. The annihilation rule shows its lift vanishes on v, so we may restrict the sum to t a unit. We may refine the remaining indicators either according to their entire last column or to their entire last row, obtaining

$$L_j[c_{e_j}]v = \sum_{y_{\mathrm{last}} \in F_S^\times} \psi_N(y_{\mathrm{last}})L_j[b_y]v = \sum_{z_{\mathrm{last}} \in F_S^\times} \psi_N(z_{\mathrm{last}})L_j[e_z]v.$$

The two sums have respectively vectors of target and source rank as their indices. By (11.14),

$$L_j[b_y] = \pi_{j+1}(K_y)L_j[b], \qquad L_j[e_z] = L_j[e]\pi_j(K_z),$$

since $b_y(\overline{K}_y X) = b(X)$ by the last-column condition and $e_z(X\overline{K}_z) = e(X)$ by the last-row condition. For example, the latter condition is equivalent to $\mathrm{lastrow}(X)\overline{K}_z = z^t = e_j^t\overline{K}_z$, hence to $\mathrm{lastrow}(X) = e_j^t$. Inserting these function identities in (11.14) gives exactly the two displayed lift identities. Applying $\pi_{j+1}(u_{j+1})$ to the column expansion gives the middle member of (11.17). Applying it to the row expansion and using (11.15) gives the final member. □

We have used precisely the two lift laws and ramification, with no averaging over the residue field. This is the proof of `IsSplitLevel.LiftLaw.u0_charDatum_eq_fourierOp`.

Proposition 11.14 (The paired chain)**.** *Suppose the consecutive lifts $L_0, L_1, \dots$ satisfy lift laws with constants $\gamma_0, \gamma_1, \dots$. Let d_i be arbitrary scalars and define the paired lift from rank k_{2i} to rank k_{2i+2} by*

$$\mathcal{P}_i = d_i\pi_{2i+2}(u_{2i+2}) \sum_{y \in F_S^{k_{2i+1}}} L_{2i+1}[c_y]L_{2i}[b_y].$$

Write $B_j = L_{j-1}[b]\cdots L_0[b]$ and $B_0 = 1$, with the standard last-column datum of the correct rank in every occurrence. For h pairs and a ramified initial vector v,

$$\mathcal{P}_{h-1}\cdots\mathcal{P}_0 v = \left(\prod_{0\le i<h} d_i q_i \prod_{0\le l\le 2i+1} \gamma_l\right) B_{2h}\mathcal{F}_0^h v, \qquad q_i = \prod_{w\in S}(N^{k_{2i+1}} - 1). \tag{11.18}$$

Proof. If y is not everywhere nonzero, b_y is invariant under central scalar units at a place with $y_w = 0$ and its lift on a ramified vector vanishes. For an everywhere nonzero y choose any K of middle rank with last column reducing to y; at every w a nonzero vector can be completed to a basis. Equivariance gives

$$L_{2i}[b_y] = \pi_{2i+1}(K)L_{2i}[b], \qquad L_{2i+1}[c_y]\pi_{2i+1}(K) = L_{2i+1}[c_{e_{2i+1}}],$$

the latter because $c_y(X) = c_{e_{2i+1}}(X\overline{K})$. Each surviving summand therefore has the same value; there are q_i of them. It follows from (11.17) that on a ramified vector $\mathcal{P}_i$ acts as

$$d_i q_i \mathcal{F}_{2i+2} L_{2i+1}[b]L_{2i}[b].$$

The ramification condition is preserved both by $\mathcal{F}_j$ and by $L_j[b]$. For the latter, use (11.14) with $A = a^{-1}I$, $B = aI$: the transformed matrix is unchanged. Hence $L_j[b]$ carries an image of $1 - \pi_j(aI)$ into the corresponding image of $1 - \pi_{j+1}(aI)$. On these vectors (11.17) gives $\mathcal{F}_j B_j = (\prod_{l<j}\gamma_l)B_j\mathcal{F}_0$, by induction on j. Starting with $h = 0$, induction on h now proves (11.18). This is the formula `IsSplitLevel.compUpTo_pairs_eq_smul`. □

Notice that each q_i is an N-adic unit. There is an important asymmetry between the constants q_i and γ_j: the former are known units directly by counting, whereas the latter require their precise Weil/kernel normalization. One must not replace the q_i by the count $N^{k_{2i+1}|S|} - 1$, which would overlook the several residue factors. In the arithmetic application the normalized scalars from the kernels are combined with the d_i in the formal-source theorem `Data.chain2_pairTranslates_eq_smul_chain` (subscript 2 in the source) of `New/Columns/TwistedHeckeType.lean`. The formula over local residue functions alone does not compute those scalars.

Lemma 11.15 (The mirabolic preservation calculation)**.** *For a fixed rank let $\mathcal{M} \subset \mathrm{GL}_k(\Lambda)$ be the matrices with last column congruent to the last basis vector. Let $g \in \mathcal{M}$ and let T be an endomorphism of a representation π. Suppose a vector v has the following property: $\pi(A)T\pi(B)v = v$ for all A, B with $AgB \in \mathcal{M}$. Then $T\mathcal{F}v = \mathcal{F}v$.*

Proof. Put $K = u^{-1}gu$. Both K and K^{-1} are integral, with reduction having last row e^t: apply (11.13) to g and g^{-1}. Thus $y \mapsto y' = \overline{K}y$ permutes the vectors with unit last entry and preserves y_{last}. For these vectors set $A = (uK_{y'})^{-1}$, $B = uK_y$. The reduction of $K_{y'}^{-1}KK_y$ fixes the last basis vector, so $AgB \in \mathcal{M}$. The hypothesis yields $T\pi(uK_y)v = \pi(uK_{y'})v$. Summing with coefficient $\psi_N(y_{\mathrm{last}}) = \psi_N(y'_{\mathrm{last}})$ proves the lemma. □

This is `IsSplitLevel.mirabolic_fix_fourierOp`. In the actual initial column space the assertion that the conjugated column level lies in this mirabolic is `Data.atPlaces_conj_mem_mirabolic` of `New/Seesaw/LiftCharacters.lean`; it includes the conjugate-place inverse-transpose condition. Using that assertion and the precise compatibility of the representation π with the smooth right action, the lemma is the preservation step of `Data.fourierOp_mem_column`. One still has to preserve integrality in arithmetic frames, which is stronger than preservation of the space of columns.

Example 11.16. If $N = 5$, two places occur in S, and a middle column has rank two, there are 25^2 residue vectors in total. Only $(25-1)^2 = 576$ enter the surviving paired sum. An individual Fourier operator of rank two, on the other hand, sums over $(5 \cdot 4)^2 = 400$ vectors with last entry a unit. Although neither count is the total field size, the paired count is prime to 5, exactly as required. This illustrates why the row/column distinction in Proposition 11.13 cannot be replaced by a sum over all vectors regardless of ramification.

11.10 Geometry and integrality in passing to a seesaw

The local calculation just proved answers a conditional question: if the data of two consecutive lifts are the indicated residue functions, what operator must be transferred to the end of the chain? Showing that the actual period can be reduced to that situation is a separate geometric and arithmetic question. We record the logical order of these reductions so that they are not concealed inside the local Fourier formula.

For a sequence of adjacent ranks, the lift $L_j[f]$ is an integral transform with a matrix parameter. The sum over global rational vectors that occurs in the theta kernel is invariant under simultaneous changes of two frames; one does not have the right to integrate each stage against a separate independent copy of a rational orbit. Instead one first uses the jointly invariant product of kernels on the full chain. Fold the odd-rank groups into the even-rank groups by integrating once over their common stabilizers. The relevant source statements are `Data.unfoldStage_even`, `Data.hasNearLines_of_unfoldStage_ne_zero`, and `Data.exists_stage_cut`, as described in Chapter 10. These are formal-source inputs here: they include the explicit global orbit classification and supports of the named Weil kernels.

At a finite cut one may replace integration over a compact class space by a finite weighted sum, provided the integrand is constant on the chosen classes. In particular if C is a compact domain with a finite disjoint partition $C = \coprod C_a$ and f is constant $f(a)$ on each measurable C_a, then

$$\int_C f = \sum_a \mathrm{vol}(C_a) f(a).$$

If a stabilizer acts with a finite orbit of size s_a on the chosen finite classes, the resulting multiplicity is s_a only when the volumes and values are equal on that orbit. More generally its contribution is the sum of those volumes times the common value. Applying this identity iteratively to the joint product of kernels requires a simultaneous cut, not one chosen independently for every appearance of a repeated class variable. The needed invariant, admissible finite cut is the formal-source input `Data.exists_isAdmissibleComb_unfolding`; its arithmetic expression by integral class frames is the input `Data.exists_isGoodDatum_of_invariant`. These are exactly the steps where one must establish finite support before redistributing scalar constants. The Gauss-sum calculation above can then be applied to the resulting chain of named finite data. It cannot replace the cut construction.

For clarity about the lattice index which remains after this reduction, suppose a coefficient module is a direct sum over a finite class set of integral polynomial coefficient representations. Let T_{last} denote the sum, over classes, monomials and embeddings contributing to the level prime, of the last-coordinate exponent. For homogeneous degree a in any given letter of rank k, the monomial basis is permuted by permuting the k coordinates; the sum of the exponents in all coordinates is a times the number of monomials. Consequently the sum in just the last coordinate is a/k times that number. The same statement holds with independent spectator letters, by multiplying their numbers of monomials. If D denotes the full dimension, f the number of relevant letters, the contribution of inverse level diagonal operators on this ambient space is bounded by faD/k. Transpose or conjugate letters are handled by taking the negative part of the diagonal exponent, so the resulting bound remains an upper estimate for denominator exponents. This is the determinant-counting mechanism behind the $1/k$ saving in the index bound for the Fourier operator.

To see how a denominator estimate composes, let L be an integral lattice of rank D over a valuation ring, and $g_1, \dots, g_s$ invertible endomorphisms of its fraction space. Suppose L_i is a lattice containing both the intermediate image and the appropriate integral frame lattice, and that the elementary divisors of the required change of frame at step i have total negative part at most b_i. A denominator-removing sublattice exists with determinant cost at most $\sum_i b_i$. Indeed, for a diagonal matrix of powers of a uniformizer, replace each basis vector whose exponent is negative by that vector multiplied by the inverse negative power; the resulting sublattice has the claimed determinant. Integral changes of bases preserve the determinant ideal, and successive intersections or inverse-images of sublattices add at most their determinant costs by the index rules of Chapter 10. This proves the same assertion for arbitrary matrices with these elementary divisors. Finite sums of frame operators must first be treated on a common cut or common lattice where all the summands act; in particular the argument is not a license to multiply the bound by no factor without first proving the simultaneous class-frame assertion.

In the source the resulting cost for the power of the level Fourier operator includes the extra change back to the reference class frame. The arithmetic statement `Data.exists_indexDvd_fourierOp_pow` in `New/Seesaw/LiftCharacters.lean` does this uniformly, using integral representatives of the CM classes and a product of inclusion/index conditions. It is a formal-source input here beyond the elementary weight-sum and composition calculation just given. The power transferred by (11.18) is $m/2$; the extra frame change is why the denominator allowance includes $m/2+1$. Likewise, the passage to the branching sublattice and its comparison with the non-isotypic column representation require the specified archimedean coefficient and complex CM character estimates, not finite-vector counting alone.

The principal improvement in this chapter is thus the proof of the abstract twisted Hecke estimate, including its nonsemisimple finite-field component and integral approximation, and the proof of the two Fourier-sliding identities. The identification of those operators with the specific marked-prime columns, and the arithmetic and geometric machinery realizing the

seesaw, are still explicitly identified formal-source inputs. Keeping the distinction makes it possible to reuse the local estimate independently while leaving the original arithmetic proof's normalization intact.

Chapter 12

Normalized CM reciprocity and the geometry of the iterated seesaw

The passage from an adelic chain with varying signature to a chain on definite groups is the essential geometric step in the upper bound. There are two different issues. First, the finite vectors in the definite chain must be related to the arithmetic theta series of Chapter 8, and the relation must hold under conjugation of the coefficients. Secondly, the data of the definite chain must be obtained with coefficients integral away from the twist conductor, not simply as an arbitrary element of a complex tensor product. We give proofs of the normalized reciprocity statement, the rational orbit calculation underlying one unfolding, the rigidity of integral lattice lines, and the finite-cut calculation including its exact multiplicities. We then explain the assembly of these results into the two-edge evaluation and iterated seesaw. The conclusions of the latter construction distinguish the elementary reductions established here from the remaining specialized comparison of finite and infinite Weil models.

12.1 A marked algebraic curve and its sections

We first treat the normalization which was left open in Chapter 8. Use $K = \mathbb{Q}(\sqrt{-d})$ with d prime and $d \equiv 7 \pmod 8$, and write $\omega = (1 + \sqrt{-d})/2$. To avoid confusion with the Weil action, denote a field automorphism fixing K by γ. For a CM lattice C let L be its period lattice, t its distinguished half-period and $a_1 = a_1(C)$ its scalar in the equation

$$Y^2 + a_1 XY = X^3 + X.$$

As in the preceding CM chapter, write θ_C for the even theta function of degree one normalized to have value one at zero. Write s_C for the canonical Hermitian sigma section, $G_C = \phi_C \theta_C^3$ and $s_C^2 = \xi_C \theta_C^2$, where ξ_C, ϕ_C are the translated elliptic coordinates. In particular the two functions s_C^2, G_C are entire. It is essential that the sigma function here is the one with the Hermitian quasiperiod, not the uncorrected Weierstrass product. We use the analytic uniformization theorem for elliptic curves and the algebraicity of morphisms of elliptic curves, in their classical characteristic-zero forms.

The marked origin. The uniformizing coordinates are

$$\xi_C(z) = \wp_L(z + t) - \wp_L(t), \qquad \phi_C(z) = \tfrac{1}{2}(\wp_L'(z + t) - a_1 \xi_C(z)).$$

The point with uniformizing coordinate zero is the affine point $(0, 0)$, while the point with coordinate t is the point at infinity. We take $(0, 0)$ as the origin of the group law, translating

the usual group law of the projective cubic by that rational point. This is an algebraic group law. The invariant differential is $dX/(2Y + a_1X)$, whose pullback is dz by the identity for ξ'_C. In particular the scalar multiplier of an isogeny is its multiplier on this differential. Conjugation of the equation and of the rational functions defining an isogeny transports this differential and changes the multiplier a to $\gamma(a)$. No application of γ to an analytic exponential is involved.

Lemma 12.1 (Marked conjugate curve)**.** *There is a CM lattice D of the same discriminant $-d$, with $a_1(D) = \gamma(a_1(C))$. If C_1, C_2 and their conjugates D_1, D_2 are so marked, an isogeny given analytically by multiplication by a from C_1 to C_2 is transported to multiplication by $\gamma(a)$ from D_1 to D_2, with the same degree. Its action on the distinguished half-periods is transported as well.*

Proof. Conjugate the cubic and its nonzero discriminant. Uniformize it, using the differential just specified to fix the analytic scale, and choose the distinguished half-period to be the point at infinity in the translated group law. The explicit conversion to Weierstrass coordinates gives the required equalities for $\wp(t), g_2, g_3$. To check the extra CM conditions, use that on a complex elliptic curve a homomorphism of tori is algebraic: its pullback on meromorphic functions is an inclusion of elliptic function fields, hence a morphism of the projective nonsingular curves. The endomorphisms $a \in \mathcal{O}_K$ conjugate to endomorphisms on D with analytic multiplier $\gamma(a) = a$. Consequently $\mathcal{O}_K L_D \subset L_D$. The two endomorphisms $\omega, \bar{\omega}$ do not kill the half-period of C by the definition of the marked lattice; they cannot kill its transported point, since the conjugation on points is bijective and preserves all group operations and identities. This proves the half-period condition on D. Conjugation of an isogeny preserves the degree, for example by the induced extension of function fields, and its differential gives its analytic multiplier. It also preserves the condition that the image of the marked point equals the origin. □

This is a convenient conventional proof of the algebraic part of `CMLattice.exists_galConjugate` and `conj_isogeny_mem_lattice`, in `Lib/CM/MainTheorem/ConjugateCurve.lean`. The formal source implements the conjugation of isogenies by comparing sections. Both approaches have to retain the specified differential and half-period; a statement just about the j-invariant does not suffice.

The canonical multiplier. Let A_L be the area of the lattice, let $H_L(z, w) = \pi z\bar{w}/A_L$, and let E_L be the integral alternating Riemann form, normalized so that its values on a positive lattice basis are ± 1. The multiplier of θ_C on the lattice is $\chi_C(l) = \mathrm{sgn}_\sigma(l) \exp(2\pi i E_L(l, t))$. Here the sigma sign is $+1$ on $2L$ and -1 on the other classes. There is a useful expression for the multiplier not involving t:

$$\chi_C(l) = (-1)^{E_L(\bar{\omega}l, l)}. \tag{12.1}$$

Indeed $\mathcal{O}_K/(2) \simeq \mathbb{F}_2 \times \mathbb{F}_2$, and $\omega, \bar{\omega}$ are complementary idempotents modulo 2 because their sum is 1 and $\omega\bar{\omega} = (d+1)/4$ is even. The lattice is a rank-one projective $\mathcal{O}_K$-module, so $L/2L$ is the sum of the two one-dimensional eigenspaces. The alternating form on this quotient is perfect. Thus the right side of (12.1) is negative precisely on the element with both nonzero components. As ωt and $\bar{\omega} t$ are not in L, that element is the class of $2t$. The sigma sign and the alternating pairing with $2t$ give the same rule on the four elements of $L/2L$. This proves (12.1).

Let multiplication by a be an isogeny $C \to C_2$ of degree N. Then

$$H_{L_2}(az, aw) = NH_L(z, w), \qquad E_{L_2}(az, aw) = NE_L(z, w).$$

Inserting $\bar{\omega}l$ and l into these equalities shows from (12.1) that $\chi_{C_2}(al) = \chi_C(l)^N$. A degree-$j$ section of C_2, pulled back by multiplication by a, is therefore a degree-jN section of C. This is the full character compatibility of `CMLattice.comp_mul_mem_sections` in `Lib/CM/Sigma/CurveTheta.lean`. Its proof does not require the isogeny to have odd degree.

Lemma 12.2 (Section transport and point evaluation)**.** *Let $S_m(C)$ be the sections of degree m. Put*

$$R(z) = \frac{\mathrm{Tay}_0\,\theta_D(z)}{\gamma(\mathrm{Tay}_0\,\theta_C(z))} \in 1 + z\mathbb{C}[[z]].$$

There is a γ-semilinear isomorphism $\mathcal{T}_m : S_m(C) \to S_m(D)$ satisfying

$$\mathrm{Tay}_0\,\mathcal{T}_m f = R^m \gamma(\mathrm{Tay}_0\, f). \tag{12.2}$$

At any conjugate pair of points $P \in C(\mathbb{C})$ and $P^\gamma \in D(\mathbb{C})$, take lifts z, z' to their period planes. A nonzero scalar λ depending on the lifts but not on m or f satisfies

$$(\mathcal{T}_m f)(z') = \lambda^m \gamma(f(z)). \tag{12.3}$$

In particular transport preserves the divisor and multiplicities of the zero set on the elliptic curve.

Proof. First divide f by θ_C^m. The quotient is elliptic and meromorphic, with its only possible pole at the marked half-period t. In elliptic coordinates every such function is a unique expression

$$P(\xi_C) + \phi_C Q(\xi_C), \qquad 2\deg P \le m, \quad 3 + 2\deg Q \le m;$$

the inequality for P is omitted if $P = 0$, and likewise for Q. To see this, functions regular away from t are the coordinate ring of the affine cubic. Its relation is monic quadratic in Y, so the displayed expressions are unique. Their pole orders at t are respectively even and odd and cannot cancel, since ξ_C has a pole of order two and ϕ_C one of order three. Conversely each displayed monomial times θ_C^m is entire and has the required multiplier. This also proves the monomial spanning statement in the source.

Apply γ to the coefficients of P, Q and use the corresponding functions of D, multiplied by θ_D^m. This defines $\mathcal{T}_m$. The formal Taylor coefficients of ξ_C, ϕ_C are recursive rational polynomial expressions in $a_1(C)$: the differential equations

$$\xi' = 2\phi + a_1\xi, \qquad \phi' = 1 + 3\xi^2 - a_1\phi, qquad\xi(0) = \phi(0) = 0$$

determine every coefficient in characteristic zero. Hence their Taylor series are carried to those on D. This proves (12.2).

Off the marked point use $\lambda = \theta_D(z')/\gamma(\theta_C(z))$. The algebraic functions ξ, ϕ have the conjugate values at the two points, so (12.3) follows. At the marked point use the alternative generators s_C^2, G_C of degrees two and three. Both are nonzero there and $G_C^2 = s_C^6$ there, by the cubic identity. Set $a = s_D(z')^2/\gamma(s_C(z)^2)$ and $b = G_D(z')/\gamma(G_C(z))$. Then $b^2 = a^3$ and $\lambda = b/a$ gives $\lambda^2 = a$, $\lambda^3 = b$. All monomial sections are products of θ_C, s_C^2, G_C, so this gives the same evaluation property at the marked point. The algebraic expressions identify the sections of $\mathcal{O}(mt)$; their divisors, including vanishing orders at t, are preserved under algebraic conjugation. □

The source analogues are `thetaQuotient`, `taylorAtZero_monomialSection`, and `exists_conjugate_section` in `Lib/CM/LocalHeight/IsogenyIntegrality.lean`. Observe why the unknown factor R^m is allowed at this stage. The isomorphism of section rings is algebraic, but the analytic trivializations by normalized theta functions may differ. The following rigidity argument is what removes this factor.

Proposition 12.3 (Rigidity of the normalized theta Taylor series)**.** *For the conjugate CM lattice of Lemma 12.1, one has*

$$\mathrm{Tay}_0\,\theta_D = \gamma(\mathrm{Tay}_0\,\theta_C).$$

Proof. For an isogeny $a : C \to C_2$ of degree N, the section $F(z) = \theta_{C_2}(az)$ belongs to $S_N(C)$. Its divisor is the inverse image of the marked point under the isogeny, with each point simple. By Lemmas 12.1 and 12.2, $\mathcal{T}_N F$ has that same divisor on the conjugate curve as $\theta_{D_2}(\gamma(a)z)$. The ratio of these two sections has no pole or zero, so is constant on a compact elliptic curve. At zero both equal one, by (12.3) with the normalized lift zero (or directly by (12.2)); the constant is one. Taking Taylor series gives $R_2(\gamma(a)z) = R_1(z)^N$.

Take $C = C_2$ and the two endomorphisms $a = 2$ and $a = \omega$. If $R \neq 1$, let $j \geq 1$ be the first positive index with coefficient $r_j \neq 0$. The first relation, $R(2z) = R(z)^4$, gives $2^j = 4$, hence $j = 2$. The second relation gives $\omega^2 = \omega\bar{\omega}$. Since $\omega \neq 0$ and $\omega \neq \bar{\omega}$, this is impossible. Notice that the degree used here is the actual norm, not a degree inferred from an action on a formal parameter. This proves $R = 1$. □

The result is `CMLattice.thetaQuotient_eq_one` in `Lib/CM/MainTheorem/ConjugateCurve.lean`. Conjugation of the normalized theta coefficients is much stronger than conjugation of the cubic. This proof has no appeal to a theorem on transcendental periods.

12.2 Simultaneous torsion translations and finite-adelic reciprocity

The next argument is needed because the upper bound does not involve only the theta function at the origin. It involves finite-adelic translates of the generating vector. For y in the complex period plane set

$$U_y\theta_C(z) = \exp\bigl(-\pi(z\bar{y} + |y|^2/2)/A_L\bigr)\theta_C(z + y).$$

The addends in the exponent are not individually algebraic, so coefficient conjugation is not to be applied to them separately.

Proposition 12.4 (One multiplier for a finite set of translates)**.** *Assume $L_C = \Omega_0\mathcal{O}_K$. Given a finite subset $S \subset K$, there is $v \in L_D$ such that for every $x \in S$*

$$\gamma(\mathrm{Tay}_0\,U_{x\Omega_0}\theta_C) = \mathrm{Tay}_0\,U_{xv}\theta_D. \tag{12.4}$$

Proof. Choose a positive integer n with $nS \subset \mathcal{O}_K$. Let $u = \Omega_0/(2n^2)$ and choose a complex lift w on D of the point conjugate to u on C. As $2n^2u \in L_C$, the conjugate point is killed by $2n^2$; hence $v = 2n^2w \in L_D$. The scalar w is chosen once for all $x \in S$. The point nxu is transported to nxw because $nx \in \mathcal{O}_K$ and the marked CM action is preserved.

Fix $x \in S$, set $y = x\Omega_0$, $y' = xv$, and compare the two functions

$$F_C(z) = U_y\theta_C(nz), \qquad F_D(z) = U_{y'}\theta_D(nz).$$

They are sections of degree n^2 for their respective curves. Indeed the ordinary theta pullback has degree n^2, and the extra multiplier from translation is $\exp(2\pi i E_L(nl, y)) = 1$ since $ny \in L_C$ (and similarly $ny' \in L_D$). Their zero divisors are $nz + y \equiv t_C \pmod{L_C}$ and $nz + y' \equiv t_D \pmod{L_D}$. These divisors correspond under γ: the points y, y' correspond because $y = 2n(nx)u$, $y' = 2n(nx)w$. Therefore $\mathcal{T}_{n^2}F_C = cF_D$ for a nonzero constant c.

To determine c, evaluate at $z_0 = -nxu$, $z_0' = -nxw$. These are conjugate points and $nz_0 = -y/2$, $nz_0' = -y'/2$. The exponential factor of the shift is exactly one there. Since θ is even,

$$F_C(z_0) = \theta_C(nz_0), \qquad F_D(z_0') = \theta_D(nz_0').$$

The ordinary pullback of theta is carried exactly to the ordinary pullback of the conjugate theta by the isogeny calculation in Proposition 12.3. Thus the common scale λ^{n^2} in (12.3) cancels, proving $c = 1$ if the value is nonzero. If it vanishes, then $y/2$ and $y'/2$ are the distinguished half-periods modulo their respective lattices. In this case $y \in L_C$, $y' \in L_D$, with classes $2t_C, 2t_D$ modulo twice the lattice. By (12.1) their multipliers are -1. Hence at zero $F_C(0) = F_D(0) = -1$. The normalized scale at zero is one, and again $c = 1$.

Proposition 12.3 makes $R = 1$ in (12.2). Taking Taylor series of the equality of sections and undoing the rescaling by n proves (12.4). Every choice made before fixing x works on the whole finite set. □

This proves `CMLattice.exists_taylorVia_twistedTranslate_theta` in `Lib/CM/MainTheorem/ConjugateCurve.lean`. The special evaluation at the negative half-translate is important: using a value at zero alone can fail if the shift moves the zero divisor there. Also note that the choice of v may depend on the finite set S; it has not been assumed to be a homomorphism of adeles.

Theorem 12.5 (Soft normalized CM reciprocity)**.** *There exist $\Omega \neq 0$ and an integer $e > 0$ with*

$$A_{L_D} = |\Omega|^2\sqrt{d}/2, qquade\Omega \in L_D.$$

For each finite $S \subset K$ one can choose $a_S \in K$ such that

$$\gamma(\mathrm{Tay}_0\, U_{x\Omega_0}\theta_C) = \mathrm{Tay}_0\, U_{x\Omega a_S}\theta_D \quad (x \in S).$$

Moreover the net a_S indexed by finite subsets ordered by inclusion has a cluster point in the finite adeles of K.

Proof. Here use the elementary odd-class-number argument of Chapter 6.2, which makes the square map on the class group bijective. Write $L_D = u\mathfrak{c}^2$ for a nonzero complex number u and a nonzero integral ideal $\mathfrak{c}$ (choose a representative of the square root of the lattice ideal class and absorb the principal factor in u). Set $e = \mathcal{N}\mathfrak{c}$ and $\Omega = eu$. The area of the lattice of an integral ideal of K is its norm times $\sqrt{d}/2$. Thus the claimed area identity follows from multiplicativity of the ideal norm. Since $e \in \mathfrak{c}$, one has $e\Omega = e^2u \in L_D$.

Use Proposition 12.4 for each S and write the resulting lattice element v as $u\beta_S$ with $\beta_S \in \mathfrak{c}^2$. Then $a_S = \beta_S/e$ works. Every a_S belongs to the fixed compact subset $e^{-1}\widehat{\mathcal{O}}_K$ of the finite adeles. The closures of the tails of the net in this compact set have the finite intersection property; their intersection is nonempty. This gives the cluster point. □

The theorem is exactly the soft statement `CMLattice.complex_multiplication_theta` in `Lib/CM/FormalGroup/FormalHeight.lean`. It does not assert an identification of the cluster point with the Artin symbol. Elliptic uniformization and the classical marked-curve comparison are standard external inputs, but the differential, section and scalar normalization are now derived. If a finite-adelic translate is locally constant in its adelic translation variable, its value on the net agrees with its value at the cluster point along a cofinal subnet: for each fixed finite set of torsion tests, the equality already holds on the corresponding tail. This is the passage needed by the normalized generating-vector construction of Chapter 8. It avoids the invalid assertion that arbitrary complex Taylor series are continuous in the finite-adelic topology before the finite Weil realization has been made. The precise comparison of the finite and infinite theta models is still used at that point.

12.3 Removing one string: the Hermitian linear algebra

We now turn to the geometry behind the passage to the definite chain. All the gradings in this section are those of `Grading` in `Lib/Weil/Periods/DegreeAndUnipotent.lean`. Temporarily let E be the CM coefficient field, with fixed field E^+ under the involution. A graded Hermitian space of string length s is

$$V = X \oplus \langle e_0, \dots, e_{s-1} \rangle_E,$$

where X is the positive-definite space with identity matrix. String vectors pair to 1 with their reflected partner e_{s-1-i}, except that the middle vector, if s is odd, has norm $-N$. All other pairings are zero, and the string is orthogonal to X. Here N is the positive split level prime from Chapter 4; the Hermitian form itself is defined over E^+. Let J_s be this matrix, and use the grades

$$\mathrm{gr}(X) = 0, \qquad \mathrm{gr}(e_i) = s - 1 - 2i.$$

The reflection of a string index is written $i^* = s - 1 - i$; reflection on X fixes each chosen basis index. For a matrix M from a graded space to another, write $\deg M \geq a$ if $M_{vu} = 0$ whenever $\mathrm{gr}(v) < \mathrm{gr}(u) + a$.

The adjoint M^* is $J_s^{-1} \bar{M}^t J_{s'}$ and sends a matrix of degree at least a to a matrix of degree at least a in the opposite direction. In fact each nonzero entry is paired with one whose row and column grades are the negatives of the original column and row grades. Products add the lower bounds on degree. These elementary properties apply both to rational matrices and to adelic ones.

Lemma 12.6 (The graded Bessel subgroup)**.** *Let U_s be the unitary group, let U_s^+ consist of the $g \in U_s$ such that $\deg(g-1) \geq 1$, and let B_s consist of those unitary matrices for which*

$$ge_i - e_i \in \bigoplus_{\mathrm{gr}(v) > \mathrm{gr}(e_i)} Ev \quad \text{if } \mathrm{gr}(e_i) \geq 0.$$

Then $B_s = U_s^+ \rtimes U(X)$, and equivalently $g \in B_s$ if $\deg g \geq 0$ and its degree-zero part acts as identity on the string and arbitrarily unitarily on X. All these identities hold also over the adeles and the finite adeles.

Proof. In the defining condition of B_s, preservation of the positive flag, and the identity on its successive lines, are immediate. The annihilators of that flag under the Hermitian form form the complementary negative filtration; unitarity implies their preservation and gives the inverse adjoint action on the negative graded pieces. If s is odd, the grade-zero quotient is $X \oplus Ee_{(s-1)/2}$. The middle string vector is fixed; its norm is nonzero, so the orthogonal complement X is preserved in that quotient. If s is even, the grade-zero quotient is simply X. Altogether this proves the description of the degree-zero part, and multiplication identifies it with a unitary matrix x acting on X. Then $gx^{-1} \in U_s^+$. Conversely any product of the stated kind satisfies the defining condition. The arguments on filtrations apply over each completion, and to the restricted product because all the relevant matrix identities preserve integrality at almost all places. □

Both groups are closed. The rational points of B_s are cocompact in its adelic group: for U_s^+ this follows by the positive-degree filtration, whose successive groups are additive vector groups over E^+, and the compactness of $E^+ \backslash \mathbb{A}_{E^+}$. One can choose a compact set of representatives successively using the filtration and lifts of compact subsets of the vector groups. The quotient for $U(X)$ is compact by anisotropy and the finiteness of its class set, as in Chapter 5. The combination gives a compact set of representatives for B_s. In the formal source

this is `Grading.isCocompact_ratBessel` in `Lib/Weil/Periods/BesselIntegral.lean`; no assertion of compactness of the full quotient of the indefinite group is made.

Write $V' = X' \oplus \langle e'_0, \dots, e'_{s-2} \rangle$ for the next space, with the same N and with parameter $c \in (E^+)^\times$ (for our chain $c = -(jg\varpi)^{-1}$). Define $T : V \to V'$ to kill X and the final string vector, and to send e_i to $t_i e'_i$, where, if $0 \le i < s-1$,

$$t_i = \begin{cases} 1, & s \text{ odd and } 2i+1 < s, \\ -Nc, & s \text{ odd and } 2i+1 = s, \\ c, & s \text{ odd and } 2i+1 > s, \\ c, & s \text{ even and } 2i+1 < s, \\ 1, & s \text{ even and } 2i+1 > s. \end{cases} \tag{12.5}$$

This is `Grading.stringMap` in `Lib/Weil/Periods/FiniteBesselGroup.lean`. In particular T is homogeneous of degree -1. Let S_s send each e_i to ce_{i+1}, with extra factor $-N$ at degree zero and $-Nc$ at degree one, and kill X and e_{s-1}. It has degree -2.

Lemma 12.7 (The two moment-map identities)**.** *For the string maps in* (12.5)*,*

$$T^*T = S_s, \qquad TT^* = S_{s-1}.$$

If $g \in B_s$ and $h \in B_{s-1}$, then $\deg(hTg - T) \ge 0$.

Proof. Put $h_i = -N$ at the middle string index and $h_i = 1$ otherwise; similarly put h'_i for the target. Directly from the Gram matrices, the adjoint sends e'_i to

$$\bar{t}_{s-2-i} \frac{h'_{s-2-i}}{h_{s-2-i}} e_{i+1}.$$

Consequently the coefficient of e_{i+1} in T^*Te_i is $t_i t_{s-2-i} h'_{s-2-i}/h_{s-2-i}$, because each t_i is fixed under conjugation. If neither of the paired indices is a middle index, it equals c. In the odd case an exceptional factor $-N$ at degree zero occurs, and in the even case the target middle index gives $-Nc^2$ at degree one. These are exactly the indicated coefficients; the same calculation on V' gives TT^*. Finally, the degree-zero parts of g, h are identity on string coordinates, and T kills all other coordinates. In expanding $hTg - T$, each surviving term therefore contains a positive-degree factor from h or g, and hence has degree at least zero. □

For example, in two consecutive ranks with string length 3 and 2, T sends $e_0 \mapsto e'_0$, $e_1 \mapsto -Nce'_1$, $e_2 \mapsto 0$. Its adjoint sends $e'_0 \mapsto ce_1$, $e'_1 \mapsto e_2$. Composing gives $ce_1, -Nce_2, 0$ on the old string, and $-Nc^2 e'_1, 0$ on the new one. These are not the moment-map identities for a string with the middle Gram entry 1. It is the middle factor N that will reappear in the local lattice restriction.

12.4 The surviving rational orbit

For maps $w : V \to V'$ consider the Heisenberg bicharacter in rational coordinates

$$\beta(w, z) = \psi\left(\tfrac{j}{2} \operatorname{tr}(z^* w)\right)$$

with the additive character and purely imaginary j used in the global Weil model. We suppress the trace to the rational adeles inside ψ. For the purposes of this calculation β has the usual skew perfect pairing after taking the ratio with its reverse. Its precise half-factor is important

in relating the phase of the two moment maps. Let d be target weights with $d(v^*) = -d(v)$. The survival condition at weights d is:

$$u \in U_s^+, \quad \deg(u-1) \geq 2, \quad \deg_d(w(u-1)) \geq 1 \quad \Longrightarrow \quad \beta(Tu, T) = \beta(wu, w). \tag{12.6}$$

Here $\deg_d$ uses the old source grades and target weights d. This is exactly `Grading.Survives` in `Lib/Weil/Periods/UnfoldingStep.lean`. We prove the linear algebra responsible for its consequences. It is useful to do so separately from the analytic Fourier expansion.

Lemma 12.8 (An Eichler test)**.** *Suppose w satisfies (12.6). Fix an old string index a and let $k = e_{a^*}$. Suppose that a vector v has support in positive grades at least $\max(1, \mathrm{gr}(a) + 2)$, that wv has support in target weights $> \mathrm{gr}(a)$, and that for every nonzero component of wk and every nonzero component of v the target weight plus source grade is positive. Then*

$$(T^*Tv)_a = (w^*wv)_a.$$

The same statement holds for adelic w, v under the indicated support assumptions.

Proof. Both v and k are isotropic: v lies in the positive-grade string and k is isotropic unless a is the middle string index. If a is the middle string index, k is not isotropic; this case is handled by the correction below. Moreover $\langle k, v\rangle = 0$ because $v_a = 0$. For any adelic scalar z, let $k_z = zk$ and put

$$M = vk_z^* - k_z v^*, \qquad u = 1 + M - \tfrac{1}{2}\langle k_z, k_z\rangle vv^*.$$

Here vectors are read as column maps and the star is the Hermitian adjoint. Direct multiplication gives $M^* = -M$, $M^2 = -\langle k_z, k_z\rangle vv^*$, $M^3 = 0$, and hence $u^*u = 1$. (This convention for the sign follows by writing $u = 1 + M + M^2/2$.) The degree of M is at least two, and that of M^2 at least four. The two assumed support properties show that wM has degree at least one relative to d; the term wM^2 has that degree as well by multiplication by M. Thus (12.6) applies.

If k is isotropic, the quadratic term is absent. If not, apply the character identity separately to z and $-z$, and divide; the term quadratic in z cancels. In either case the identity for all z implies that the difference $\langle wk, wv\rangle - \langle Tk, Tv\rangle$ vanishes: the linear character in z is trivial for every adelic z only if its coefficient is zero, by the perfect adelic additive pairing. More explicitly the trace of $(w^*w - T^*T)(vk_z^* - k_z v^*)$ is the difference of $\bar{z}\langle k, (w^*w - T^*T)v\rangle$ and its conjugate, and multiplication by the nonzero purely imaginary j detects both components when z varies. The entry in the reflected position is paired by the nonzero Gram value with the coordinate at a. This proves the claim. □

The rank-one test is the mechanism of `Grading.eq_of_survives`. The correction in the formula for u is necessary if the reflected basis vector has nonzero norm. Our applications include precisely this possibility.

Define $\mathrm{Agree}(w, T)$ to mean that $w_{ba} = T_{ba}$ whenever the source grade $\mathrm{gr}(a)$ is positive and the target grade $\mathrm{gr}(b)$ is smaller. Notice that w need not agree with T at positive or zero relative degree, nor at all entries in a negative source column.

Proposition 12.9 (Rational normalization of the surviving orbit)**.** *Let w be a rational map $V \to V'$ satisfying (12.6) for the zero target weight. Then a rational unitary matrix h on V' makes hw agree with T.*

Proof. Write V_+ for the strictly positive-grade string subspace. First $w|_{V_+}$ is injective. If not, take a nonzero $v \in V_+$ killed by w, and let i be its largest string index with nonzero coefficient. Set $a = i + 1$; this string index still exists since e_i has positive grade. The support of v lies at grade at least $\mathrm{gr}(a) + 2$, and the zero target-weight assumptions of Lemma 12.8 hold (the condition on wv is vacuous, while all components of v have positive grade). By Lemmas 12.8 and 12.7, the a-component of T^*Tv is zero. But it is the nonzero coefficient of the shift of v_i. This is a contradiction.

For two positive string indices i, l, apply the lemma to $v = e_i$, $a = l^*$. The target-weight condition on wv is automatic since $\mathrm{gr}(a) < 0$. The support condition on wk follows as all target weights are zero and v has positive grade. We obtain

$$\langle we_l, we_i \rangle = \langle Te_l, Te_i \rangle.$$

The restrictions of w and T are injective (the latter is clear from its nonzero string coefficients), so they give an isometry between the two subspaces of V'. The classical Witt extension theorem over the field E, valid also for degenerate subspaces of a nondegenerate ambient Hermitian space in characteristic zero, extends it to a rational unitary matrix. In particular all the positive string columns can be made equal to those of T, which is stronger than agreement. □

Lemma 12.10 (Stabilizer of agreement)**.** *Suppose w agrees with T, $g \in B_s$. Then hwg agrees with T if and only if $h \in B_{s-1}$.*

Proof. On a positive string column e_i, g sends it to itself plus higher positive string columns. Thus wge_i is $t_i e'_i$ plus a vector of grade at least $\mathrm{gr}(e_i)$: this uses (12.5), whose grade drop is one. If $h \in B_{s-1}$, the degree property of Lemma 12.6 immediately gives agreement. Conversely start from the highest nonnegative target string grade, and proceed downwards over the nonnegative target string indices k. Each e'_k corresponds to a positive source string index k, with grade one greater, and $t_k \neq 0$. By induction, h already fixes every higher target string column modulo still higher grades. Applying the agreement of hwg to the source column corresponding to k, the higher terms therefore cannot contribute at grades at most $\mathrm{gr}(e'_k)$. The coefficient t_k is nonzero, so $he'_k - e'_k$ is supported in grades strictly greater than $\mathrm{gr}(e'_k)$. This is exactly the defining condition of B_{s-1}. □

Proposition 12.11 (The final surviving class)**.** *If w agrees with T and satisfies* (12.6) *for the actual target grades $d = \mathrm{gr}'$, then $w - T$ has degree at least zero.*

Proof. All negative-degree entries in positive source columns already agree. For the remaining entries prove by induction on increasing target grade $r < 0$ that $w_{ba} = T_{ba}$ for source grade $\mathrm{gr}(a) > r$ and target string row b of grade r. (There are no other rows of negative grade.) Use the target reflected row b^* of grade $-r$ and let the source string vector v be the one with grade $-r + 1$ sent by T to te'_{b^*}, $t \neq 0$. By agreement, $wv = te'_{b^*}$ plus terms of grade greater than $-r$.

Fix a source column a with $r < \mathrm{gr}(a) \leq 0$. If it is a string column, put $k = e_a$, and in the Eichler test use the reflected index a^*. If it is in X, use the same calculation with the corresponding norm-one basis vector $k = e_a$. The latter vector also has grade zero, so the proof of the lemma with its quadratic correction applies unchanged. The vector v has grade $-r + 1 \geq -\mathrm{gr}(a) + 2$; every component of wv has target grade at least $-r > -\mathrm{gr}(a)$. By the induction hypothesis, all entries of wk in grades less than r equal those of Tk. Such entries of Tk vanish, since T either kills the column or drops its grade by one. Thus every nonzero component of wk has grade at least r, and its grade plus that of v is at least one. The Eichler test gives $\langle wk, wv \rangle = \langle Tk, Tv \rangle$. Terms of wv above grade $-r$ pair only with grades below r,

where wk vanishes. Thus the equality of pairings forces equality in row b, because t and the paired Hermitian Gram entry are nonzero. This completes the induction. □

These propositions derive the rational-orbit part of the one-step unfolding. Their source anchors are `Grading.exists_agreesWithT`, `Grading.agreesWithT_mul_iff`, and `Grading.hasDegreeGe_of_survives`, all in `Lib/Weil/Periods/UnfoldingStep.lean`. The two uses of survival, at zero weights and at true weights, are different: applying only the first gives too many matrix entries and too large a stabilizer.

12.5 The one-step period transfer

We explain the analytic argument in sufficient detail to state precisely which Weil realization it uses. A rational symplectic vector space in the construction is $\mathrm{Hom}(V, V')$, with alternating form given by the trace and j; it is paired by the reflection of both indices. The maps of strictly positive relative degree form an isotropic subspace. Its annihilator is the subspace of nonnegative degree, because each coordinate of relative degree a pairs with a coordinate of degree $-a$. This continues to hold for target weights d which are odd under reflection. Choose a rational symplectic splitting into these positive coordinates, their dual negative coordinates, and the degree-zero part.

In the resulting mixed Schrödinger model, write $F(x, z)$ for a Schwartz function with x in the negative-coordinate space and z in a remaining Lagrangian of the degree-zero part. The mixed-model theta functional uses the rational sum over both sets of coordinates. The constant term at the rational matrix w, denoted $C_d(f; w)$, is the integral of the theta functional of the translated datum over the adelic positive-coordinate group modulo its rational points. Fourier expansion gives

$$\Theta(f) = \sum_{[w]} C_d(f; w), \tag{12.7}$$

where $[w]$ runs over rational maps modulo nonnegative-degree maps. To verify this, first use $w = 0$: integrating the mixed-model theta sum in the positive variables kills every nonzero dual rational frequency. At general w, its negative-coordinate part shifts that frequency; the coset label is precisely the list of negative coordinates. The degree-zero translation acts by the rational Heisenberg operator and does not change the theta functional, and the nonnegative positive translation changes the phase trivially on a rational w. This proves that $C_d(f; w)$ is constant on the indicated rational coset and establishes (12.7).

If $w - w'$ has strictly positive relative degree (now not necessarily rational), the Heisenberg multiplication law and a change of variable in the compact period integral give

$$C_d(f; w') = \beta(w, w') C_d(f; w). \tag{12.8}$$

For rational elements $g \in B_s$, h in the target parabolic preserving the weight filtration, conjugation of the positive-coordinate integration variables by their action on Hom gives

$$C_d(\omega(g, h) f; w) = C_d(f; h^{-1} w g). \tag{12.9}$$

The quotient measures on additive groups here have total mass one; hence no determinant Jacobian is left by rational substitution. The splitting characters of the rank twists are trivial on rational determinants. The identical formula for adelic grade-raising unipotents on the two sides follows from the mixed-model action: their determinant on the isotropic subspace is one and their remaining diagonal action is trivial. These are the standard mixed-model identities, corresponding to `Grading.constTerm_congr_rat`, `Grading.constTerm_eq_homBi`

`character_mul`, and `Grading.constTerm_mul` in `Lib/Weil/Periods/UnfoldingStep.lean`. They express the needed normalization of the Schrödinger model, not a statement of analytic continuation of a theta series. The finite–archimedean identification of the particular datum of our chain with this realization is a separate matter.

One can also take partial Fourier expansions using different target weights. In particular take zero target weights and actual target grades. In either expansion retain only the terms agreeing with T in the earlier sense. The fixed negative-coordinate positions are precisely those satisfying simultaneously

$$0 < \mathrm{gr}(a), \qquad \mathrm{gr}'(b) < \mathrm{gr}(a).$$

Those positions occur in both sets of negative coordinates. Their dual positive-coordinate positions form a common isotropic subgroup of the two choices. Summing over all other negative coordinates in either model is the same as averaging the theta functional with the character prescribed by T only on this common isotropic subgroup, by ordinary Fourier inversion on the compact rational quotient. This proves that the two retained partial sums agree; the observation is the content of `Grading.tsum_constTerm_agreesWithT`. It is often a useful shortcut to avoid choosing a model in which all Fourier coordinates are written out at once.

Let P be a continuous function invariant under $B_s(E^+)$ and transforming under degree-at-least-two unipotents by

$$P(ug) = \beta(Tu, T)P(g) = \psi\left(\tfrac{j}{2}\,\mathrm{tr}(S_s u)\right) P(g).$$

It is called generic here only relative to this specific string and character. The theta kernel is evaluated on a fixed datum Φ by the twisted Weil action $\omega(g, h)$. Define

$$P'_x(h) = \int_{[B_s]} P(rx)\, C_{\mathrm{gr}'}(\omega(rx, h)\Phi; T)\, dr, \qquad [B_s] = B_s(E^+)\backslash B_s(\mathbb{A}_{E^+}). \tag{12.10}$$

This is the function `Grading.unfoldStep`.

Theorem 12.12 (One step of the global period unfolding)**.** *In the normalized mixed model just described, let y be a continuous rapidly decreasing automorphic function on $U(V')$. With the compatible quotient measures of the theta integral and subgroup periods, the identity is*

$$\begin{aligned}\int_{[B_s]} P(rx) \left(\int_{[U(V')]} \Theta(\omega(rx, h)\Phi) y(h)\, dh \right) dr \\ = \int_{B_{s-1}(E^+)\backslash U(V')(\mathbb{A}_{E^+})} P'_x(h) y(h)\, dh.\end{aligned} \tag{12.11}$$

The unfolded function is generic for the next string and its own shift.

Proof. Use the expansion (12.7) first at zero target weights. Partition its rational labels into orbits of $B_s(E^+)$. For a particular label w with failure of the survival condition choose a degree-at-least-two unipotent u as in (12.6) for which the ratio of the two characters is not one. The integral of that label, after unfolding its rational orbit, is an integral on the quotient by its rational stabilizer. Left translation by u multiplies the integrand by

$$\beta(Tu, T)/\beta(wu, w)$$

by genericity and (12.8)–(12.9). One must justify this step, since left multiplication by an arbitrary adelic element need not be a map on a rational stabilizer quotient. The following

general argument does so. For a continuous, nonzero, left-Γ-invariant integrand F on a locally compact group G, let H consist of the elements h for which $F(hz) = \chi(h)F(z)$ for some $|\chi(h)| = 1$ and all z. By evaluating at a point where F is nonzero, H is seen to be a closed subgroup with continuous unitary character χ; it contains Γ. The orbit quotient integral is absolutely integrable. Disintegrate it in stages, first on $\Gamma\backslash H$ with left Haar measure, then on $H\backslash G$ with the compatible quotient density. On each fibre the map $h \mapsto hu$ is well defined on $\Gamma\backslash H$ and scales its measure by a positive real modular factor. But it scales the integrand by $\chi(u)$ because $F(hux) = \chi(h)\chi(u)F(x)$. As $\chi(u) \neq 1$ and has absolute value one, it cannot equal that positive modular factor. The fibre integral is zero, hence so is the total integral. This proves the required vanishing without a normalizer assumption on u. It is the quotient-measure argument of `integral_quotMeasure_eq_zero_of_twist` in `Lib/Adelic/QuotientMeasure/CocompactIntegrable.lean`.

By Proposition 12.9, the remaining rational labels are carried to the set of labels agreeing with T by rational matrices of the target unitary group. Lemma 12.10 identifies the stabilizer within that group as the rational Bessel subgroup. Unfold the rational cosets of the target automorphic integral: a sum over translates indexed by the rational Bessel subgroup quotient of the rational target group converts the integral on $[U(V')]$ to the integral on $B_{s-1}(E^+)\backslash U(V')(\mathbb{A}_{E^+})$. After this unfolding it remains to sum the agreed labels at zero target weights. Replace the partial Fourier sum by the equal partial sum at actual target grades, as explained above. The same unipotent character argument applies to that partial sum. Proposition 12.11 shows that its sole surviving class is T modulo the nonnegative-degree subgroup. This is exactly the integrand of (12.10).

The invariance of P'_x under rational B_{s-1} follows from (12.9) and Lemma 12.7. For a degree-at-least-two unipotent v of the target, $vT - T$ has positive relative degree. The phase of (12.8) for that change is the character of the next shift: the identity $TT^* = S_{s-1}$ and cyclicity of trace give $\psi(j\,\mathrm{tr}(S_{s-1}v)/2)$. This is next-stage genericity.

The integrals and sums involved are absolutely convergent under the usual Schwartz-theta growth estimates and rapid decay stated here. Indeed a partial Fourier transform takes Schwartz functions to Schwartz functions, and their absolute rational sums converge uniformly for data in bounded sets of Schwartz seminorms, since the finite support restricts to a lattice of bounded denominators and archimedean Schwartz decay beats its polynomial counting bound. One may use a compact set of representatives for $[B_s]$ when bounding the factor depending on r. On the remaining quotient, the rational theta sums, and the sum of absolute values of their partial Fourier coefficients, are bounded by a power of the adelic reduction height of h, by the standard moderate-growth estimate for the Weil representation on Schwartz seminorms. Multiplication by a rapidly decreasing y makes this integrable using reduction theory for unitary groups. These are the same analytic hypotheses as the single-step result of B. Wang Peng Jun, *Global theta lifting and automorphic periods associated to nilpotent orbits*, J. Number Theory 271 (2025), 122–149, Proposition 3.2 and Theorem 4.5, in its unitary hook setting. The proof uses this classical moderate-growth input rather than any quantitative bound in the varying parameters of GRH. The estimates justify the changes of order made above. For a purely algebraic finite-place test function the corresponding summations are finite already in the coefficient variables. □

This is the structure of `Grading.integral_besselGroup_unfold` and `isGeneric_unfoldStep`. The geometric orbit computation, including the Witt step and the single surviving class, has been reconstructed above. The theorem uses the ordinary normalized mixed-model identities and its qualitative moderate-growth theorem. The quantitative uniformity in the upper bound will instead come from the integral coset cuts below. In particular, the one-step period identity alone does not imply the iterated seesaw with its good finite datum.

12.6 The glued lattice and integral Witt extension

Here is the algebra of the lattices used in that quantitative step. Write $\widehat{O} = \prod_w O_{E,w}$ for the integral finite adeles. A finite adelic lattice of full rank is a matrix translate of $\widehat{O}^n$ in rational coordinates; its dual is taken with respect to the given Hermitian form. The field in the construction contains a primitive 5040th root of unity, the field hypothesis named `StablyUnram`. In particular the quadratic CM involution is unramified or split locally: if it were a nontrivial inertia element at a place above a rational prime p, it would be trivial on roots of unity of order prime to p. One can choose an odd prime dividing 5040 and different from p, whose primitive root is not fixed by complex conjugation. This contradiction proves the assertion. Norm-one units and matrices should nevertheless be treated separately at a split pair of places; the involution there exchanges the two factors rather than acting within either one.

Choose a finite idele π with $\pi\bar{\pi} = N$, supported with valuation one at precisely one place over each split pair above N. For a finite idele a and a unitary idele u define the basis $B_s(a, \pi, u)$ as follows. If s is even it is the identity on X and multiplies the first half of the string by u/a and the second half by $u\bar{a}$. If s is odd it multiplies every string vector by u, and, writing e for the last vector in X and z for the middle string vector, replaces e by

$$\bar{\pi}^{-1}e + uz/N.$$

It does not change the other X-vectors. This is the matrix `gradedLatticeBasis` in `New/Seesaw/LatticeAndTorsionDefs.lean`.

Lemma 12.13 (The glued basis)**.** *The Gram matrix of this basis is integral and unimodular. Its string diagonal entries in the upper triangular flag order are u, for odd s, and $u/a, u\bar{a}$ on the respective halves for even s.*

Proof. In the even case every pair of reflected string vectors has pairing $\overline{u/a}u\bar{a} = 1$, just as before. In the odd case the only changed Gram block is the pair consisting of the new X-vector and the new middle string vector. Their Gram matrix is

$$\begin{pmatrix} 0 & -1 \\ -1 & -N \end{pmatrix}.$$

Indeed the new X-vector has norm $N^{-1} - N/N^2 = 0$, and its pairing with uz is -1. This block has determinant -1 and integral inverse. All other Gram blocks are unchanged. The described diagonal entries follow directly from the basis. For the flag order put all the positive strings first, then the middle string if present, then the X-block, then the negative strings; the added middle entry in the last X-column lies earlier in the flag order. □

At a nonsplit unramified place a unimodular Hermitian lattice admits an orthonormal basis. To see this, some diagonal value of its reduction is nonzero, or a pair of off-diagonal entries is nonzero and an appropriate linear combination of the two basis vectors has unit norm, by surjectivity of the residue trace. Split off a unit line and induct. Rescale the resulting unit diagonal entries by norms of local units, using the surjectivity of the norm of an unramified quadratic extension on units. At a split pair the two factors are in perfect duality, so choose a basis on one and the dual basis on the other. Taking these choices at all places gives a finite adelic frame x_f with $x_f^* J_s x_f = 1$ spanning the glued lattice. At an archimedean place the signature of J_s has exactly $\lfloor (s+1)/2 \rfloor$ negative directions. The reflected pairs are diagonalized by sum and difference divided by $\sqrt{2}$, and the middle vector, if present, is rescaled by $N^{-1/2}$. This gives a real frame of the prescribed signs; for $N = 1$ its matrix in the original

coordinates can be chosen unitary for the Euclidean inner product. These are the existence statements of `isUnimodular_gradedLatticeBasis_gram` and `exists_frame_gradedGram` in `New/Seesaw/FramesForY.lean`. They apply to the intermediate spaces as well as the endpoints.

Lemma 12.14 (An integral Hermitian extension lemma)**.** *Suppose F and F' are n by r matrices over $\widehat{O}$ whose columns admit integral left inverses, and $F^*F = (F')^*F'$ for the standard finite Hermitian form. There exists an integral unitary h with $hF = F'$.*

Proof. We give a local proof, since it clarifies the need for the unramified condition at residue characteristic two. In the nonsplit case the residual extension is quadratic over finite fields. The reductions of both frames are injective and define an isometry between their subspaces, including possibly degenerate subspaces. The Hermitian Witt extension theorem over finite fields gives a residual unitary carrying one to the other. It is enough to lift it to an exact unitary with the same prescribed frame action.

Use a uniformizer ϖ_+ of the fixed field, also a uniformizer of the extension. Suppose a matrix h_n has $h_nF \equiv F'$ and $h_n^*h_n \equiv 1 \pmod{\varpi_+^n}$. Let $\Delta = (F' - h_nF)/\varpi_+^n$ and $D = (1 - h_n^*h_n)/\varpi_+^n$, read modulo the uniformizer. We seek $h_{n+1} = h_n(1 + \varpi_+^n A)$ with

$$A\bar{F} = \bar{h}_n^{-1}\bar{\Delta}, qquad A + A^* = \bar{D}.$$

The compatibility condition is

$$\bar{F}^*\bar{D}\bar{F} = \bar{F}^*(\bar{h}_n^{-1}\bar{\Delta}) + (\bar{h}_n^{-1}\bar{\Delta})^*\bar{F},$$

which follows by expanding equality of the two Gram matrices mod ϖ_+^{n+1}. Pick a residual scalar τ with $\tau + \bar{\tau} = 1$ by surjectivity of the trace (this is valid even in characteristic two). Then $A_0 = \tau\bar{D}$ has the required symmetrization. The remaining prescribed action of $A - A_0$ on the column space of $\bar{F}$ is skew-adjoint there by the displayed compatibility. Such a partial skew-adjoint map extends to the entire ambient space: choose any linear complement to the column space and define the corresponding sesquilinear form on pairs of column-space and complementary vectors by the prescribed map and skew symmetry; on the two complementary variables take it to be zero. The nondegenerate ambient form identifies this extended form with a skew-adjoint endomorphism. Thus an A exists. Lift its entries arbitrarily and iterate; completeness gives the exact unitary matching the two frames.

At a split pair the module is a pair of free modules R^n, R^n on which the involution exchanges factors. The residual action is an arbitrary invertible g on one and g^{-t} on the other. With reduced frames $F = (P, Q)$ and $F' = (P', Q')$, equality of Gram means $Q^tP = (Q')^tP'$. We need $gP = P'$ and $(Q')^tg = Q^t$. Here P, P' are injections and $Q^t, (Q')^t$ are surjections. The prescribed map between the images of P, P' identifies their intersections with the kernels of the respective surjections. Extend that identification to an isomorphism between the two full kernels; choose representatives of the remaining quotient vectors compatible with the common image of Q^tP in the quotient. This gives an invertible g with both properties. The same successive lifting argument works over the split ring, now using trace-one element $(1, 0)$ and complementary bases in the two factors (the linear compatibility is the equality of the prescribed paired values). Alternatively one can make the kernel and quotient basis choices integrally by taking primitive frames and matching the common elementary divisors of Q^tP. This proves the local assertion at every place. All the resulting h are integral, so their product is a finite-adelic unitary element. □

In this proof the residual use of the Hermitian field Witt theorem is elementary classical linear algebra. The lifting of the frames is essential; reducing a rational-unitary Witt matrix

modulo a prime would not prove the required existence of an integral one. This is a conventional proof of `witt_Ohat_of_orthonormal` in its more general primitive form the integral Witt lemma for $\widehat{O}$ in `Lib/Adelic/Witt/SpaceOperations.lean`; we will refer to the latter also as the integral Witt extension lemma.

There is a further useful congruence inclusion. Embed $U(X)$ by fixing all string vectors. If s is even, every integral unitary on X preserves the glued lattice. In the odd case suppose the last row of the matrix h on X is congruent to the last row of the identity modulo π. The inverse satisfies the same row congruence (multiply $(h-1)$ on the right by h^{-1}), and by unitarity this is the congruence of the last column of h modulo $\bar{\pi}$. In the odd glued basis the X-part has the last coordinate scaled by $\bar{\pi}^{-1}$; the Hermitian conjugate row is thus scaled by π^{-1}. For the endomorphism $h-1$ on V, all the entries of

$$B_s^* J_s (h-1) B_s$$

are then integral: the only one in which both denominators occur is the last diagonal X-entry, divisible by $\pi\bar{\pi} = N$ since these ideals are coprime at each local place. The contributions involving the extra string coordinate vanish as $h-1$ kills the string. Because $B_s^* J_s B_s$ is unimodular, $B_s^{-1} h B_s$ and its inverse are integral. This proves the inclusion of the subgroup fixing the last row modulo π into the lattice stabilizer. Restricting further at the chosen primes gives the subgroup `Data.columnLevel`. It proves the assertion of `Grading.inclXFin_mem_Kf` in `New/Seesaw/StandardLatticeLevel.lean`.

12.7 Adapted bases and equality of lattice lines

Use the flag order used above, grouping all indices of X into a single block. If $L_x = x_f x \widehat{O}^n$ is the lattice associated to $x \in U(V)(\mathbb{A}_f)$ in framed coordinates, an adapted basis B is a matrix of generators upper block triangular in this flag order. Its string diagonal entries are finite ideles β_i. Such a basis exists: locally intersect the lattice with every subspace of the ordered flag. These intersections are saturated in the lattice, hence the successive quotients are free over the discrete valuation ring. Choose and lift bases of these quotients; this gives an integral basis adapted to the flag. At almost all places the lattice is already standard, so one obtains a compatible finite-adelic basis. In particular $B^{-1} x_f x$ is integral invertible. The valuation of β_i is independent of the adapted basis, since the change of bases is triangular and integral invertible, and string diagonal entries multiply under triangular products. More intrinsically it is the valuation of the rank-one lattice induced on the ith string graded piece. If $L_{x'} \subset L_x$ then $\operatorname{ord}_w \beta_i' \geq \operatorname{ord}_w \beta_i$ (inclusions induce inclusions on each flag quotient); for equality of local lattices the valuations agree. The source uses ideleLogVal $= -\operatorname{ord}_w$.

Let $h_i = 1$ except for the middle string entry of odd length, where $h_i = -N$. The following formula controls the Gram of the lifted flag:

$$(B^* J_s B)_{ij} = h_i \bar{\beta}_i \beta_j \, [j = s-1-i], \qquad j \leq s-1-i \tag{12.12}$$

for string indices i, j. Indeed in the product the ith column is supported only in string indices $\leq i$ and possibly in X if i lies later in the flag. If $i + j < s-1$ no paired string indices can contribute; if $i + j = s-1$ only the diagonal pair contributes. At least one of the two indices in the displayed range lies in the nonnegative-grade string and therefore has no X-component. This proves the formula even when the other string vector lies after the X-block.

Proposition 12.15 (Bessel double cosets determined by the nonnegative lines)**.** *If two such lattices $L_x, L_{x'}$ have equal string diagonal valuations at all places on the nonnegative-grade string, then $x' k x^{-1}$ belongs to the finite Bessel subgroup for some framed integral unitary k.*

Proof. Rescale the corresponding adapted basis columns of B' by integral units to make their diagonal ideles exactly those of B. This is possible because equality of valuations means the ratios of diagonals are integral idelic units. The matrices $F = (x_f x)^{-1}B$ and $F' = (x_f x')^{-1}B'$ are integral invertible. Retain their columns from the nonnegative string. The resulting two frames have integral left inverses and their Gram matrices coincide by (12.12), since the finite frame identifies J_s with 1 and x, x' are unitary.

Use Lemma 12.14 to send one frame to the other by an integral unitary k. The matrix of $x'kx^{-1}$ in rational coordinates consequently sends every retained column of B to the corresponding column of B'. By the flag-triangular form with the same diagonal, it fixes each of these standard basis vectors modulo higher string grades. One sees this inductively along the columns: earlier string columns have already been treated and the coefficient of the current line is the same. Lemma 12.6 identifies the resulting element as a member of the Bessel subgroup. □

The source anchors are `Grading.exists_isAdaptedAt`, `Grading.IsAdaptedAt.gram_string_entry`, `Grading.IsAdaptedBasis.ideleLogVal_le`, and `Grading.IsAdaptedAt.exists_mem_besselFin` in `New/Seesaw/EqualLinesOrbit.lean`. This proposition is stronger than a mere compactness criterion: it gives a representative for the double coset from just finitely many string-line valuations, not from all the off-diagonal entries of an arbitrary basis.

12.8 Local supports of zero terms

We recall the finite-coordinate normalization used in the support estimate. The product finite lattice $O^n = \widehat{O}^n$ is self-dual for the additive pairing $\psi_0(\bar{x}^t y)$ of the Heisenberg representation in its standard finite coordinates. This means that if $\psi_0(\bar{x}^t y) = 1$ for every integral x, then $y \in O^n$; products over places reduce the assertion to self-duality for the local additive character of conductor O_w. The frame of each graded spaces is related to rational coordinates by

$$w \longmapsto \alpha^{-1} x_f'^{-1} w x_f$$

on the matrix Hom space (in vector notation). The finite idele α incorporates the different and character normalizations; see Chapter 6. For the following support lemma all finite vectors are already in these normalized coordinates.

Let L_0 denote the finite part of a zero-term map along an isotropic frame k in the source space of the finite Heisenberg representation. We use its standard mixed-model property $L_0\rho(kz) = L_0$ for all adelic z. Denote by M^* the conjugate transpose in finite coordinates.

Lemma 12.16 (A support criterion from finite translation invariance)**.** *Let M, M' be matrices with $M(M')^* = 1$, and let the finite datum f be fixed by $\rho(Mx)$ for all $x \in O^n$. If $L_0\rho(y)f \neq 0$, then there is $\ell \in O^n$ with*

$$k^*(y - M'\ell) = 0.$$

Proof. The matrix $(M')^*k$ is a full-rank matrix over the finite adeles. Choose an invertible column matrix D such that $X = (M')^*kD$ is an integral primitive frame. Locally this is just a basis of the saturated lattice in the column space, and almost everywhere the given frame is already integral. For any integral vector z, one has $MXz = kDz$. The operator $\rho(kDz)$ fixes f by the assumed invariance, and disappears after applying L_0. Commuting it past $\rho(y)$, and using the finite Heisenberg relation, gives

$$L_0\rho(y)f = \psi_0((kDz)^*y)L_0\rho(y)f.$$

Since the displayed vector is nonzero, the phase is one for every $z \in O^r$. But $kD = MX$, so this says $X^*M^*y \in O^r$ by self-duality. The primitive frame X has an integral left inverse; hence its adjoint maps O^n onto O^r. Choose $\ell \in O^n$ with $X^*\ell = X^*M^*y$. Since $X^* = D^*k^*M'$ and $M'M^* = 1$, the equality is $D^*k^*(y - M'\ell) = 0$. The matrix D is invertible, giving the claim. □

This is exactly `zeroTermFin_support` in `New/Seesaw/EdgePair.lean`. A zero term of a finite datum need not vanish whenever its translated argument is nonintegral: only its coordinates dual to the isotropic frame are tested. Replacing this criterion by the stronger and false full integrality condition would destroy the sharp estimates.

There is also a zero-term formula for the vacuum. Write i for the finite embedding of the degree-zero part, and let $B = ig + kc$ be an integral matrix, where g is unitary on the degree-zero part and c is arbitrary. The vacuum v_0 on this part is determined up to scalar by invariance under the standard integral finite shifts (one-dimensionality of its fixed line). The vector $\omega(g)^{-1}L_0v_{\mathrm{vac}}$ is on that line, by the mixed-model equivariance and the integrality of B: translating by $iz + kcg^{-1}z$ has the same zero-term action as translating by iz, and the first vector fixes the incoming vacuum when z is integral in the transformed coordinates. Hence $L_0v_{\mathrm{vac}} = C\omega(g)v_0$ for some scalar C. For a general translate $\rho(y)v_{\mathrm{vac}}$ the support lemma with $M = M' = 1$ supplies $\ell \in O^n$ with $y - \ell$ orthogonal to k if the result is nonzero. Decompose this difference as $iw + kb$ in the mixed-model splitting. Its zero term is the translate of L_0v_{vac} by w. Because $B^*(y - \ell) = g^*w$, and $B^*\ell$ is integral, the finite shifts and their commutation phases now give

$$L_0\rho(y)v_{\mathrm{vac}} = C\zeta\,\omega(g)\rho(B^*y)v_0, \qquad \zeta = 0 \text{ or a root of unity.} \tag{12.13}$$

The root of unity is the product of a finite number of values of the finite additive character and semicharacter. In particular its magnitude after any algebraic conjugation is one when nonzero, and it is integral at all finite places. The precise source is `zeroTermFin_finHeisOp_vacElem` in `New/Seesaw/EdgePair.lean`. This proof uses the fixed vacuum line and the elementary mixed-model covariance, but not a computation of the real vector in the zero term.

12.9 A one-step bound for the lattice lines

Let the old adapted basis have string diagonals β_i and the next one β'_i. Let Q be the chosen split places of the twist conductor (one from each conjugate pair occurring), and write

$$A_w = \mathrm{ord}_w\, a, qquad D_w = A_w - A_{\bar{w}}.$$

For an integer n define the baseline for a string by

$$B_{s,n,i}(w) = nD_w + \begin{cases} 0, & s \text{ odd}, \\ -A_w, & s \text{ even},\ i < s/2, \\ A_{\bar{w}}, & s \text{ even},\ i \geq s/2. \end{cases} \tag{12.14}$$

These are the additive valuations of the diagonals of $B_s(a, \pi, (a/\bar{a})^n)$. In particular

$$B_{s,n,i}(w) + B_{s,n,s-1-i}(\bar{w}) = 0. \tag{12.15}$$

The condition called `Grading.HasNearLines` is that the difference between the actual valuation and (12.14) is bounded by k above $Q \cup \bar{Q}$ and is zero everywhere else. At stage zero it follows immediately when nonzero support lies in the Bessel orbit of the standard lattice: multiplying its adapted basis by a Bessel element leaves all diagonal entries unchanged. This is the base case `Grading.hasNearLines_zero` in `New/Seesaw/StandardLatticeLevel.lean`; the first Poincaré datum supplies the needed support condition.

Lemma 12.17 (Inequality extracted from a nonzero constant term)**.** *Let f be a finite datum invariant by all finite shifts with parameter $b\widehat{O}^n$, with b a finite idele. Suppose the constant term at T of its Weil translate at two adjacent lattice points is nonzero. Then, with $s' = s - 1$ and $0 \le i < s'$, one has for every place w*

$$\operatorname{ord}_w \beta'_i \le \epsilon_i \operatorname{ord}_w c + \operatorname{ord}_{\bar w} b - \operatorname{ord}_w \alpha - \operatorname{ord}_{\bar w} \beta_{s-1-i}, \tag{12.16}$$

$$\epsilon_i = \begin{cases} 1, & s' \text{ odd and } i \le (s'-1)/2, \\ 1, & s' \text{ even and } i > (s'-1)/2, \\ 0, & \text{otherwise.} \end{cases}$$

Proof. Use the finite support lemma, with the finite unitary action of the pair of lattice points and with $M = bg$, $M' = \bar b^{-1} g$. The product $M(M')^*$ is the identity. The isotropic frame consists of the strictly positive relative degree coordinates, so its annihilator imposes that the negative-degree part of the rational frequency agrees with the negative-degree part of an integral matrix in the translated frames. On returning to the rational Hom coordinates this gives a matrix W satisfying

$$\deg(W - T) \ge 0, qquad \alpha^{-1}\bar b\,(x'_f x')^{-1} W x_f x \in \operatorname{Mat}(\widehat{O}).$$

The factor α^{-1} is exactly the inverse scale in the finite Hom coordinates used in Lemma 12.16. The involution on b comes from $M' = \bar b^{-1} g$. Replacing either frame by an adapted basis does not change integrality. Thus

$$\alpha^{-1}\bar b\,(B')^{-1} W B \in \operatorname{Mat}(\widehat{O}).$$

Look at row i, column i in the string-to-string block. Since the source column has grade one more than this target row, the equality $W - T$ of degree at least zero gives $W_{ii} = t_i$. The matrices B and $(B')^{-1}$ are upper triangular for the flag. In the product $(B')^{-1} W B$, the contributing intermediate source grades are at least the grade of e_i and the intervening target grades are at most that of e'_i; this forces all such terms to lie in the negative-degree part of W, where it equals T. The only possible diagonal string contribution is then at row i, column i. More exactly, the integral dual basis of the lattice B is $J_s^{-1}(B^{-1})^*$; its diagonal at the ith string coordinate is $h_i^{-1}\bar\beta_{s-1-i}^{-1}$ with h_i the paired Gram coefficient. As the lattice is self-dual, we may use this basis for the domain. This gives the required integral scalar

$$\frac{\bar b}{\alpha}\,\frac{t_i}{h_i \beta'_i \bar\beta_{s-1-i}}. \tag{12.17}$$

In the odd source case the coefficient t_i/h_i equals 1 on the first half and c on the second half including the old middle index (the two factors $-N$ cancel). For even source it is c on the first half and 1 on the second. In both cases this is c^{ϵ_i}. Taking additive valuations of (12.17) proves (12.16). □

For clarity about the middle index in this calculation, the dual basis may be taken with its reflected column order, namely the columns of $J_s^{-1}(B^{-1})^*$ indexed by reflected strings. This is again upper triangular in the original string order by Hermitian reflection. Both these bases generate the same lattice, since $B^* J_s B$ is integral unimodular. Formula (12.17) is the scalar calculated in `Grading.HasNearLines.step` of `New/Seesaw/LinesStep.lean`.

Proposition 12.18 (The velocity of adapted lines)**.** *Assume that the source lattice satisfies the near-line condition with error k and exponent n, that $n' = n$ when s is odd and $n' = n - 1$ when s is even. Assume also*

$$\operatorname{ord}_w c = A_w + A_{\bar w}, \qquad \operatorname{ord}_w b \le \operatorname{ord}_{\bar w}(\alpha/a) + [w \in Q],$$

and suppose the next lattice has

$$\operatorname{ord}_w \beta'_i + \operatorname{ord}_{\bar{w}} \beta'_{s'-1-i} \geq 0. \tag{12.18}$$

Then a nonzero constant term at T forces the next lattice to satisfy the near-line condition with exponent n' and error $k+1$.

Proof. In (12.16) substitute the bound on the valuation of b and the near-line condition at the reflected index of the source. Using (12.15) gives

$$\operatorname{ord}_w \beta'_i \leq B_{s,n,i}(w) + \epsilon_i(A_w + A_{\bar{w}}) - A_w + (k+1)[w \in Q \cup \bar{Q}].$$

The first three terms are exactly $B_{s',n',i}(w)$. Indeed, for odd s the subtraction changes the first half of the next baseline by $-A_w$ and the second by $A_{\bar{w}}$, while $n' = n$. For even s the old half contributions cancel with the relevant two terms to give $(n-1)D_w$ in both halves. This proves the upper bound. Apply the same upper bound at the reflected index and conjugate place, then use (12.18) and (12.15) to obtain the opposite bound. Off the two supports all errors vanish. □

Most of (12.18) follows from (12.12): the Gram on an adapted basis is integral and a reflected pair has coefficient 1. The only apparent exception is the middle string index for odd length at a place dividing N, where integrality yields just the lower bound -1. The source eliminates this exceptional case by using ramified columns, rather than accepting a step at the level prime.

For the actual modulus in the construction, $a = \alpha/\bar{\lambda}$ where λ is the chosen finite idele generating the chosen-prime ideal, and b generates that ideal times the product of the places in Q. Therefore the valuation of b is exactly the valuation of $\overline{\alpha/a} = \lambda$ plus $[w \in Q]$. This is the inequality used in Proposition 12.18. Moreover the finite residue datum is a sum of translated vacua with translation parameters t satisfying $\bar{b}t \in \widehat{O}^h$. Each such vector is fixed by shifts $\rho_f(bx)$, x integral: the vacuum is fixed since b is integral, and the commutation phase with the translate is trivial by integrality of $\bar{b}t$ and self-duality of the finite pairing. The same holds for the sum. Thus the invariant-shift hypothesis of the finite support lemma is compatible with the level modulus; it does not require the translate itself to be an integral vector. These are the idele and modulus calculations of `New/Trace/GoodIdempotent.lean` and the residue expansion used in `hasNearLines_of_unfoldStage_ne_zero`.

Lemma 12.19 (The exceptional middle line)**.** *For s odd, if the middle pair sum of valuations equals -1 at a split place w of N, there is a representative in the same Bessel–integral-unitary double coset whose lattice admits the embedded full local maximal compact subgroup of $U(X)$ at w. Consequently, if a nonzero generic stage function has the property that it vanishes on all such representatives, then* (12.18) *holds.*

Proof. At the pair $(w, \bar{w})$ write the lattice in one local factor; its other local factor is the dual under the Hermitian form. Form a diagonal lattice at w as follows. On the positive and middle string lines, choose the valuations of the given adapted basis at w. For a negative string index j, choose the negative of the original valuation on its positive reflected partner at $\bar{w}$. Take the identity lattice on X. The dual at $\bar{w}$ has the correct valuations on the nonnegative string by these choices; for the middle string its valuation is minus the valuation at w minus one, since that Gram entry has valuation one. By the assumed defect this is exactly the original middle value at $\bar{w}$. Elsewhere keep the old lattice. This specifies a valid self-dual lattice, hence a point in the finite unitary group (use local dual bases at the split pair). Now use Proposition 12.15, since the nonnegative lines agree. The diagonal lattice admits the claimed local compact subgroup by its identity lattice on X. The asserted vanishing condition excludes this possibility. □

The ramification condition on the initial column is precisely designed to obtain the vanishing used in the lemma. At an even stage, an unfolded value is a coefficient of a paired definite lift. The action of a local scalar z at a place of N on the input column is intertwined with its action on this coefficient by the pair Weil equivariance; the relevant splitting twists are unramified there. Every ramified column lies in the image of $1-\rho(z)$ for each such local generator. If the output test value is invariant under z because the full local compact is contained in the stabilizer, it vanishes on this image. Thus, once the even-stage coefficient identity and its equivariance have been established, Lemma 12.19 applies at the odd-length stages. This is the exact use of `exists_eq_sub_centralGen`, `scalarFin_centralGen_mem` in `New/Columns/ChainValuation.lean` and `hasNearLines_of_unfoldStage_ne_zero` in `New/Seesaw/UnfoldingStages.lean`. The geometric part of the implication is proved here; the particular even-stage coefficient identity is discussed below.

12.10 Finite cuts with prime-supported multiplicity

There is an entirely group-theoretic argument which explains why this bound on lines is enough. Let H be a locally compact group, $J \subset A \subset B$ subgroups, where J, A are compact open, and let S be a closed subgroup. Let D be a compact subset of B such that for every $d \in D$ there is an integer $n_d \geq 0$ with

$$(dud^{-1})^{q^{n_d}} \in A \quad (u \in J). \tag{12.19}$$

Let a desired support $W \subset H$ satisfy $W \subset SD$. For a finite subset T of H/J stable under $S \cap A$ define

$$\mathrm{mult}_T(xJ) = \#\{r(S \cap A) \in S/(S \cap A) : r^{-1}xJ \in T\}.$$

It is constant on S-orbits.

Theorem 12.20 (Finite cut lemma)**.** *Under the preceding hypotheses there is a finite $S \cap A$-stable $T \subset B/J$ such that W/J is covered by ST, and the multiplicity on T and on W/J is nonzero and divides a power of q.*

Proof. The image of D in the discrete set H/J is finite. Choose one representative in this image of each distinct S-orbit meeting it. Call the resulting finite set C, and take $T = (S \cap A)C$. This is finite because the image of a compact group on a discrete set is finite. It lies in B/J and is stable as required. Every point of W/J is in the S orbit of one of the chosen representatives.

For a representative $dJ \in C$, put $U = S \cap dJd^{-1}$. Its multiplicity is exactly

$$[U : U \cap A].$$

Indeed the only points of T in that S-orbit are $(S \cap A)dJ$. The elements r with $r^{-1}dJ \in T$ form $U(S \cap A)$, and passage to right cosets of $S \cap A$ gives the asserted index. It is finite because U is compact and A open. Let V be the normal core of $U \cap A$ in U. The target index divides the order of the finite group U/V. By (12.19), for every $v \in U$ its q^{n_d}th power belongs to $A \cap U$, and the same is true for every conjugate of v in U. Thus $v^{q^{n_d}} \in V$. All element orders in U/V divide a power of q. By Cauchy's theorem every prime factor of its group order is a prime factor of q; hence its order divides some power of q. This proves the assertion, including the case $q = 1$ when the multiplicities must all equal one. □

This is `exists_cut_coverMultiplicity_dvd` in `New/Seesaw/IntegralAsCosetSum.lean`. Notice that $U \cap A$ was not assumed normal. Passing directly to a quotient group

by it rather than its core would invalidate the argument. The same cut gives a partition of unity along the S-orbits with their Haar measure. Normalize Haar measure dr of S by $\mathrm{vol}(S \cap A) = 1$, and define

$$\phi_T(x) = \frac{\mathbf{1}_T(xJ)}{\mathrm{mult}_T(xJ)}$$

with zero value off T. Then

$$\int_S \phi_T(r^{-1}x)\, dr = 1 \quad (x \in W). \tag{12.20}$$

Indeed the integrand is constant of value $1/\mathrm{mult}_T(xJ)$ on a union of exactly that many disjoint left cosets of $S \cap A$; invariance of the multiplicity along S was used. The formula requires no assertion of normality of the small level in H. It is this partition of unity, not counting of representatives without a weight, that enters the quotient periods.

Construction of the required compact set. For stage $k > 0$ let W be the support constrained by the near-line bound. At each place outside $Q \cup \bar{Q}$, the adapted valuations of any such point equal the fixed baseline, which is the valuation pattern of $x = 1$. At the remaining finitely many places they lie in fixed intervals of length $2k + 1$. The number of possible tuples of nonnegative string valuations at those places is therefore finite. For each tuple that actually occurs, choose one finite group point having that tuple, and replace its components outside $Q \cup \bar{Q}$ by the identity. This operation preserves the unitary condition, since the indicated set of places is stable under the involution. The adapted valuations at a given place depend only on the local point, not on choices elsewhere: they are the valuations of the graded flag quotients of the local lattice. Thus the modified point has the desired tuple at those places and the baseline elsewhere. By Proposition 12.15, every element of W belongs to the Bessel orbit of one of these finitely many points times K_0. Taking D to be their union of right K_0-cosets gives a compact set integral away from q. In particular this argument does not require a diagonal representative on the middle string line at a twist prime; such a representative need not have the two desired valuations. When the twist is trivial one simply takes $D = K_0$. For stage zero the initial Poincaré support lies in the S-orbit of the chosen level above g, so take that level as D and A.

Above a twist prime the subgroup J at that stage belongs to the first principal congruence group in the local lattice coordinates, which is pro-p at a residue characteristic p dividing q. For every fixed representative d, conjugation of sufficiently deep congruence subgroups by d is contained in the local maximal compact A; the finite quotient of the first principal group by any such deeper congruence subgroup is a finite p-group. Therefore some power of q as in (12.19) sends all elements at these places into the desired subgroup. Away from the twist primes no power is needed by integrality and the choice of A (the first stage has A equal to the chosen level and all other stages can use the maximal compact). At the first stage $d \in A$, $J \subset A$, and the condition holds even with exponent one. For stage $k > 0$, the embedding of the indicated column subgroup in A was proved above using the glued lattice. This supplies the compact set and prime support conditions of `Data.exists_stage_cut` from the near-line estimate, including the first-stage support and equality-of-line arguments. There is no need to bound the number of its cosets uniformly for nontrivial twists: only finiteness and prime support matter.

The measure factor. Suppose an integrable stage product $F = Py$ on the rational Bessel quotient of the ambient group is supported in W in its finite variable. In the compatible quotient disintegration, insert the partition of unity (12.20). This replaces integration over

the ambient finite quotient of the Bessel group by integration over the finite cut, with weight ϕ_T. Suppose the periods over the Bessel group on the cut factor as

$$\int_{[B_s]} F(r(x_\infty, x_f))\, dr = [K_0 : J], a(x_\infty) b(x_f J).$$

The compact-open K_0 is normalized to have mass one. Every J-coset has measure $[K_0 : J]^{-1}$ by Haar invariance, so this index cancels on summing over the finite cut. Integrate a against the induced real quotient measure to obtain a scalar κ. The result is

$$\int F = \kappa \sum_{t \in T} \frac{b(t)}{\mathrm{mult}_T(t)}. \tag{12.21}$$

This is `Grading.exists_integral_eq_sum_over_cut`. In the general case the Bessel group maps to its finite factor with kernel its real factor, so one uses a cutoff for the real subgroup in taking the quotient by stages. The integration of a against that cutoff produces the same scalar independently of b, J, T. The partition identity shows why neither the finite factor κ nor the reciprocal multiplicities can be changed by choosing a different representative of a J-coset. In particular the only variable denominators of (12.21) are supported at the rational primes dividing q.

12.11 Two edges at once: removal of the extra plane

The even stages of the chain are special. Start with an odd string length s, then $s-1$ (even), then $s-2$ (odd or, at the end, zero). The two maps of (12.5) are denoted $T : V \to V'$ and $T' : V' \to V''$. On the orthogonal sum $V^- \oplus V''$, where the form on V^- is the negative of that on V, let

$$R = \begin{pmatrix} T^* \\ T' \end{pmatrix} : V' \longrightarrow V^- \oplus V''.$$

By the moment-map identities,

$$R^* R = -TT^* + (T')^* T' = 0. \tag{12.22}$$

Thus the two frequencies of the successive periods fit into a single isotropic frequency. The identification of the two Hom spaces with the Hom from V' to $V^- \oplus V''$ uses the adjoint in the first component and the identity in the second. It is a symplectic identification, because the adjoint reverses the trace pairing and the form on the first outer space has been negated. Each relative degree is preserved by it. The identification is `arrowProduct`, and R is `stackedT`, in `New/Seesaw/EdgePair.lean`.

For $s > 1$ write z, z'' for the two middle vectors in the outer spaces; their norms in the sum are N and $-N$. The degree-zero part of this sum is $X^- \oplus X''$ plus their two-dimensional plane. A basis of that plane is

$$p = (z - z'')/(2N), \qquad v = z + z''.$$

Both vectors are isotropic and their pairing is 1. The source's two-step splitting first takes the degree-zero part of the Hom space and then removes this additional hyperbolic plane. When $s = 1$, the final space has no string and the degree-zero part instead contains the one vector z/π in the negative outer space; its norm is 1, and it is adjoined as an extra definite line. This is precisely the pair of constructions `middlePair` and `lastPair`.

The reason for removing the extra plane can be made entirely explicit. Every grade-raising unipotent in the first odd string has degree at least two, and its phase cancels the generic

character by (12.7). In the even intermediate string, degrees at least two cause no phase in the product of the two constant terms, by (12.22). There are also degree-one terms because the string grades are odd and the X'-grade is zero. The degree-zero part of the product $R(g-1)$ for a grade-raising unipotent g of the intermediate space can only arise from a source X'-column going to the grade-one string vector, or from a grade-minus-one vector going to X'. But R kills X', and on the grade-one vector its value is $c(z+z'')$: for the first component, the old middle coefficient of T was $-Nc$, divided in the adjoint by the old Gram coefficient $-N$; the second component has coefficient c by (12.5). Thus all such extra degree-zero terms translate along the isotropic line of the additional plane, times arbitrary X'-coordinates. The integral over the rational quotient of this additive group is exactly the additional zero term. The map from the degree-one part of the intermediate unipotent to these X'-coordinates is onto: each elementary map from the X'-block to the grade-one vector, together with its negative adjoint from the reflected vector to X', exponentiates to an element of the unipotent, and its higher-degree terms do not affect the projection. The uniformizing measure on that additive quotient is fixed up to a constant independent of the datum. When $s = 1$, the intermediate space is definite so this additional integration is absent.

In the mixed model the two-step zero term consequently gives an ordinary theta integral on the three definite spaces. More explicitly, the product of the two theta functionals in the orthogonal sum is a fixed nonzero multiple of the theta functional for their product datum. This follows in the classical Schrödinger model by taking the rational theta sum on a direct sum of Lagrangians: on elementary tensors the sum is the product of the two sums, and finite linear combinations of tensors, and their Schwartz completions, give the general case. The adjoint and the sign reversal use the naturally isomorphic opposite symplectic model. Taking constant terms along the two positive Hom coordinate groups commutes with this product identification by Fubini, because they become the direct sum of the positive coordinate groups. The subsequent folding just calculated removes the intermediate unipotent integrals. Thus the result of integrating the two original constant terms against a generic P is, up to a single normalization independent of the finite datum, the integral of P restricted to the first definite group against the double theta lift with the evaluated datum. This is the calculation in `ArrowPair.integral_constTermAtT_mul`, `New/Seesaw/DoubleLiftValue.lean`.

The finite part of its evaluated datum is

$$\mathrm{ev}_f(\rho_f(R_f)(f_1 \boxtimes f_2)),$$

where ev_f is the two-stage zero term just described followed by identifying with the tensor of the two definite Hom spaces, and R_f is R in normalized finite coordinates. We denote the whole evaluation by L. A tensor product of the two finite vectors is sufficient here: arbitrary real polynomial vectors are allowed after taking the real zero term because the real and finite realizations factor. More precisely, the zero term of the tensor of a real vector and a finite vector in the mixed model is the tensor of its real zero term and its finite zero term. This assertion is used with the real vector possibly depending on the two archimedean test points; the finite operator does not depend on those test points or the twist. On polynomial Fock vectors, integration against a theta kernel produces a linear-form coefficient of the ordinary definite theta lift; applying this identity twice and using linearity gives the corresponding coefficient for the double lift. General real zero-term vectors follow by the density of polynomial Fock vectors and continuity of the integrals on the compact definite quotients; the resulting functionals lie in a finite-dimensional space and thus their limits remain in that space. These are the content of `zeroTermFin_spec` in `New/Seesaw/EdgePair.lean` and `integral_theta_eq_lift` with subscript 2 (the two-lift identity) in `New/Seesaw/DoubleLiftValue.lean`.

The preceding argument establishes the two-step operation in the classical normalized factorized Weil representation. To apply it to the particular algebraic modular form spaces of

the GRH argument, one must still identify the real joint harmonic vector and the branching coefficient used to interpret the resulting integral as exactly the source's linear lift map. This entails a compatibility of the kernel and its normalization, rather than a new orbit computation: the forms must be the indicated type, the twist must occur with the indicated exponent on both pairs, and the algebraic model must use the same class measure. We retain that specialized representation-theoretic identification, as in `Data.isArrow_datumY`, `Data.unfoldStage_even` and `Data.exists_last_stage_map`, when making the final application. The finite denominator calculation can be done independently of it, as follows.

12.12 The split neighbour lattice in the paired evaluation

We need the exact finite-place operation produced by removing the additional plane, not just its valuation. Denote by P_N the chosen places above N and write J for the finite idempotent equal to one on P_N and zero elsewhere. Fix the integral idelic unit η of norm -1 from the Weil model. Let e, e' be the two last coordinates of the two definite blocks in the finite tensor product (all other coordinates will stay unchanged). There is a unitary matrix Γ, identity outside $P_N \cup \bar{P}_N$, whose two-by-two block at a chosen split place w is

$$\Gamma_w = \begin{pmatrix} \pi & (\pi-1)\eta \\ 0 & 1 \end{pmatrix}, \qquad \Gamma_{\bar{w}} = \begin{pmatrix} \bar{\pi}^{-1} & 0 \\ (\bar{\pi}^{-1}-1)\bar{\eta} & 1 \end{pmatrix} \tag{12.23}$$

with the indicated components read at their own places. Conjugating the second block and taking its transpose gives the inverse of the first one, so the paired element is unitary for the standard finite Hermitian form. This is the matrix γ in `exists_isAdmissibleEval_std` of `New/Seesaw/AdmissibleEvaluation.lean`.

Here is the elementary identity of characteristic functions that makes this matrix useful. In the factor at $\bar{w}$, put $p = \bar{\pi}$, which is a uniformizer; N/p is a unit. Put $r_0 = e + \bar{\eta} e'$. For each residue $y \in \mathbb{F}_N$ let $\tilde{y}$ be its integer representative and take translation $n(y) = \eta(\tilde{y}/N) r_0$ at $\bar{w}$ and zero off the conjugate places. In the finite coordinates of a rectangle there is one such residue variable per intermediate row and per chosen place. At a single such coordinate, the new pair of coordinates of $\Gamma x + n(y)$ is

$$\left(p^{-1}x_e + \eta\tilde{y}/N, \quad x_{e'} - \bar{\eta}x_e + \bar{\eta}(p^{-1}x_e + \eta\tilde{y}/N)\right).$$

If these two coordinates are integral, then multiplying the first by p shows x_e integral, and subtracting its contribution to the second shows $x_{e'}$ integral. Conversely, given integral $x_e, x_{e'}$, precisely one residue y makes the first coordinate integral (since both p/N and η are units), and then the second is integral too. Thus the characteristic function of the integral box on the conjugate factor is the sum of the characteristic functions of the corresponding shifted boxes after action by Γ. Taking products over coordinates and places gives the box identity with multiplicity exactly one.

In the split Schrödinger realization, the action on box functions is substitution by the linear transformation and a constant Weil multiplier. The above identity therefore proves that, up to one scalar independent of the translate, the operation $\omega_f(\Gamma)$ on the projection to the relevant integral box is the sum of finite shifts indexed by $y \in (\mathbb{F}_N)^{r \times P_N}$. Under the adjoint identification of the two Hom spaces, the shifts are the source's

$$\tau_N = \sum_y \rho_f(-y/N \text{ in the last column}) \otimes \rho_f(-\bar{y}/N \text{ in the last row}), \tag{12.24}$$

where y/N denotes the adele supported at P_N with the indicated representatives, not the diagonal rational element. To check the orientation, the first component of the product datum

is converted by η times conjugate transpose; the shift $\eta(y/N)r_0$ on the conjugate support becomes $-y/N$ in that first component since $\eta\bar{\eta} = -1$. In the other component the second coordinate of r_0 gives $-\bar{y}/N$. The shift group is supported on disjoint conjugate factors, so their products have only the finite character phases already accounted for in the finite translation law. This is exactly `translateSum` of `New/Seesaw/NamesAtSplitLevel.lean`. It explains why one must use the two correlated shifts in (12.24), not the product of two independent sums.

12.13 Weighted integrality of evaluations

We now check the exceptional lattice columns (the neighbour vectors) and denominators. Let β be the integral finite idele associated to the chosen primes above g, with $|\widehat{O}/\beta\widehat{O}| = g^f$ in the source normalization. Let δ be the unitary ratio $(\bar{\alpha}/\alpha)(a/\bar{a})$, equal to one above N. In this application it is the ratio of the chosen idele above g and its conjugate: $\bar{\beta}\delta$ is integral and has the same valuations as β. For a vector in the product Hom model attach the coefficient weight β to a coordinate in the first block before conjugate transpose, and $\bar{\beta}$ in the second block. A vector is allowed if its coordinate multiplied by its weight is integral away from the twist conductor q.

The columns of Γ are neighbour vectors in the following precise sense. For each standard basis column j there is c_j such that

$$\begin{gathered}\delta c_j = c_j, \qquad \bar{\pi}\,\Gamma_{ij} \in \widehat{O}, \qquad \Gamma_{ij} \in \widehat{O}\ (i \neq e, e'),\\ c_j - \pi^{-1}\Gamma_{ej} \in \widehat{O}, \qquad \bar{\eta}c_j - \pi^{-1}\Gamma_{e'j} \in \widehat{O}.\end{gathered} \tag{12.25}$$

To verify these assertions, use the formula in (12.23). At $w \in P_N$, choose $c_j = \pi^{-1}\Gamma_{ej} - h_j$ with integral h_j so that $\bar{\eta}c_j - \pi^{-1}\Gamma_{e'j}$ is integral: modulo the ideal π, the last two entries of the first row of Γ are $0, -\eta$, and those of the second are $0, 1$. Their relation is precisely $\Gamma_{e'j} \equiv \bar{\eta}\Gamma_{ej}$ (mod π). At $w \in \bar{P}_N$, π is a unit, and the two possibly nonintegral entries occur with $\Gamma_{e'j} \equiv \bar{\eta}\Gamma_{ej}$ (mod $\widehat{O}$), by subtracting the multiple appearing in (12.23). Again choose $c_j = \pi^{-1}\Gamma_{ej}$ up to allowed integral error. Elsewhere take $c_j = 0$ since π is a unit and $\Gamma = 1$. Thus c_j is supported above N, where $\delta = 1$, and the remaining integrality assertions are clear. This proves the neighbour condition of `IsNeighbourVector` in `New/Seesaw/AdmissibleEvaluation.lean`.

In the additional plane removed in a middle pair, the corrected embedding of a neighbour vector has zero coordinate on the first hyperbolic basis vector and has second coordinate proportional to c_j (on the isotropic line $z + z''$). We check it belongs to the two glued lattices. In an odd glued lattice, a vector with X part x and only middle string part uh is integral in its framed coordinates precisely when all ordinary X-coordinates are integral, $\bar{\pi}x_e$ is integral, and $h - \pi^{-1}x_e$ is integral. Indeed express it in the basis of Lemma 12.13; the coefficient of the last X vector is $\bar{\pi}x_e$ and the coefficient of the middle string is the stated difference. In the first outer space of the pair the corrected vector has $h = c_j$. In the second it has $h = \bar{\eta}\delta c_j = \bar{\eta}c_j$, because the two middle string scalings differ by $a/\bar{a}$ and the sign-reversing idele has product with that ratio equal to $\bar{\eta}\delta$. Both membership tests are precisely (12.25). At the last pair there is no hyperbolic plane; the old middle coordinate is instead the added last definite coordinate π^{-1} multiplied by the same sign and scaling ratio. In the old lattice the needed difference now reads

$$\bar{\eta}^{-1}\pi^{-1}\Gamma_{e'j} - \pi^{-1}\Gamma_{ej} \ \in \widehat{O}$$

and follows by subtracting the two congruences in (12.25). All other X''-coordinates are unchanged. Thus the same membership holds also at the terminal pair. These checks are the

lattice part of `Data.isAdmissibleEval_pairEvalY`, in `New/Seesaw/GoodDatumExists.lean`.

After converting to the normalized finite frames the integral matrix of this embedding, including the correction along the removed isotropic plane, is $B = ig + kc$ in (12.13), with $g = \mathrm{diag}(\kappa)\Gamma$ and $\kappa = 1$ except for the special terminal coordinate, where $\kappa = \delta$. Tensoring with the X' coordinates preserves integrality: in the even intermediate glued lattice the X' components are not rescaled. The adjoint B^* carries weighted translations to weighted translations with the corresponding weights, including κ. To check this local statement, if an entry of B is nonzero, either the source and target coordinates lie in blocks of the same weight, or they lie above N where both $\beta, \bar{\beta}$ are units. The additional terminal coordinate is the only change of block weight away from N, and multiplying its target weight $\bar{\beta}$ by δ gives exactly the valuations of the old weight β. Therefore in each nonzero summand of a coordinate of B^*y, its target weight has the valuation of the source weight at the corresponding conjugate place. Each term is integral under the weighted hypothesis on y, and so is their sum. This is the proof of the weighted matrix lemma `mulVec_mem_weightedTranslations` of `New/Seesaw/AdmissibleEvaluation.lean` in the present application. Also $B^*R_f = 0$, since the map R kills the definite grade-zero subspace X' and the embedding is supported in the two grade-zero outer blocks; the nonzero frequency thus gives no additional translation on the definite component. This last observation uses the removal of nonzero relative degrees before forming the finite frame matrix B.

To bound the remaining coefficients we use the standard vacuum sum in finite coordinates. For an invertible unitary g_0 let $L_0 = \widehat{O}^h$ be the standard lattice; put

$$I = [L_0 : L_0 \cap g_0^{-1} L_0].$$

With the fixed finite Weil multiplier $\Delta_f(g_0)$ the formula is

$$\omega_f(g_0)v_{\mathrm{vac}} = \frac{\Delta_f(g_0)}{I} \sum_{a \in L_0/(L_0 \cap g_0^{-1}L_0)} \rho_f(g_0 a) v_{\mathrm{vac}}.$$

Indeed project onto the subspace fixed by the standard integral lattice by averaging the finite Heisenberg translation representation over Haar measure one on L_0. This is the vacuum line. Applying the analogous projection for $g_0 L_0$ to the original vacuum yields a nonzero vector on the translated vacuum line. Its factor with the normalized vacuum is by definition the fixed multiplier with $|\Delta_f(g_0)|^2 = I$, and the average over L_0 is the sum over the indicated finite subgroup cosets divided by I. This is `omega_vacElem_sum` in `Lib/Weil/Construction/Projections.lean`. After moving a translate of the vacuum, the additional coefficients are roots of unity from the Heisenberg law. In particular when the fixed scalar $\Delta_f(g_0)$ is factored out, the coefficients have all conjugate absolute sum at most one. If $\bar{\beta} g_0$ is an integral matrix, the intersection lattice contains $\bar{\beta} L_0$, so I divides $|\widehat{O}/\beta\widehat{O}|^h$. The source uses this argument for the diagonal $g_0 = \mathrm{diag}(\kappa)$, while the special Γ operation is treated by the split identity (12.24); otherwise one would introduce an unnecessary denominator at N.

Combine this sum, the weighted matrix condition and the zero-term formula (12.13). For the evaluation map L of a pair of consecutive lifts and each admissible y, the evaluated vacuum translate is a finite combination of the vectors $C\tau_N(v_t \otimes v_{t'})$ with

$$\bar{\beta}t, \ \bar{\beta}t' \quad \text{integral away from } q,$$

for one constant C independent of q and y. The denominators of the coefficients divide $|\widehat{O}/\beta\widehat{O}|^h$ for h equal to the product of the intermediate definite rank and the sum of the two

outer definite ranks (with the terminal extra line). When $q = 1$ their conjugate absolute values sum to at most one. If the source finite vectors themselves are the indicators of residue conditions on this weighted lattice, their standard finite projector formula expresses them as the corresponding sums of these vacuum translates; the normalized finite multipliers at the two directions of a pair cancel as in the formula of the preceding chapter, leaving the same fixed scalar and bound. Thus the operation on a pair is the admissible evaluation of `IsAdmissibleEval` in `New/Seesaw/AdmissibleEvaluation.lean`, including its explicit exceptional neighbour columns. The specialized assertion that the initial finite residue functions have the precise chosen projector realization is still the one discussed in the CM and Weil-model chapters; the neighbour calculation establishes what happens once that realization is made.

The paired projections at the level prime. The two translated factors before applying τ_N are integral in both conjugate supports of N, because q and β are prime to N. Project both factors to the subspace fixed by the standard integral shifts above those places. In the sum (12.24), a nonzero residue y has a nonintegral translation in at least one factor. Its projection is zero by the self-dual pairing: commuting with an integral shift detecting its residue multiplies the purported projection by a nontrivial character. Only $y = 0$ survives, giving back the unshifted pair. Conversely every term of the τ_N expansion has its translation integral after multiplication by $N\bar{\beta}$ away from q. On a pair with intermediate rank r, the number of terms is N^{fr}; all their phases have conjugate absolute value one (or their coefficient is zero). This proves the projection-and-expansion assertions of `Data.projOnFinPart_pairTranslates` and `Data.isAdmissibleComb_reducedVacuum` in `New/Seesaw/GoodDatumExists.lean`, from the finite projection law and the now-proved neighbour formula.

12.14 The admissible combinations and the iterated identity

Let $m = 2\varrho + 2$ be the number of lifts and n_m the largest rank parameter of the chain, and put

$$K_{mcut} = g^{fn_m^2}, \qquad D_{meval} = K_{mcut}^{2m}.$$

In this section a combination of vectors with coefficient function $r(t)$ of finite support is called admissible with denominator d if every $dr(t)$ is integral under all finite coefficient embeddings at primes not dividing q. Under $q = 1$ also require $\sum_t |\sigma r(t)| \le B$ for every automorphism of the complex numbers. We record three simple closure properties, since losing one of these bounds in the iteration would lose the gain at the small nonsplit primes.

First, if $v = \sum_t r(t)v_t$ is such a combination and for each t the vector v_t is itself a combination of the w_u with denominator d' and size B', then the resulting coefficients are $c(u) = \sum_t r(t)r_t(u)$. All $dd'c(u)$ are integral away from q by closure of each valuation ring under sums and products. The complex conjugate sizes satisfy

$$\sum_u |\sigma c(u)| \le \sum_t |\sigma r(t)| \sum_u |\sigma r_t(u)| \le BB'.$$

Secondly, tensor products of combinations multiply their denominators and size bounds by the same argument, since their coefficients are products. Applying any complex-linear map to each vector leaves coefficients unchanged. Thirdly, replacing any summand by a root-of-unity multiple of an admissible vector leaves all the integrality and size conditions unchanged. These are precisely the rules `IsAdmissibleComb.bind`, `pure`, `map`, `of_eq_isOfFinOrder_smul` in `New/Seesaw/AdmissibleComb.lean`.

Use the one-step transfer successively. At each stage its finite cut is independent of the column; every finite support is in the orbit covered by the cut because of the near-line bound. Formula (12.21) writes the integration over that stage as one scalar times the weighted finite sum over the corresponding finite quotient. Fix the last adelic point to be the identity. For a linear map of finite tensor data $S(y)$ indexed by the remaining tuple of finite cosets, define the modified function in the next step of the induction by

$$S_{<k}(y) = \sum_{t \in T_k} \mathrm{mult}_{T_k}(t)^{-1} S(y_0, \dots, y_{k-1}, t, y_{k+1}, \dots).$$

The multiplicity factors are rational integers' reciprocals with primes dividing q only. Thus repeated operations produce an admissible combination with denominator 1. When $q = 1$, every multiplicity is one, and the stage cut has at most K_{cut} terms: each representative is integral and its quotient by the chosen principal level embeds in the general linear group over the residue ring at the chosen primes. Indeed reduction gives the map into that finite group, its kernel on K_0 (or the stage-zero chosen subgroup) contains the principal subgroup of the chosen level, and J_k contains that subgroup in the trivial twist case. The order of the general linear group is at most the number $(g^f)^{n_m^2}$ of all matrices. Consequently the m cuts have total size at most K_{cut}^m. The invariance of the combinations under the embedded column levels follows because those elements act on T_k by a permutation preserving its multiplicity, while at the last stage no update is made. This is the induction of `Data.exists_isAdmissibleComb_unfolding` in `New/Seesaw/SeesawTheorem.lean`.

At an even stage use the two-edge calculation of Section 12.11. The function left after that stage is a definite matrix coefficient of the paired-lift chain, with finite datum obtained by the product of the evaluated edge pairs. This is an induction on the number of pairs: for the first pair the initial restriction of the Poincaré series to the first definite group gives its coefficient, and the two-step zero-term calculation acts on it by the double theta lift. The same statement for one additional pair follows by applying that calculation to the preceding coefficient and composing the two definite lift maps. Thus the overall coefficient is linear in the finite tensor data.

By the preceding admissibility rules and weighted-evaluation calculation, each such tensor is a combination, up to the shared factor C, of reduced vacua $V_{mred}(t)$, one correlation operator τ_N on each pair, with denominator D_{eval}. Indeed on a pair the product of the relevant ranks h is at most $4n_m^2$, so its denominator divides $(g^f)^{4n_m^2} = K_{\mathrm{cut}}^4$. There are $m/2$ pairs, yielding D_{eval}. The size before this denominator is cleared is at most one for trivial twist; hence the combinations obtained by the preceding cut induction have size at most K_{cut}^m. Every t is integral after multiplication by $\bar{\beta}$ away from q. Expanding the τ_N factors then gives a combination of ordinary translated vacua integral after multiplication by $N\bar{\beta}$ away from q, and multiplies the size by at most $N^{fn_m m/2}$. Multiply the coefficients by D_{eval} to clear the common denominator and absorb its inverse into the final fixed scalar. The trivial-twist size bound becomes

$$K_{\mathrm{cut}}^{3m} N^{fn_m m/2}. \tag{12.26}$$

This is below $\exp(x^{60})$ by the numerical bounds of the data chapter: the first log term is $3mfn_m^2 \log g$ and the second $fn_m m \log N/2$, precisely the two crude size terms in the good datum construction of the preceding chapter. The important quantitative feature is that the conductor q has not appeared in (12.26) for the trivial twist, and N has not appeared in the denominator at any twist.

The evaluated tensor sum is invariant under the product of the intermediate column levels fixing the last component. This follows from the permutation invariance of the cuts and the equivariance of finite evaluation: simultaneous translation of the two endpoints of an

evaluated pair intertwines the rotation of the Hom factors. For shifts in the final expansion, apply the two projections at N to the reduced vacuum: only the untranslated factors survive, as in the preceding paragraph, and reapplying τ_N recovers the original reduced vacuum. By linearity this identity also holds for the entire tensor sum. Thus translating the invariant evaluated tensor into a sum of individual translated vacuum data produces the joint invariance and compatibility with the projected pair expansion required in `Data.IsGoodDatum`. This is the finite linear-algebraic content of `Data.exists_isGoodDatum_of_invariant`. The final operation of restricting the last definite space to the smaller definite group is the usual branching restriction. Taking the Petersson integral of every resulting coefficient against every vector of that last type determines the equality of the two lift chains. Its normalization, chosen in the source's `Data.exists_last_stage_map` of `New/Filling/DegreeBounds.lean`, is the fixed scalar independent of the twist. If the product of all the common scalars is zero, all the corresponding indefinite-chain values vanish and the zero combination is a good datum; otherwise divide by that product to choose the common nonzero rescaling. This explains the use of a nonzero unit in the source's conclusion.

12.15 The remaining representation-theoretic comparison

This chapter has supplied conventional proofs of the normalized CM conjugacy, the simultaneous torsion-translation and finite-adelic cluster law, the geometric orbit argument and its elementary mixed-model Fourier unfolding, the glued lattice and the integral Witt step, the velocity inequality and its exceptional middle-line condition, the cut lemma and its prime-supported denominators, and the finite neighbour-lattice evaluation of the paired datum. The latter gives the explicit denominator and size propagation in the iterated calculation. These are the specialized constructions underlying `Data.seesaw` in `New/Seesaw/SeesawTheorem.lean`. In particular the geometry of the cut is no longer a black box once the single-stage supports and coefficient identities are available.

For the application to all columns of the particular GRH construction, the remaining dependencies in this argument are as follows. The initial Bessel Poincaré series must have the prescribed generic character and finite support with its stated normal form; these are the filling identities in `New/Seesaw/BesselGeneric.lean` and the filling chapter. The explicit realization of the abstract Fourier–Weil representation by the theta vectors of the arithmetic construction must preserve their finite vacuum line and exact multipliers and the real joint harmonic types. In particular the coefficient formula for the concrete definite lift maps, and its specialization to each even stage in `Data.unfoldStage_even`, must be verified with its root factors and class measure, together with the final real branching pairing `Data.exists_last_stage_map`. These are classical-model comparison and exact representation choices, not further arbitrary orbit or lattice support assertions. They are used here as identified inputs, and the new arguments do not assert to reconstruct the entire analytic representation theory of those maps independently. The elementary induction of the final section proves the stated good-datum consequences subject to them. The distinction is important: a general period transfer theorem without the precise finite evaluation and coefficient normalization would not suffice to obtain the arithmetic upper bound.

Part III

Spanning: Bessel and Poincare series

Chapter 13

Formal degrees, elliptic trace comparison and the row dimension

There is a global representation-theoretic issue in comparing the row space (on a unitary group which is not definite) with a space of algebraic modular forms on a definite group. A comparison of dimensions of the full spaces would be false without hypotheses; in particular the discrete series at infinity cannot simply be replaced by its compact counterpart. The source uses the central projection onto the good finite types, and compares only the corresponding Hecke traces. We give the exact formal-degree normalization, explain the trace comparison by changing the signs of elliptic classes, and prove the finite algebraic step that isolates the good traces. The construction of the actual discrete-series kernel and its reproducing property, the classical discrete-series character formula on the elliptic Cartan, and the Galois-theoretic compatibility of the automorphic labels are distinguished from the deductions proved here.

13.1 Parameters and the compact-dual Haar measure

For a rank n unitary group put $J = \mathrm{diag}(I_p, -I_q)$, with $p + q = n$, and write $G_J = U(J)$; its maximal compact subgroup K_∞ is $U(p) \times U(q)$. The notation K_∞ in this chapter is not the auxiliary imaginary quadratic number field. Let $\lambda = (\lambda_i)_{i=1}^n$ be the Harish–Chandra parameter, and $l = 2\lambda$ the doubled integer vector used in the source. For an admissible regular parameter one has $l_i \equiv n - 1 \bmod 2$. Its decreasing rearrangement $\lambda_1^\downarrow > \cdots > \lambda_n^\downarrow$ is used for the compact form, with highest weight

$$\mu_i = \lambda_i^\downarrow - \rho_i, \qquad \rho_i = \frac{n + 1 - 2i}{2}.$$

In particular μ_i are integers and $\mu_i \geq \mu_{i+1}$, since the positive differences of the λ_i are integers. The compact representation has dimension

$$\dim V_\mu = \prod_{i<j} \frac{\lambda_i^\downarrow - \lambda_j^\downarrow}{j - i} = \frac{\prod_{l_i > l_j} (l_i - l_j)/2}{\prod_{a=0}^{n-1} a!}. \tag{13.1}$$

Indeed the Weyl character of μ on the diagonal torus is the quotient $\det(z_i^{\mu_j + n - j}) / \det(z_i^{n-j})$. Its limit at the identity is the displayed ratio: set $z_i = e^{t u_i}$ for distinct u_i and expand the two determinants in powers of t. In the numerator the first nonzero term has total degree $0 + 1 + \cdots + (n-1)$ and equals the product of the Vandermonde determinants of (u_i) and of $(\mu_j + n - j)$, divided by $\prod_{a=0}^{n-1} a!$. The corresponding factor for (u_i) and the factorials cancel in the ratio. Finally $\prod_{i<j}(j - i) = \prod_{a=0}^{n-1} a!$. This is the polynomial `weylDimensionPoly`

in `Lib/DiscreteSeries/FlenstedJensen/RepresentationDefs.lean`. A simultaneous additive shift of all parameters changes the determinant character but leaves (13.1) unchanged.

We must identify the scalar implicit in the Haar measure. On G_J the form $Jg^{-1}dg$ is skew-Hermitian. Its Lebesgue density uses the n diagonal imaginary coordinates and the real and imaginary coordinates of the upper triangular off-diagonal entries. Divide this left-invariant density by the volume of the compact group $U(n)$ in the corresponding coordinates at $J = I_n$. The result is the compact-dual measure of the source. Explicitly a Cayley chart is

$$g(Y) = (I + \tfrac{1}{2}J^{-1}Y)(I - \tfrac{1}{2}J^{-1}Y)^{-1}, \qquad Y^* = -Y.$$

Putting $A = I - \frac{1}{2}J^{-1}Y$, its logarithmic derivative in a skew-Hermitian direction dY is

$$Jg^{-1}dg = A^{-*}\, dY\, A^{-1}.$$

This follows by differentiating $g(Y) = 2A^{-1} - I$, and using $Y^* = -Y$ and $J^* = J = J^{-1}$. The real determinant of congruence $Y \mapsto C^*YC$ on the real space of skew-Hermitian matrices has absolute value $|\det C|^{2n}$. To see this, use the singular-value decomposition of C: the unitary factors preserve the Euclidean form on that space and have determinant of absolute value one. For $C = \operatorname{diag}(c_1, \dots, c_n)$ with $c_i > 0$, the diagonal coordinates are multiplied by c_i^2, while each off-diagonal pair of real coordinates is multiplied by c_ic_j. The product is $\prod c_i^{2n}$. The Jacobian in the Cayley chart therefore equals $|\det A|^{-2n}$, which is precisely `cayleyDensity` of `Lib/DiscreteSeries/Haar/CompactDualMeasure.lean` for $J^2 = 1$. The chart omits the locus where $g + I$ is singular, a proper real-analytic locus of measure zero. The volume normalization is thus an exact comparison of two Haar measures, not merely a proportionality up to an unknown real scalar.

For $J = I_n$ the volume of $U(n)$ in these Lie algebra coordinates equals

$$\operatorname{vol}_0 U(n) = \prod_{a=1}^{n} \frac{2\pi^a}{(a-1)!}.$$

One may compute it successively using the first-column map $U(a) \to S^{2a-1}$, whose fiber is $U(a-1)$. At the identity, the first-column tangent coordinates consist of one diagonal imaginary coordinate and the $a - 1$ pairs of off-diagonal real coordinates; these are the ordinary Euclidean tangent coordinates of the unit sphere, of density one with respect to the indicated Haar volume element. The fiber coordinates are exactly the smaller skew-Hermitian coordinates. Invariance under $U(a)$ and the invariant quotient measure then give the recurrence multiplying by the ordinary sphere area $2\pi^a/(a-1)!$ (from polar coordinates in the Gaussian integral on $\mathbb{C}^a$). This matches `volU` and `isProbabilityMeasure_haarU` in `Lib/DiscreteSeries/Haar/HaarMass.lean`. In particular the measure of the compact form in this normalization is one.

13.2 Evaluation of the constant in the formal degree

For an irreducible unitary discrete series π let v be any unit vector. Its formal degree with respect to the preceding Haar measure is

$$d_\pi = \left(\int_{G_J} |\langle v, \pi(g)v \rangle|^2 \, dg \right)^{-1}.$$

The independence of v follows from Schur orthogonality for a square-integrable irreducible representation: integrating a rank-one operator against conjugation gives an operator commuting with the representation, and polarization gives the stated coefficient orthogonality. In

the compact case with probability measure this reduces to the elementary formula $d_\pi = \dim \pi$. The formal degree changes by the inverse of the scalar by which Haar measure is changed.

We use the classical Harish–Chandra formal-degree theorem: for a fixed real unitary group and Haar measure, its regular discrete series of parameter λ satisfy

$$d_\pi = C_J \prod_{i<j} |\lambda_i - \lambda_j|,$$

with a constant independent of the parameter (Harish–Chandra, *Discrete series for semisimple Lie groups II*, Acta Math. 116 (1966), in particular the formal-degree formula). Instead of relying on a convention for C_J, we evaluate it using one scalar holomorphic discrete series. This is also a check of the factor two implicit in the doubled parameter. The classification of the scalar holomorphic discrete series and the lowest K_∞-type formula $\Lambda = \lambda + \rho_n - \rho_c$ are the standard classical statements used in this calculation.

For $p, q > 0$ use the graph coordinate $Z \in M_{q,p}(\mathbb{C})$ on G_J/K_∞, where $I_p - Z^*Z > 0$. If g has blocks $(A, B; C, D)$, its action is $Z' = (C + DZ)(A + BZ)^{-1}$, and

$$I_p - Z'^*Z' = (A + BZ)^{-*}(I_p - Z^*Z)(A + BZ)^{-1}.$$

For the compact dual Grassmannian $U(n)/K_\infty$, the same graph coordinate uses $I_p + Z^*Z$ instead; it covers the complement of the subspaces whose projection to $\mathbb{C}^p$ is noninvertible, a subset of measure zero. Differentiate the graph action to obtain

$$dZ' = (D - Z'B)\, dZ\, (A + BZ)^{-1}.$$

A block determinant calculation gives $\det(D - Z'B)\det(A + BZ) = \det g$; one can multiply g on the right by the block lower triangular matrix with $(2,1)$ block Z and diagonal identities and then use the Schur complement. The absolute complex determinant of g is one. Thus the real Jacobian of $dZ \mapsto dZ'$ has absolute value $|\det(A + BZ)|^{-2n}$. Consequently the invariant quotient densities on the two graph charts are respectively

$$c \det(I - Z^*Z)^{-n}\, dZ, \qquad c \det(I + Z^*Z)^{-n}\, dZ.$$

Here the constant c is the same for both: at $Z = 0$ the horizontal tangent matrices in the two groups correspond to the same complex off-diagonal block dZ, differing only by the sign in the opposing block, hence with the same volume in the skew-Hermitian coordinates of the preceding subsection. The fiber K_∞ has the same tangent density on both sides, and we use its probability measure for both disintegrations. Since the compact quotient has total mass one,

$$c^{-1} = \int_{M_{q,p}(\mathbb{C})} \det(I + Z^*Z)^{-n}\, dZ.$$

Take an integer $k > n$, and the scalar weighted Bergman representation on the ball $I - Z^*Z > 0$, with squared norm proportional to

$$\int |f(Z)|^2 \det(I - Z^*Z)^{k-n}\, dZ.$$

The usual holomorphic action on functions is $f \mapsto \det(A' + B'Z)^{-k} f(g^{-1} \cdot Z)$, where $(A', B'; C', D')$ are the blocks of g^{-1}. Preservation of this norm follows from the two transformation formulas just given. The constant functions form the K_∞-type $(\det)^k \boxtimes 1$. Rotation $Z \mapsto e^{it}Z$ in the second block of K_∞ projects a holomorphic function onto its constant term by averaging: expand its power series in the scalar variable e^{it}, or apply Cauchy's integral formula on each complex line. Every nonzero invariant closed subspace contains the constants,

since one may move a point where a nonzero function is nonzero to the origin and then average over these rotations. Applying the same argument to an invariant orthogonal complement proves irreducibility. With the constant function normalized to norm one, projection onto that function is evaluation at zero. Its coefficient at g is therefore $\overline{\det A}^{-k}$ (up to the convention on which argument of the inner product is linear). In particular its squared absolute value is $|\det A|^{-2k} = \det(I - Z^*Z)^k$ for $Z = g \cdot 0$.

We evaluate the two matrix integrals explicitly. For real $s > p + q - 1$ and $t > -1$ put

$$I_+(s) = \int_{M_{q,p}(\mathbb{C})} \det(I + Z^*Z)^{-s}\, dZ, \qquad I_-(t) = \int_{Z^*Z<I_p} \det(I - Z^*Z)^t\, dZ.$$

Writing the last column of Z as z and the rest as Z_0, let $B = I \pm Z_0 Z_0^*$. The determinant identity

$$\det(I \pm ZZ^*) = \det B\,(1 \pm z^*B^{-1}z)$$

follows by factoring out B and taking the determinant of a rank-one perturbation of the identity. The domains correspond to all z in the plus case and to $z^*B^{-1}z < 1$ in the minus case. Substitute $z = B^{1/2}u$, whose real Jacobian is $\det B$. In the plus case the remaining exponent is $-(s-1)$ and the u integral is $\pi^q\Gamma(s-q)/\Gamma(s)$. In the minus case the remaining exponent is $t+1$ and the u integral is $\pi^q\Gamma(t+1)/\Gamma(t+q+1)$. These scalar integrals follow by radial integration and the beta integral, with the sphere area computed above. Iterating column by column gives

$$I_+(s) = \pi^{pq} \prod_{j=0}^{p-1} \frac{\Gamma(s-q-j)}{\Gamma(s-j)}, \tag{13.2}$$

$$I_-(t) = \pi^{pq} \prod_{j=0}^{p-1} \frac{\Gamma(t+1+j)}{\Gamma(t+q+1+j)}. \tag{13.3}$$

The chosen ranges guarantee absolute convergence at every stage. We used $\det(I \pm Z^*Z) = \det(I \pm ZZ^*)$, which follows by applying the determinant formula to the two block decompositions of a matrix with identity diagonal blocks.

In the Bergman realization the squared coefficient integral is finite and equals $I_-(k-n)/I_+(n)$, by disintegration along G_J/K_∞. Thus the representation is square-integrable and its degree is

$$\frac{I_+(n)}{I_-(k-n)} = \prod_{i=1}^{p} \frac{(p-i)!}{(n-i)!}\frac{(k-i)!}{(k-q-i)!}.$$

For this holomorphic discrete series the parameter with positive noncompact roots $e_i - e_{p+j}$ is, by the displayed lowest-type relation,

$$\lambda_i = k + \frac{p-q+1}{2} - i \quad (1 \le i \le p), \qquad \lambda_{p+j} = \frac{n+1}{2} - j \quad (1 \le j \le q).$$

Here ρ_c is half the sum of positive roots within each of the two blocks, and ρ_n is half the sum of the cross-block roots. Since $k > n$, these λ_i are indeed strictly decreasing in the indicated order. The differences within the two blocks are exactly those of ρ. The cross-block differences are $k - q - i + j$, and their ratios to the corresponding differences of ρ multiply to

$$\prod_{i=1}^{p}\prod_{j=1}^{q} \frac{k-q-i+j}{p+j-i} = \prod_{i=1}^{p} \frac{(k-i)!}{(k-q-i)!}\frac{(p-i)!}{(n-i)!}.$$

This is exactly the degree just computed. Harish–Chandra's parameter-independent proportionality now determines C_J for all regular discrete series of G_J: it is $(\prod_{i<j}(j-i))^{-1}$. If $p = 0$ or $q = 0$ the conclusion follows instead directly from compact Schur orthogonality and (13.1).

Proposition 13.1 (Exact formal-degree normalization)**.** *For a regular discrete series of $U(p,q)$ with doubled parameter l, with respect to the compact-dual measure used in the source, one has*

$$d_l = \dim V_\mu = \mathrm{WeylDim}(l),$$

where μ and the right hand side are given by (13.1)*. Thus the formal degree is independent of the signature for the corresponding compact infinitesimal character.*

For this use of the classical formal-degree theorem, it remains essential that the actual source coefficient `HCParam.dsCoeff` at the regular parameter be the normalized coefficient of the indicated discrete series (the Flensted–Jensen identification in `Lib/DiscreteSeries/FlenstedJensen/`). The preceding calculation fixes the constant in the classical degree formula; it does not prove that an unrelated function belongs to this representation. The degree equality is the first part of `Orbital.harish_chandra_degree_and_orbital` in `Lib/DiscreteSeries/Character/FormalDegree.lean`.

13.3 The exact elliptic orbital character identity

For a normalized unit vector v in the discrete series, set

$$e_l(g) = d_l\,\overline{\langle v, \pi_l(g)v\rangle}.$$

The source makes this choice with its highest coefficient; in particular it is a *minimal* convolution idempotent, not the character distribution of the representation. Indeed the coefficient orthogonality with degree d_l gives $e_l * e_l = e_l$ and $e_l(g^{-1}) = \overline{e_l(g)}$. These two identities also explain why taking the trace of a projection onto this type counts one coordinate per occurrence of the representation, rather than its full K_∞-type dimension.

We invoke the following precise classical orbit-character statement on the regular elliptic compact Cartan. For $z = (z_1, \dots, z_n)$ with pairwise distinct complex numbers of modulus one, write t_z for the corresponding diagonal element. Use on the centralizer torus $U(1)^n$ its probability Haar measure. Let

$$O_l(z) = \int_{U(1)^n \backslash G_J} e_l(g^{-1} t_z g)\, dg.$$

For the very regular parameters used in the source this integral exists. The stable sum is

$$(-1)^{pq} \sum_{\sigma \in S_n} O_l(z \circ \sigma) = p!\, q!\, \overline{\chi_\mu(t_z)}. \tag{13.4}$$

Here χ_μ is the compact Weyl character and the Haar measure on G_J is the compact-dual measure. This is the second assertion of `Orbital.harish_chandra_degree_and_orbital`; the relevant classical result is Harish–Chandra's discrete-series character formula on the compact Cartan (Acta Math. 116 (1966), Theorem 16, with the orbital-integral character formula of part I). We retain (13.4) in this precise normalization as the elliptic orbit-character input. Replacing the integral by a character value at each individual permutation would be false. Its *sum* has the asserted compact character and the sign $(-1)^{pq}$.

The integrand has been stated in the actual source convention. Haar integration over the whole noncompact group against a centralizer-invariant function would in general not be suitable; at regular elliptic elements the centralizer is compact of volume one, so the integral may equivalently be written on the whole group, as in the source. The stable sum divides by $p!\,q!$, because within each of the two signature subspaces the roots may be permuted by the compact subgroup. That division will be exactly the number of orderings of each signed orbit in the trace calculation below.

13.4 Absolute convergence and a direct ellipticity argument

The trace comparison in the formal source has explicit decay hypotheses. This is important for using ordinary integral kernels instead of an abstract pseudo-coefficient and a truncated trace formula. Write $F = E^+$, $f = [F : \mathbb{Q}]$, and

$$R(g) = \prod_{\tau \in \Phi} \|g_\tau\|_{\mathrm{op}} \quad (g \in G_\infty).$$

Then $R(g) \geq 1$, $R(g^{-1}) = R(g)$ and $R(gg') \leq R(g)R(g')$ by the identity $g_\tau^{-1} = J_\tau g_\tau^* J_\tau$. Let $\mathcal{O} \subset G_f$ be compact, with the condition that every $h \in \mathcal{O}$ has irreducible characteristic polynomial at some finite place of E. We consider the real product kernel $e = \prod_\tau e_{l_\tau}$. The precise `IsCuspKernel` hypothesis of order

$$B = 16[E : \mathbb{Q}](n+1)^3$$

provides $e * e = e$, a bound $|e(g)|R(g)^B \leq C$, integrability of $|e(g)|R(g)^{B+1}$, and a compact bump reproducing e. The specialized proof of this final kernel construction for the actual large parameters is part of the real discrete-series analysis; here we give the deductions from these properties.

Lemma 13.2 (Cyclic vector and power-packing bound)**.** *Let V be a framed rational Hermitian space of rank n, with fixed finite standard lattice and real norm as above. For every adelic point x and $T \geq 1$, the number of rational unitary elements γ such that*

$$(x^{-1}\gamma x)_f \in \mathcal{O}, \qquad R((x^{-1}\gamma x)_\infty) \leq T$$

is at most $C_{\mathcal{O},E,n} T^{[E:\mathbb{Q}](n+1)^3}$, independently of x.

Proof. Compactness of $\mathcal{O}$ in the finite adeles supplies a single positive rational integer c with ch carrying the standard lattice into itself for all $h \in \mathcal{O}$. Such an integer is obtained by bounding the finitely many denominators of the coordinates, since a compact set in the restricted product lies in a finite union of restrictions with all remaining coordinates integral. Use the adelic frame of V together with x to pull this standard finite lattice back to a lattice L_x in E^n, and pull the real Euclidean product norm back at infinity. Under that real norm the discrete full lattice in $\mathbb{R}^{[E:\mathbb{Q}]n}$ has a positive first minimum δ_x; choose a nonzero vector u of norm at most $2\delta_x$.

For each indicated γ, the vectors $c^j\gamma^j u$ $(0 \leq j \leq n)$ belong to L_x and have norm at most $2\delta_x(cT)^n$. The characteristic polynomial of γ over E is irreducible: a factorization of a monic polynomial over the field would remain proper after extension to every local field, contradicting the hypothesis on $\mathcal{O}$. Consequently $u, \gamma u, \dots, \gamma^{n-1}u$ form an E-basis. Indeed the annihilator of the nonzero cyclic vector in $E[X]$ is generated by a divisor of the irreducible characteristic polynomial (use Cayley–Hamilton), so must equal it. If two such γ, γ' give the same tuple $(c^j\gamma^j u)_{j=0}^n$, they act identically on that basis and are equal.

The number of lattice points in a Euclidean ball of radius $2\delta_x(cT)^n$ is at most $(1 + 4(cT)^n)^{[E:\mathbb{Q}]n}$: place disjoint open balls of radius $\delta_x/2$ around the points and compare total Euclidean volume to the enlarged ball. There are at most the $(n+1)$-st power of this number for the tuples. The resulting exponent $[E : \mathbb{Q}]n^2(n+1)$ is at most the exponent in the assertion. All constants are independent of the first minimum and hence of x. □

Summation in dyadic ranges of T now shows that $\sum_\gamma |k((x^{-1}\gamma x)_\infty)|\, \mathbf{1}_{\mathcal{O}}((x^{-1}\gamma x)_f)$ is uniformly bounded in x for any function k with $|k(g)|R(g)^B$ bounded. The finite-volume

assertion for a rational unitary quotient is the standard arithmetic-group covolume theorem (Borel–Harish-Chandra; the center here is anisotropic). Thus the sum is absolutely integrable over the rational quotient. Grouping by rational conjugacy classes, and replacing a sum over the translates of a class by the integral over its rational centralizer quotient, is justified by absolute convergence. This proves the relevant convergence assertion `Space.orbitalIntegral_unfold` with its explicitly larger exponent B. In particular it applies not just to e but also to the real majorant R^{-B}.

We next prove the special Selberg principle needed here from convolution idempotence and that convergence. Let γ have irreducible characteristic polynomial and an eigenvalue λ at an archimedean place with $|\lambda| \neq 1$. Fix a nonzero left eigenvector (a row) Q there. Every element a of its rational centralizer acts on that row by a scalar $\chi(a)$: the relevant eigenline in the dual space is one dimensional since the characteristic polynomial is separable. Put $H(x) = \log \|Qx_\tau\|$ for an adelic point expressed in the chosen real frame. Then $H(ax) = H(x) + \log|\chi(a)|$, so differences $H(xy) - H(x)$ for $y \in G_\infty$ descend to the centralizer quotient. Moreover

$$|H(xy) - H(x)| \leq \log R(y) \leq R(y), \qquad H(xd) - H(x) = \log|\lambda|,$$

where $d = (x^{-1}\gamma x)_\infty$ acts only in the real factor (the last equality uses only the real coordinates). The first inequality follows by the operator norm bound, applied in both directions using $R(y^{-1}) = R(y)$.

Let I be the orbital integral of $e(d)\mathbf{1}_{\mathcal{O}}((x^{-1}\gamma x)_f)$ on the centralizer quotient. In the integral write $e(d) = (e * e)(d)$ and insert the identity

$$\log|\lambda| = (H(xy) - H(x)) + (H(xd) - H(xy)).$$

The integral associated to the first term is

$$J_2 = \int_x \int_y e(y)e(y^{-1}d)(H(xy) - H(x))\, \mathbf{1}_{\mathcal{O}}((x^{-1}\gamma x)_f)\, dy\, dx.$$

It is absolutely convergent: by submultiplicativity the integrand in absolute value is bounded by

$$CR(d)^{-B}|e(y)|R(y)^{B+1}\mathbf{1}_{\mathcal{O}}((x^{-1}\gamma x)_f),$$

which is integrable by the packing bound and the moment hypothesis. The second term has the same integral by the change of variables $(x, y) \mapsto (xy, y^{-1}d)$. This transformation preserves product measure: first translate x on the right by y, then invert and translate y by a fixed group element depending on the new x. All relevant groups are unimodular and the rational centralizer quotient has invariant measure. The finite indicator is unchanged. Thus

$$\log|\lambda|\, I = 2J_2.$$

Alternatively use the triple convolution $e * e * e = e$ and the three successive increments along x, xy, xyz, xd. The first increment has integral J_2 after integrating the convolution in z; its absolute convergence follows from the bound

$$CR(d)^{-B}|e(y)|R(y)^{B+1}|e(z)|R(z)^{B}\mathbf{1}_{\mathcal{O}}((x^{-1}\gamma x)_f).$$

The cyclic change of variables $(x, y, z) \mapsto (xy, z, z^{-1}y^{-1}d)$ preserves measure by the same argument, and permutes the increments. It therefore proves convergence of the other two and equality of their integrals with J_2. Hence $\log|\lambda|\, I = 3J_2$. Subtracting the two identities gives $J_2 = 0$ and then $I = 0$. This proves precisely the required assertion: a nonzero orbital integral on this irreducible finite support can only come from a rational element whose complex eigenvalues all have modulus one. It is `Space.elliptic_of_integral_ne_zero` in `Lib/LocalGL/TraceFormula/TraceComparison.lean`. No noncompact orbital integral is silently equated to a compact orbital character.

13.5 Changing the signatures of the rational elliptic classes

The source changes real signatures while keeping the finite local groups and the parameter fixed. The elementary global algebra behind this operation uses the norm theorem, not strong approximation. Here are the details.

Let $\gamma \in \mathrm{GL}_n(E)$ have irreducible characteristic polynomial and suppose it preserves a rational Hermitian form J and has only eigenvalues on the unit circle at infinity. Its rational centralizer is the multiplicative group of the field $B_\gamma = E[\gamma]$: the cyclic-vector argument identifies a commuting linear map with multiplication by its value on a cyclic vector. The adjoint involution with respect to J on this centralizer induces an involution $\dagger$ on B_γ extending conjugation of E and satisfying $\gamma^\dagger = \gamma^{-1}$; this does not depend on the particular preserved form because these prescriptions determine it. At each embedding of B_γ into $\mathbb{C}$ extending the chosen CM type, the image of γ^{-1} is the complex conjugate of the image of γ. It follows on generators that $\dagger$ is complex conjugation on every embedding. Thus B_γ is a CM field, with totally real fixed subfield B_γ^+. The rational unitary centralizer is precisely $\{a \in B_\gamma^\times : a^\dagger a = 1\}$, independent of the preserved form.

A frame class may be described by a pair (x_f, γ) consisting of a finite adelic invertible matrix and the rational matrix, modulo simultaneous rational conjugacy (left transformation of the frame and conjugation of γ) and right multiplication of x_f by the fixed standard finite unitary group. The pair is admissible if some rational J preserved by γ satisfies $x_f^* J x_f = I$. This condition determines J uniquely when it exists. For fixed γ, after choosing numbered roots at every complex place, a frame class has a sign on every eigenline: eigenlines with distinct eigenvalues are mutually orthogonal, and each is nonisotropic because the Hermitian form is nondegenerate. Denote the total sign assignment by $\epsilon_{\tau,i}$.

Lemma 13.3 (Norm change of a signed elliptic frame)**.** *Any prescribed changes $\delta_{\tau,i} = \pm 1$ of those signs whose global product is $+1$ can be realized by a bijection of these frame classes. It commutes with their centralizers, preserves their finite positions $x_f^{-1}\gamma x_f$, and preserves the finite centralizer embeddings into the standard finite group.*

Proof. The sign changes are indexed by the real embeddings of B_γ^+. The global quadratic norm theorem for B_γ/B_γ^+ gives $a \in (B_\gamma^+)^\times$ with exactly these signs and a finite idele z of B_γ with $az^\dagger z = 1$ in the finite adeles. Indeed the infinite idele of the prescribed signs with finite part 1 has quadratic Artin symbol equal to their product, so lies in the principal-times-norm subgroup. The real norms are positive, yielding the stated sign of the principal term. This argument is the same as in Corollary 6.5, and does not require B_γ/B_γ^+ to be unramified.

View a and z as multiplication operators in the centralizer algebra and put

$$J' = Ja, \qquad x'_f = z x_f.$$

The first is Hermitian and preserved by γ, since a is adjoint-fixed and commutes with γ. At the finite places,

$$(x'_f)^* J' x'_f = x_f^* J z^\dagger a z x_f = x_f^* J x_f = I.$$

At infinity a acts by the scalar of sign $\delta_{\tau,i}$ on each eigenline. Hence the resulting form has exactly the required modified signs. The rational matrix γ is unchanged, the finite position $(x'_f)^{-1}\gamma x'_f$ equals the original one, and every element of its centralizer commutes with z, so its finite position also remains unchanged.

For the assertion about classes, fix the characteristic polynomial and identify all its cyclic rational matrices with one companion matrix by rational conjugation. The multipliers a, z can then be chosen once for each fixed sign assignment. The conjugations identifying a rational

matrix with the companion differ by an element of its centralizer, which commutes with both multipliers. Thus the construction is well-defined on frame classes. Its inverse uses a^{-1} and z^{-1}, proving bijectivity. □

This is the mechanism of `exists_hasSignType_mul_iff` in `Lib/LocalGL/TraceFormula/ChangingSigns.lean`. In the comparison of signatures (p, q) constant over Φ, every choice of q negative eigenlines at each real place differs from the positive assignment by a sign vector of product $(-1)^{qf} = +1$, since f is even. Lemma 13.3 therefore gives all the signed classes from the positive ones with their finite positions and centralizers preserved. Conversely every contributing signed class is obtained uniquely from its positive counterpart for each specified assignment. Equivalence of rational Hermitian spaces with the same local signatures and frames follows from Landherr's theorem, as explained in Chapter 6. This identifies those classes in any of the fixed global spaces used for the trace computation. In particular it is not necessary to assume that all indefinite unitary spaces with arbitrary local invariants have the same automorphic spectra.

13.6 The trace of the elliptic finite test functions

Let $\mathcal{V}_s(K_f)$ be the principal-level row-type space on a framed space of signature s, with all real parameters fixed. Here K_f denotes the principal level at the chosen places (a compact open subgroup of the standard finite maximal group). For $h \in G_f$, consider the compression of right translation by h to this fixed-level space; it is represented by the finite test function $1_{hK_f}/\operatorname{vol}(K_f)$. If h normalizes K_f this is the ordinary translation on the fixed-level space; more generally one may compose with the orthogonal projection to the level space. We use the source's reproducing-kernel property of the actual cusp kernel: the real idempotent and the finite-level averaging operator have a finite-rank automorphic reproducing kernel whose range is precisely $\mathcal{V}_s(K_f)$ (see `Lib/Adelic/CuspForms/ReproducingKernel.lean`, `Space.exists_kernel_form`). This is an identification of the kernel with the actual rapid cusp forms, not merely an application of the formal-degree polynomial.

The ordinary kernel trace, on multiplying by $\operatorname{vol}(K_f)$, is

$$\int_{[G_s]} \sum_{\gamma \in G_s(F)} e((x^{-1}\gamma x)_\infty) \mathbf{1}_{hK_f}((x^{-1}\gamma x)_f)\, dx.$$

For example it follows by expressing a reproducing kernel in any orthonormal basis of its finite-dimensional range, composing with translation, and integrating its diagonal. The equality with the rational sum is the periodized kernel identity; the normalization by $\operatorname{vol}(K_f)$ comes from averaging over the level. Off the diagonal the periodizing sum converges on compact sets: compact finite support bounds the rational matrix entries in a fixed denominator lattice on any such set, and the archimedean norm bound reduces counting to a Euclidean lattice point estimate of dimension $[E : \mathbb{Q}]n^2$. The exponent B above is more than sufficient for this estimate. For the diagonal and the integral above, the stronger uniform estimate of Lemma 13.2 applies under the irreducibility hypothesis on the coset. Thus regrouping by rational conjugacy classes and applying the ellipticity argument of the preceding subsection is legitimate.

Here is the remaining comparison for the elliptic part. Fix a positive definite frame class with its rational γ and the finite frame. Its centralizer at each real place is $U(1)^n$, and the finite centralizer embedding and finite position are the same for all the sign assignments obtained from Lemma 13.3. The nonarchimedean factors of the orbital integral are therefore unchanged. More formally, take product Haar measures on the centralizer adelic group, choosing at the

real places its probability torus measures; its rational subgroup and the quotient by it are independent of the signed frame by the centralizer description above. Disintegrating over the adelic centralizer and its quotient factors every geometric integral as a common centralizer volume times its finite orbital integral times a product of real orbital integrals. This is just invariant quotient integration: on changing the real eigenbasis, the conjugation identifies the centralizers with the same probability compact torus. Thus none of these factors introduce signature-dependent real constants. The disintegration is justified by the absolute convergence proved above.

At each real place, the sum over the choices of q negative eigenlines is the sum over S_n divided by $p!\,q!$: permutations within the positive and negative blocks are represented by the compact group and give the same orbital integral. Formula (13.4) makes this sum equal to $(-1)^{pq}\overline{\chi_\mu(t_z)}$, the compact orbital integral with the indicated sign. Across the f real places the total factor is $(-1)^{pqf} = 1$. Hence summing the orbital contributions in the signed space yields the very same number as summing them in the positive space. This argument proves signature independence of the trace for every such finite test coset under the stated cusp kernel and elliptic character input:

$$\operatorname{tr}_{\mathcal{V}_s(K_f)} T_h = \operatorname{tr}_{\mathcal{V}_{s'}(K_f)} T_h. \tag{13.5}$$

The proof follows the mechanism of `Space.exists_trace_eq` in `Lib/LocalGL/TraceFormula/TraceComparison.lean` (compare Labesse, *Changement de base CM et séries discrètes*, Thm. 5.4). The two genuine representation-theoretic inputs here are the real orbital character identity (13.4) and the identification and decay of the actual reproducing cusp kernel. We have given independently the algebraic movement of the framed orbits, the convergence mechanism, and the idempotent proof of elliptic support. The trace identity applies to the exact finite irreducible-coset support, not to all finite test functions indiscriminately.

13.7 Primitive labels and the finite cubic correction

The passage from (13.5) to equality of traces on the good subspaces is a separate finite argument. We detail it, including why it does not amount to assuming equality on the whole principal-level space.

Recall the regular labels c for $\mathrm{GL}_r(k)$: monic squarefree degree r polynomials with nonzero constant term. The Green character χ_c and its central idempotent are written

$$e_c = \frac{\chi_c(1)}{|\,\mathrm{GL}_r(k)|} \sum_h \overline{\chi_c(h)}\, h.$$

If c is not a label the source puts $e_c = 0$. When c is irreducible of degree r, χ_c is the cuspidal Green character of the corresponding regular finite-extension character, and c is primitive when its roots generate $\mathbb{F}_{q^r}^\times$, $q = |k|$. We use the classical finite-general-linear-group character facts: these regular labels give pairwise orthogonal central projectors; for an irreducible label its cuspidal character vanishes unless the matrix characteristic polynomial is a power of an irreducible polynomial; twisting that character by a determinant character amounts to scaling the roots of its label by the corresponding base-field character name. These are Green's character table facts (J. A. Green, *The characters of the finite general linear groups*, 1955), used in `Lib/LocalGL/Green/FrobeniusOrbits.lean` and `New/Trace/BadShare.lean`. The important elementary consequence is the following.

Lemma 13.4 (Cubic avoidance with unchanged nonelliptic coefficients)**.** *Suppose $r = 3^a$ for $a \geq 1$, $9 \mid q-1$, and S, T are two sets of primitive labels, each containing at most one*

element. Given any label c there is a label c' such that the coefficients of e_c and $e_{c'}$ agree at every h whose characteristic polynomial is reducible, and if c is primitive then c' is primitive and outside $S \cup T$.

Proof. If c is not primitive, take $c' = c$. Otherwise let β be a primitive root of c in $\mathbb{F}_{q^r}$, and put $D = q^r - 1$. The three elements $a\beta$ for $a^3 = 1$ in $\mathbb{F}_q^\times$ are all generators of $\mathbb{F}_{q^r}^\times$: express a as $\beta^{jD/3}$, so $a\beta$ has exponent $1 + jD/3$ relative to the generator, which is coprime to D because $9 \mid D$. Indeed every prime divisor of D divides $D/3$. In particular each has degree r over k. Their Frobenius orbits are distinct. If two were equal, β^{q^i-1} would be a nontrivial cube root of unity for some $0 \le i < r$. But $q^i - 1$ is either zero or lies strictly between zero and $D/3$ (use $q > 3$); it cannot be congruent to $D/3$ or $2D/3$ modulo D. Thus these give three distinct primitive labels. They are the cubic determinant twists of the original label, by the classical compatibility just stated. One of them is not in the at most two excluded labels.

At any h of reducible characteristic polynomial where the original cuspidal character is nonzero, that polynomial is P^b for an irreducible P and $b > 1$ dividing r. Therefore $3 \mid b$. Its determinant $(-1)^r P(0)^b$ is a cube, the sign being a cube as well. The cubic determinant character is thus one at h. At all other reducible characteristic polynomials the original character and its twist both vanish. The dimension is unchanged by twisting, so the coefficients of their central idempotents agree at every such h. □

This proves the content of `exists_labelIdem_eq_off_irreducible` used in the source apart from the standard Green character formulas. Notice the distinction between the *label* polynomial of a representation and the characteristic polynomial of a chosen matrix h. The equality in Lemma 13.4 concerns the latter being reducible, not the former.

At a fixed Hecke eigensystem α write $S_w(\alpha, V)$ for the primitive labels occurring at a chosen place w on the principal-level space of the framed group V. This means the corresponding e_c is nonzero on the α eigenspace. The source proves that this set has at most one element for the actual row and single-lift source spaces, not for arbitrary unitary groups. We record why the unusual primitive condition allows weak Galois compatibility rather than full equality of local Langlands parameters.

Lemma 13.5 (A finite rigidity implication)**.** *Let c, c' be primitive labels of degree r over $\mathbb{F}_q$, and suppose a common attached Galois representation has an inertia element whose operator is annihilated by both polynomials obtained from c, c' by adjoining further linear factors. Assume all the further roots have order dividing M, with $q - 1 > M$, under the same faithful multiplicative embedding of the residue algebraic closure into the coefficient field. If the determinant of the operator is the product of the roots of the first enlarged polynomial, then $c = c'$.*

Proof. Two distinct irreducible polynomials of the same degree have disjoint sets of roots. If $c \neq c'$, every eigenvalue of the common operator, since it is killed by both polynomials, must therefore be among roots of order dividing M. Consequently its determinant has order dividing M. On the other hand the product of the roots of c is the norm of a generator of $\mathbb{F}_{q^r}^\times$, and has exact order $q - 1$ in $\mathbb{F}_q^\times$: it is $\beta^{(q^r-1)/(q-1)}$. Multiplying by roots of order dividing M cannot change it into an element of order dividing M, since that subgroup is closed under division. This contradicts the asserted determinant identity. □

In the actual application the primitive roots may first be scaled by a scalar of order dividing $M = M_\omega$. They stay primitive when $M^2 \mid q - 1$: relative to a generator of $\mathbb{F}_{q^r}^\times$ the exponent of such a scalar is divisible by $(q^r - 1)/M$, which in turn is divisible by every prime divisor of $q^r - 1$. Thus the exponent of the scaled generator is coprime to $q^r - 1$, as in the proof

for cubic scaling. Scaling by this same scalar is injective on labels. The further roots of the attached polynomial arise from determinant characters of additional blocks and have order dividing M. The local-global statement used by the source is `Attached.isLabelledAt` in `Lib/GaloisRep/LocalGlobal/LabelCompatibility.lean`; it gives a killing polynomial and the determinant identity at a suitable common tame exponent away from the coefficient prime, not necessarily equality of characteristic polynomials. The tame exponent one is used in `Data.occurringLabels_subsingleton` of `New/GaloisRep/LabelUniqueness.lean`. Weak attachment and the equality of the semisimple attached Galois representation for the two eigenforms follow there from the spherical Hecke system and its density compatibility. The actual injective auxiliary theta lifts and the parahoric intertwining are needed so that this local statement applies to the primitive label on the original space. These geometric and Galois-theoretic inputs are retained at their stated strength; Lemma 13.5 proves the finite implication once they have been established. In particular the condition `HasRigidLabels` for `spaceY0` and `spaceWj` is not a formal consequence of the existence of an abstract Hecke system.

13.8 Equality of all good Hecke traces and determinants

Let $\mathcal{G} = \prod_w \mathrm{GL}_r(k_w)$ be the chosen finite-level group. The good idempotent of the data is

$$e_0 = \sum_{\mathbf{c}\ \text{good}} \bigotimes_w e_{c_w}.$$

A good family consists of allowed regular labels at every chosen place and at least one primitive label above each of the two chosen rational primes. It is an orthogonal central projector. Let V, V' be two of the spaces of principal-level forms of the proof, with the same archimedean parameter and (possibly) different signatures. Denote their actions of the full away-from-chosen-places Hecke algebra $\mathcal{H}$ by $P(a), P'(a)$. In the source this algebra may be larger than the spherical commutative algebra: it consists of families of endomorphisms induced, across the groups, by finite group-algebra combinations of translations trivial at the chosen places (`Data.heckeAlgebra`, `New/Selmer/FitTwists.lean`). It commutes with the chosen finite group action. Its spherical operators at the split places away from the level commute and act normally, so their joint eigenspaces give a direct sum on the finite-dimensional spaces. The following argument does not require every element of the larger algebra to act diagonally.

Let $\mathcal{F}$ be the union of the finitely many occurring joint eigensystems on V, V'. For each $\alpha \in \mathcal{F}$ choose, for each distinct $\beta \in \mathcal{F}$, a spherical Hecke operator $T_{x(\beta)}$ on which they have distinct eigenvalues. The product

$$p_\alpha = \prod_{\beta \neq \alpha} \frac{T_{x(\beta)} - \beta(x(\beta))}{\alpha(x(\beta)) - \beta(x(\beta))} \quad \text{in } \mathcal{H}$$

acts as identity on the α-part and zero on the other parts in both spaces. This explicit interpolation proves the finite simultaneous projection `Data.exists_hecke_proj_eigenPart`. It is harmless that one eigensystem might occur on just one of the spaces.

Assume the rigidity property $|S_w(\alpha, V)|, |S_w(\alpha, V')| \leq 1$. At each w and for each local label c use Lemma 13.4 with these two sets to choose $c'(\alpha, w, c)$. Form the correction

$$E_\alpha = e_0 - \sum_{\mathbf{c}\ \text{good}} \bigotimes_w e_{c'(\alpha, w, c_w)}.$$

For each good tuple, a primitive factor is replaced by one not occurring on the α-parts of *either* space; the corresponding tensor operator kills both parts. Hence on their α eigenspaces E_α

and e_0 have the same action. On the other hand every coefficient of E_α at a tuple of matrices having reducible characteristic polynomials at all chosen places vanishes by Lemma 13.4, applied factor by factor to the corresponding summands. Thus its support consists of matrices having an irreducible characteristic polynomial at at least one chosen place. These are two completely different properties: the first concerns which representations the replacement labels act on, while the second concerns the geometric support of the group-algebra element.

For any $a \in \mathcal{H}$, cyclicity of trace and the fact that a commutes with the chosen-level action yield

$$\mathrm{tr}_{e_0 V}(P(a)) = \mathrm{tr}_V(P(a)e_0) = \sum_{\alpha \in \mathcal{F}} \mathrm{tr}_V(P(a p_\alpha)E_\alpha),$$

and likewise for V'. In the first equality the trace on the full space is that of the action on the image of the projection, with zero on its kernel. For the second one, sum the projections p_α to the identity and then use the equality of E_α and e_0 on that part; both commute with p_α. No commutation of a with all p_α is needed for this formula.

Each operator on the last sum is a finite linear combination of translations by a product of an away-from-chosen-places adelic element and a lift of a tuple in the support of E_α. The full right coset by the principal level has irreducible characteristic polynomial at one of the chosen places where the tuple does. Indeed its reductions there have the same monic irreducible characteristic polynomial, so the lifted characteristic polynomial is irreducible over the local field by reduction of monic integral factors (Gauss's lemma over the discrete valuation ring). The away operator is trivial there and cannot affect this fact. Thus every summand may be treated by the elliptic trace comparison (13.5). We obtain

$$\mathrm{tr}_{e_0 V}(P(a)) = \mathrm{tr}_{e_0 V'}(P'(a)) \qquad (a \in \mathcal{H}). \tag{13.6}$$

This is the argument of `Data.trace_heckeOnGood_eq` in `New/Trace/GoodFormTransfer.lean`. It proves the finite projection step exactly, assuming the explicit hypotheses of the geometric comparison and the actual rigidity assertion, rather than making a stronger unsupported transfer claim.

Taking $a = 1$ proves equality of dimensions. Taking all powers of a in (13.6) proves equality of the characteristic polynomials (by Newton's identities over $\mathbb{C}$) and hence

$$\det_{e_0 V}(P(a)) = \det_{e_0 V'}(P'(a)). \tag{13.7}$$

In particular if a joint spherical eigensystem is absent from one good space it is absent from the other: interpolate the projector onto its part on the union of the finite eigensystems and take its trace, which is the dimension of that part. These implications reproduce `det_heckeOnGood_eq` and `goodForms_inf_eigenPart_eq_bot`. The determinant equality, not a mere numerical dimension comparison, is also used to compare the Rallis operators of the row chain and the auxiliary single lifts in the lower bound.

13.9 The precise row-dimension inequality

We finally apply these comparisons to `Data.finrank_rows_le_mass`. For the row group the doubled parameter at $\tau \in \Phi$ is

$$l_\tau(i) = 2\,\varsigma_\tau(\mathrm{basisNumber}(i)),$$

where ς_τ is the row parameter in the data, a strictly decreasing sequence. Thus the decreasing position of i in l_τ is exactly basisNumber(i) (the number of entries larger than it), independent

of the ordering chosen for the basis indices. This is `decrPosition_paramW` in `New/Filling/RowDimension.lean`. If the corresponding dominant diagram is κ_τ with central shift e_τ, the Harish–Chandra parameter associated to that highest weight is

$$l_\tau^{\downarrow}(i) = 2(e_\tau + \kappa_\tau(i)) + r - 1 - 2i, \qquad 0 \le i < r.$$

Equivalently $\kappa_\tau(i) + e_\tau$ is $l_\tau^{\downarrow}(i)/2 - \rho_i$ with zero-based $\rho_i = (r-1-2i)/2$. This recovers `hcParameter` and `IsHCParam.lowestKType_eq` in `Lib/Adelic/CuspForms/WeightParameter.lean`. The compact tensor coefficient representation has dimension

$$d_\infty = \prod_{\tau \in \Phi} \dim V_{\mu_\tau} = \prod_{\tau} d_{l_\tau}.$$

The last equality is the exact compact-dual degree formula of Proposition 13.1. The local Bessel system in the source is required to use these same degrees at each real place (the parameter-matching component of `D.filling_fits`); thus d_∞ equals the number called `Filling.besselDegree`, and not a multiple of that number.

For the definite group with these compact real representations, forms of the minimal real type correspond to algebraic modular forms with this tensor coefficient representation. One can see the isomorphism directly: if v_μ is the unit highest vector in that irreducible representation and f is a vector-valued finite automorphic form, the scalar coefficient at (x_∞, x_f) is obtained by pairing $f(x_f)$ with the translate of v_μ by x_∞ in the appropriate contragredient convention. The rational automorphy follows from the transformation of f by the coefficient representation, and compact Schur orthogonality gives the reproducing identity. Conversely integrate the scalar form against these translates of the highest vector and multiply by d_∞ to recover the vector-valued form. The two maps are inverse by Schur orthogonality and cyclicity of the highest vector; the latter follows because its orbit linearly spans a nonzero invariant subspace of the irreducible representation. Finite-level invariance is preserved. This is the compact equivalence `Space.algFormEquiv` in `Lib/Adelic/CuspForms/AlgebraicForms.lean`.

At the chosen principal level in rank r all rational stabilizers are trivial by the neatness calculation; the index into the standard finite maximal group is $|\mathcal{G}|$. The mass of its finite class groupoid is $|\mathcal{G}|\mathfrak{m}_r$, by the level-index identity proved in Chapter 6, and since the stabilizers are trivial this is exactly the number of classes. An algebraic form may be assigned an arbitrary coefficient vector independently at each such class. Thus the dimension of the full definite principal-level space equals

$$|\mathcal{G}|\,\mathfrak{m}_r\,d_\infty.$$

Its good subspace can have smaller dimension. By (13.6) for $a = 1$, the dimension of the actual row good space W equals the dimension of that definite good subspace. Therefore

$$\dim_{\mathbb{C}} W \;\le\; |\mathcal{G}|\,\mathfrak{m}_r\ \mathrm{besselDegree}\,. \tag{13.8}$$

This is precisely the inequality stated in the source as `Data.finrank_rows_le_mass`. In the notation of the spanning chapter it is the bound by the mass times the local formal degrees times the number of level translates. No factor of dimension of the lowest K_∞-type appears: the chosen archimedean projection is the rank-one matrix coefficient, as explained above. Nor does one need equality of the full indefinite and definite principal-level spaces.

To summarize the remaining dependencies of this particular chapter: the elementary coordinate measure, matrix integral, compact degree calculation and finite projection arguments are independent conventional proofs. The formal-degree proportionality and discrete-series elliptic character identity are classical real representation-theoretic inputs; the precise match of the actual source kernel to them, with its uniform decay and reproducing property, must

be verified for the constructed row parameters. The weak Galois attachments, local label compatibility and actual auxiliary injective lifts supplying the rigidity hypothesis must also be established. They are used only at the stated interfaces here. Once those inputs are supplied, the signature comparison and exact dimension inequality follow by the arguments above without introducing an additional global dimension-transfer assumption.

Chapter 14

Quantitative spanning by projected Bessel–Poincaré series

Chapters 15 and 16 supply the kernel, Schur, Laplace, finite-volume, Poisson, real-peak and sieve arguments underlying the estimates used here. Statements that also require uniform geometric input retain that input unless it is proved in these companion chapters; the elementary rank deduction is independent of the details of this analysis.

The preceding bound for the columns is useful only if the resulting rows span a sufficiently large space. A merely qualitative nonvanishing or spanning theorem does not suffice: the integral-index bound is proportional to the number of translates used. We give the quantitative statement, develop the relative trace calculation on which it rests, and prove the finite-dimensional deduction. This section separates out the analytic and arithmetic estimates that are specialized to the large ranks, levels and weights of this proof. The term *Formal-source input* below denotes a result extracted from the indicated formal declaration; it is not asserted to have been independently reproved here. We retain the actual constants when they are used in another estimate. The Hilbert space arguments and the finite-field sieve deductions are given directly.

14.1 Statement and finite-dimensional framework

Let C denote the column space of algebraic modular forms, C^{ram} its subspace ramified at the chosen auxiliary prime, W the row space, and $\mathfrak{G}$ the finite group of level translations. Put $c = \dim_{\mathbb{C}} C$ and $n = \dim_{\mathbb{C}} W$. The maps $P_s : C^{\mathrm{ram}} \to W$ are the projected translates defined above, before the common rational rescaling of Theorem A. This rescaling has no effect on the image of any map.

Theorem 14.1 (Quantitative spanning, Theorem B)**.** *For admissible data $\mathcal{D}$ we have $c, n > 0$. There exists a finite subset $T \subset \mathfrak{G}$ for which*

$$|T|c \le 5n \log(1 + b^2), \qquad \dim\Big(\sum_{s \in T} P_s(C^{\mathrm{ram}})\Big) \ge (1 - 4b^{-2})n.$$

This is `Data.enough_forms` in `New/Endgame/ShortAndEnough.lean`. We will deduce it from a local-to-global moment calculation and a finite-dimensional variance inequality, rather than take `Data.enough_forms` as an input. Notice the first inequality uses c, the dimension of the entire column space: the actual index estimate of A costs at most cB per translate. It is not enough to bound $|T| \dim C^{\mathrm{ram}}$.

For an arbitrary subspace $J \subset W$, with fraction $\alpha = \dim J/n$, the operative intermediate result is

$$\dim(P_s D_N)^{-1}(J) \le ((1+\alpha)/2 + b^{-2})c \quad \text{for some } s \in \mathfrak{G}. \tag{14.1}$$

Here $D_N : C \to C^{\mathrm{ram}}$ is the ramifying operator of the previous section, viewed as a map onto its image. When J has dimension appreciably less than n, the dimension of the new image modulo J is correspondingly large. There are two analytical tasks: construct the original Poincaré map on C, and control its translates after discarding bad finite types. The ramification operator imposes only a small additional cost on the source rank.

14.2 Finite level, central projection, and the transfer bound

Write E^+ for the maximal totally real subfield of E. The row forms before projection belong to a finite-dimensional Hilbert space V of cusp forms on $U(Y_0)$ of the prescribed archimedean type and of principal level at the chosen split places of g_0, g_1. Its inner product is the Petersson pairing. The group of finite level translations is

$$\mathfrak{G} = \mathfrak{G}_0 \times \mathfrak{G}_1, \qquad \mathfrak{G}_i = \prod_{v \mid g_i \text{ chosen}} \mathrm{GL}_r(\mathbb{F}_v),$$

acting via a unitary representation ρ. In each product there are precisely f factors with residue cardinality g_i. Conjugate places are omitted from the indexing because the local unitary condition recovers the data at the conjugates. For an irreducible complex character ξ of a finite group G the element

$$e_\xi = \frac{\xi(1)}{|G|} \sum_{g \in G} \overline{\xi(g)}[g]$$

is a self-adjoint central idempotent. If $e = \sum_{\xi \in F_0} e_\xi$ for a set of irreducible characters, its coefficient at the identity is the Plancherel mass $\mu(F_0) = \sum_{\xi \in F_0} \xi(1)^2/|G|$. Consequently $\mu(F_0)$ is the fraction of the regular representation occupied by these isotypic spaces; it need not be the analogous fraction of V. Distinguishing these assertions is essential in the filling proof, which bounds the trace of the discarded part on the particular image of the Poincaré series, rather than deducing it from the discarded proportion in an arbitrary representation.

At a chosen place with residue field $\mathbb{F}_q$ the regular labels used in the proof are monic squarefree polynomials of degree r with nonzero constant term. Let M_ω denote the excluded torsion order, as in the definition of the level characters. A label is allowed if it has no root $x \in \mathbb{F}_q$ of $x^{M_\omega} = 1$. A primitive label is an irreducible polynomial with one (and hence all) roots of multiplicative order $q^r - 1$. The good tuples are tuples of allowed labels for which at least one label is primitive. The definitions are `IsPrimitiveLabel` and `goodTuples` in `New/Trace/BadShare.lean`; the corresponding idempotents are `labelIdem`, `tupleIdem` and `Data.e0` in the finite character theory and `New/Seesaw/IndexTypes.lean`. Let e_i be the good central projection in $\mathbb{C}[\mathfrak{G}_i]$ and $R = \rho(e_0 \otimes e_1)$. (Here e_0 denotes the first factor's projection; the code uses `Data.e0` for their product.) Our row space is $W = \operatorname{im} R$. The two factors are both used: the goodness requirement contains a primitive label above *each* g_i.

Lemma 14.2 (Plancherel deficit)**.** *Let $1 < r < q$ and $M_\omega \ge 1$. Suppose $0 \le \varepsilon \le 1$ and $L_0 \ge 0$ satisfy*

$$(1 + M_\omega + 2r)/q \le \varepsilon, \qquad 2r(2r \log q + 1)L_0 \le f.$$

Then the coefficient at the identity of $1 - e_i$ is at most $f\varepsilon + e^{-L_0}$. In particular, for our data, the complementary coefficient is at most $b^{-8}/2$ for each i.

Proof. The representation-theoretic input is the degree estimate of `labelChar_degree_bounds` for the Green characters, used in `bad_share_le`: each regular label contributes Plancherel mass at least

$$\mu_0 = (1 - r/q)q^{-r}.$$

We take this degree formula as the *Formal-source input* from the finite general linear group character theory (see `New/Trace/BadShare.lean` for the use of this bound). The rest is counting. There are q^r monic polynomials of degree r. For each of at most $1 + M_\omega$ prohibited linear roots (including zero), there are q^{r-1} polynomials divisible by the corresponding linear factor. The nonsquarefree polynomials belong to the union of the families Q^2P, where Q, P are monic and $1 \le \deg Q \le r/2$, $\deg P = r - 2\deg Q$. For fixed degree j there are at most q^{r-j} choices, so a crude bound for the total is rq^{r-1}. Hence allowed labels have total uniform lower mass at least

$$(1 - r/q)(1 - (1 + M_\omega + r)/q) \ge 1 - \varepsilon.$$

Primitive roots are precisely the $\varphi(q^r - 1)$ generators of the multiplicative group of the degree-r field. Such a generator is not contained in a proper subfield, so each of their minimal polynomials is irreducible of degree r and corresponds to exactly r generators. The number of primitive labels is therefore $\varphi(q^r - 1)/r$. The elementary totient lower bound used in the source is $\varphi(u) \ge u/(2\log u + 1)$ for $u \ge 2$ (`le_totient_mul_log`). Thus the uniform mass q^{-r} assigned to these labels is at least $1/[2r(2r\log q + 1)]$. The mass of nonprimitive labels with respect to the weight μ_0 is at most their mass with respect to q^{-r}, whence it is at most $1 - L_0/f$. On the product of f factors, the mass lower estimate for all allowed labels is $(1-\varepsilon)^f$, and the uniform lower mass of the allowed tuples that are not primitive anywhere is bounded above by $(1 - L_0/f)^f$. Thus the mass of the good tuples is at least $(1-\varepsilon)^f - (1 - L_0/f)^f$. Subtracting from one, and using Bernoulli's inequality and $(1-u)^f \le e^{-fu}$, proves the claimed bound.

For the last specialization one takes $L_0 = 9\log b$. The size inequalities used in `Data.ramifiedColumns_fill_rows`, together with $g_i \ge b^{10}$ and f sufficiently smaller than the level, give $f\varepsilon + b^{-9} \le b^{-8}/2$. These exact numerical inequalities on the data are the *Formal-source input* `bad_share_le` as instantiated in `New/Filling/FillingTheorem.lean`. The counting proof explains the two separate contributions: excluded labels are individually rare, while failure to contain even one primitive label is exponentially rare across the many factors. □

The second input is a bound on the actual number of target rows. Let $d_\infty = \prod_{v\in\Phi} d_v$ be the product of the local formal degrees (defined through the Bessel systems below) and let $\mathrm{mass}_r(E)$ be the mass of the rank-r positive definite unimodular genus, with finite stabilizer weights. Define

$$n_0 = |\mathfrak{G}| d_\infty \,\mathrm{mass}_r(E). \tag{14.2}$$

Proposition 14.3 (Dimension transfer (conditional derivation in Chapter 13))**.** *One has $n \le n_0$. More precisely the trace of every element of the common split unramified Hecke algebra on good forms of the same rank, level and compatible archimedean Harish–Chandra data is independent of the unitary form, in the rigid-label situation of this proof.*

The statements are `trace_heckeOnGood_eq`, `det_heckeOnGood_eq` and `finrank_rows_eq_finrank_goodForms` in `New/Trace/GoodFormTransfer.lean`, and `Data.finrank_rows_le_mass` in `New/Filling/RowDimension.lean`. For the dimension computation, the primitive-label conditions and the local rigidity lemma in `New/Trace/LabelRigidity.lean` are needed in the trace formula, as well as the existence of projectors onto the finite Hecke eigenspaces. The proof transfers the row forms to the definite member of rank r, where good forms are a subspace of the full space of algebraic modular forms. Here the class stabilizers

are trivial at the chosen level. The class sum weighted by the reciprocal group orders is $|\mathfrak{G}|\,\mathrm{mass}_r(E)$ by the level-index identity; the dimension of the finite-dimensional algebraic representation at infinity is d_∞ in the transfer normalization. Multiplying gives (14.2). The nontrivial passage between signatures is not a formal consequence of dimension equality and cannot be omitted; Chapter 13 proves the trace comparison and dimension deduction from its explicitly stated elliptic orbital and finite rigidity hypotheses, including the precise formal-degree normalization. The same trace equality is used again to compare auxiliary lifts with the main chain in the lower-bound argument.

14.3 Hermitian geometry of the Bessel subgroup

We describe the objects for the unfolding calculation. In this section use $h = m/2 - 1$ and $s_0 = h + 1$ (not the dimension n of row forms). The admissible rank $r = k + 2h + 1$ is the rank of the split extension of the definite rank-k space X_0. It has ordered basis of h hyperbolic pairs e_i, e_i^* $(0 \le i < h)$, an orthonormal basis of X_0, and one middle vector z with norm $-N$. Thus the Gram matrix in the rational basis is the sum of the hyperbolic planes, I_k, and $(-N)$. Denote this form by Y. Let U be the unipotent radical of the stabilizer of the isotropic flag spanned by successive initial e_i, with its standard Bessel character, and let H_0 be the subgroup isomorphic to $U(X_0)$ acting as the identity on the frame and its dual. The Bessel subgroup is $S_B = UH_0$. Its frame is $F = (e_0, \dots, e_{h-1}, z)$; U^0 is the pointwise frame stabilizer in U. The coordinate form of these groups, including the flag, is the `Shape` construction of `New/Filling/EichlerFlag.lean` and the groups in `Lib/Adelic/Poincare/BesselSubgroup.lean`.

At finite places put $N = \pi\bar{\pi}$, interpreted as an idelic factorization if necessary, and modify the basis of $X_0 \oplus Ez$ by replacing its selected vector e_* with $v_* = e_*/\bar{\pi} + z/N$. Its self-pairing is $1/N - 1/N = 0$ and its pairing with z is -1; the 2×2 middle Gram block becomes

$$\begin{pmatrix} 0 & -1 \\ -1 & -N \end{pmatrix}$$

with determinant -1. All other middle basis vectors retain their original pairings; the hyperbolic planes are unchanged. Thus the lattice Gram matrix is unimodular. These are the formulas of `latMiddleGram` and `gram_besselLatticeBasis` in `New/Filling/EichlerFlag.lean`. Let K_f be the stabilizer of this lattice away from the infinite places, and $K(g)$ the kernel of reduction to the above finite group. The column level is the inverse image of $K(g)$ under inclusion of H_0; this is one of the hypotheses in the filling record. Let $\mathcal{C} = H_0(E^+)\backslash H_0(\mathbb{A}_f)/K_c$ and h_u be the chosen images of representatives of classes $u \in \mathcal{C}$.

Proposition 14.4 (Formal-source input: global filling datum)**.** *The admissible data give a filling system in the sense of* `Filling` *and* `Filling.IsAdmissible` *of* `New/Filling/Setup.lean`*, with the prescribed r, k, h, N, g_i, b, whose column stabilizers are trivial, reduction onto $\mathfrak{G}$ is surjective, and whose Bessel coordinate spaces at each infinite place identify with the column weight representation. The resulting column coordinates give an isomorphism*

$$C \simeq \bigoplus_{u \in \mathcal{C}} \mathbb{C}^I, \qquad I = \prod_{v \in \Phi} I_v.$$

In particular $c = |\mathcal{C}||I| > 0$.

The coordinate identification is `Data.columnCoords` in `New/Seesaw/BesselGeneric.lean` (the coordinate model is built in `New/Filling/ColumnCoords.lean`); the filling construction

is `New/Filling/FillingData.lean` and the admissibility check is `New/Filling/FillingAdmissible.lean`. Positivity follows since the class of the identity and every weight coordinate set are nonempty. The dimension formula also follows directly from the double-coset description when stabilizers are trivial. For the local-to-global estimates below we record several of the other numerical hypotheses, rather than replace them by unspecified implied constants. If ϖ is the purely imaginary generator of the absolute different and $\kappa_v = \Re(gv(\varpi)/(2\pi i))$, then

$$1 \le |\kappa_v| \le b^{23}, \quad |v(\varpi)| \le b, \quad N \le b^{60 s_0}, \quad b^{10} \le g_i \le b^{11}.$$

Also $h \ge 10^5$, $1000(2h+2)^9 \le k \le f$, $s_0^7 \le k$ and $(200 f k \operatorname{rd}_E)^{1000} \le b$. These are much stronger than the hypotheses of the elementary rank argument. They enter below in estimates which would be false uniformly for arbitrary families of lattices of varying dimension.

The normalized archimedean basis at v is scaled by $\kappa_v^{h-i}/\sqrt{N}$ on e_i, its inverse on e_i^*, and $1/\sqrt{N}$ on z. On X_0 the basis is unchanged. The middle norm is thereby normalized to -1. Put

$$l_{vi} = \frac{\kappa_v^{2(h-i)}}{N} \quad (0 \le i \le h), \qquad \mathcal{L} = \prod_{v \in \Phi} \prod_{i=0}^{h} l_{vi}^{k+2(h-i)}.$$

For a rational element γ of $U(Y)$ define the frame-product matrix $Z(\gamma) = (\langle F_i, \gamma F_j \rangle)_{0 \le i,j \le h}$, and write $D_i(\gamma)$ for its leading principal minor of order i. The discriminant separating the open and smaller cells is

$$\mathfrak{d}(\gamma) = D_{h+1}\overline{D_{h+1}} - N^2 D_h \overline{D_h}. \tag{14.3}$$

The normalized archimedean matrix Z_v has entries $\sqrt{l_{vi} l_{vj}}\, v(Z(\gamma)_{ij})$ up to signs of the real scales (the scales themselves are chosen real, not assumed positive); its normalized discriminant is $|\det Z_v|^2 - |D_h(Z_v)|^2$. Direct expansion shows that it equals $(\prod_i l_{vi})^2 v(\mathfrak{d}(\gamma))$, since the ratio between the square of the products of the full scalings and of the first h scalings is $l_{vh}^2 = N^{-2}$. The definitions are `Filling.frameProducts`, `frameScaleSq`, `termDisc` in `Setup.lean`, and `frameDisc` in `EichlerFlag.lean`. The nonnegativity of this discriminant for actual frame products and the strict inequality in the full open cell are geometric consequences of the positive definiteness of the complement of the hyperbolic planes, recorded in `NearTerms.lean`.

This particular frame is relevant both to the archimedean kernels and to counting at the finite level: the leading minors act as pivots. It is important that one does not have to count all pairs of rational points of the full rank-r group indiscriminately. Coherent pairs are sorted by common frame data and by a linear relation between the two rank-k middles; then the deepest strata admit a local orbital calculation. The smaller strata are controlled by means of the same discriminant and congruences.

14.4 The archimedean Bessel system

For each $v \in \Phi$ the group at the normalized real place is the unitary group $U(k+h, h+1)$ in the split basis just described. Denote it by G_1. We use Haar measures in the normalization of `BesselSystem.lean`; in particular the measure on $U\backslash G_1$ is the quotient by the specified unipotent measure and is not a probability measure. The compact type of $H_0(\mathbb{R}) \simeq U(k)$ is identified with the prescribed column representation (of dimension $|I_v|$). The archimedean parameter of the target representation has large gaps: the local construction uses a gap at least r^7 and each doubled parameter has absolute value at most $3a$. The finite-dimensional column type has weight $(\kappa_v^{\mathrm{col}}, 0, \dots, 0)$ in the convention of the proof: its only potentially nonzero entry is the one corresponding to the marked first column of order a at the tops. This weight should not be confused with the real scaling constant κ_v just introduced.

The representation is described in Lean via a reproducing kernel rather than via a choice of realization of a discrete-series Hilbert space. An invariant kernel of formal degree d_v, acting by convolution, reproduces the functions of its type. Its integral against the Bessel character constructs a relative kernel; the associated map on the compact type is normalized to have prescribed L^2 norm. The source constructs a full function and a cutoff with unusually strong uniform bounds despite the changing representation and rank. Here are the exact properties that will be used.

Definition 14.5 (Quantitative Bessel system)**.** For a real place v, let I_v be a nonempty finite set, let $D_v = |I_v|$, and let $d_v > 0$. Let W_v (full) and $W_v^\flat$ (cut) be continuous functions on G_1 with values in $\mathbb{C}^{I_v}$, equivariant on the left under U through the fixed character and under H_0 through the specified irreducible unitary representation. Define the relative size

$$\mathrm{bs}(y) = \inf_{s \in S_B} (1 + \|sy\|_{\mathrm{op}}).$$

The following conditions are part of the quantitative Bessel system:

1. The full function is integrable on $U\backslash G_1$, its squared L^2 norm is D_v, and the cut has squared norm at most D_v and vanishes where $\mathrm{bs}(y) > b^{6s_0+2}$. Moreover
$$\|W_v^\flat - W_v\|_2^2 \le b^{-10r^2} D_v.$$

2. There are a measurable left U-invariant majorant A and a finite measure μ on G_1 such that each of the two norms is bounded by A, and in the region where the functions differ their norms are bounded by $b^{-3r^3 f} A$. In addition
$$A(y) \le \mu\{g : \mathrm{bs}(yg) \le b^{6s_0+2}\}, \qquad \mu(G_1) \int_{U\backslash G_1} A(y)\, dy \le b^{2r^3} D_v.$$

3. For each word α in the matrix-coordinate directions, with $|\alpha| \le 2s_0^2$, there is a continuous derived Bessel function $W_{v,\alpha}$. Its coordinate right differentiation relation is the ordinary chain rule along a curve $\gamma'(t) = \gamma(t)X$, namely the derivative of $W_{v,\alpha}(\gamma(t))$ is the sum of the coordinates X_{ij} times $W_{v,(i,j)\alpha}(\gamma(t))$. For every $P \ge b^5$ and every integrable scalar ϕ on G_1, the derived Schur bound is
$$\int_{U\backslash G_1} \left\| \int_{G_1} \phi(y) W_{v,\alpha}(\gamma y)\, dy \right\|^2 d\gamma \le P^{2|\alpha|} \frac{D_v}{d_v} \|\phi\|_2^2. \tag{14.4}$$

4. The Bessel correlation of the cut, after the changes of coordinates below, belongs to a tower of kernels with explicit Fourier bounds, called a sieve system. Its bottom constant is $(1 + b^{-10})/d_v$.

These are precisely `IsCutPair`, `HasSchurBound`, `BesselSystem` and `HasBandLimits` of `New/Filling/BesselSystem.lean`. The word "derived" refers here to differentiation, not to derived categories. Condition (2) is stronger than L^2 approximation alone: after summing translates over a growing lattice, a small L^2 error does not automatically imply a small contribution to the second moment.

Proposition 14.6 (Formal-source input: existence of local systems)**.** *There exist quantitative Bessel systems with the prescribed formal degrees, type, and column-coordinate representations at all $v \in \Phi$. The cutoff system satisfies the Fourier bounds of* (14.6) *below, and the open-cell estimate of* (14.7) *below.*

The existence assertion is `Data.exists_besselSystem` in `New/Filling/FillingData.lean`, obtained from `HasPeakAt.exists_cut_isCutPair` and `IsBesselFunction.hasSchurBound_of_convolve_eq` in `New/RealPlace/BesselCut.lean` and `hasBandLimits_of_isCutBessel` in `New/RealPlace/BandLimits.lean`. The local construction obtains a peak of a Bessel convolution by lifting between unitary pairs with separated parameters; localization and lower bounds on the peak are what permit normalization by the peak value. The gap and the bound on the highest parameters are used both in the tail of the kernel and to control iterated right derivatives. For example the derivative estimate for a reproducing kernel of rank r, highest doubled parameter bounded by A_0, permits a Schur constant P at least $(r+1)^4(A_0+r+1+2s_0^2)$; the admissible data absorb this in b^5. In reproducing-kernel notation the reason for the Schur estimate is that convolution is an orthogonal projection up to the factor d_v. After normalization and differentiation, its operator norm is bounded by $d_v^{-1/2}$ times the norm of the differentiated vector. The nontrivial uniform part is bounding those vectors and obtaining the stated peak and cutoff, parts of the above formal-source input; neither follows merely from the existence of a discrete-series representation.

Lemma 14.7 (Orthonormality of the full coordinates)**.** *For each v and all $i, j \in I_v$,*

$$\int_{U\backslash G_1} W_{v,i}(y)\overline{W_{v,j}(y)}\,dy = \delta_{ij}.$$

Proof. The integral matrix $T = \int W_v(y)W_v(y)^* dy$ is a Hermitian positive semidefinite operator on the coordinate space. Left invariance of the quotient measure under the compact group H_0 and Bessel equivariance show that T commutes with its irreducible unitary representation. Schur's lemma gives $T = aI$ with $a \geq 0$. Taking traces and using condition (1) gives $aD_v = D_v$, so $a = 1$. This is the argument of `BesselSystem.orthonormal` in `New/RealPlace/BesselOrthonormal.lean`. □

We record the sharper local correlation bounds because they fix the constant in the global moment. For any vector-valued W on G_1 let

$$J_W(\gamma) = \int_{G_1} W(y) \otimes W(\gamma^{-1}y)^*\,dy$$

when this integral is defined. For $W = W_v^\flat$, the full open-cell Weyl representative with head $p = (p_0, \ldots, p_h)$ is denoted w_p. Here $p_i \neq 0$ are the successive pivot ratios in elimination of the frame-product matrix; for $i = h$ we have $p_h = \langle z, \sigma z\rangle$ for the residual middle element $\sigma \in U(X_0 \oplus \mathbb{C}z)$. Let $V(U(j))$ be the volume of the compact unitary group in the source normalization and put

$$c_G = \frac{V(U(k))}{V(U(r))}\frac{2\pi^k}{(k-1)!}, \qquad c_{\mathrm{band}} = \frac{V(U(k))}{V(U(r))}\pi^{(2h+1)k}.$$

For matching top entries $|p_h| > 1$ define the normalized correlation tensor

$$T_v(\sigma, p) = \frac{\sqrt{c_G}\,(|p_h|^2-1)^{(k-1)/2}}{D_v}\prod_{i<h}|p_i|^{k+1+2(h-i-1)}\,J_{W_v^\flat}(w_p\sigma). \tag{14.5}$$

The normalization of the tensor has a particular Jacobian meaning: the powers in the product are the half-densities of the successive rank-one steps in the open cell. In the last step the residual middle group contributes the volume of a sphere of dimension $2k-1$ and the radial factor $(|p_h|^2-1)^{k-1}$. This is the real factor that will compensate the global finite-place quadric count.

The top kernel of the local sieve system is the inner product of the tensors (14.5), when the two top entries both come from σ. Hence the top kernel is positive semidefinite for common top entry. The frequency bound for the system has the following explicit majorant. Write

$$\lambda_a(u) = u^a \min(1,u)^6, \qquad \Lambda_i(p) = \prod_{j<i} \lambda_{s_0^6}(|p_0 \cdots p_j|^{-1}).$$

In the integral of one common entry at level i the nonzero Fourier coefficient, at frequency ξ, is bounded in absolute value by

$$\beta_v(\xi)\Lambda_i(p)\Lambda_i(q), \qquad \beta_v(\xi) = c_{\mathrm{band}} b^{2000 s_0^9} \lambda_{k-s_0^7}(|\xi|^{-1}). \tag{14.6}$$

The same bound in the unequal-entry relation (with the next head weights) is used below. These are the actual constants in `HasBandLimits`. The large exponent $k - s_0^7$ is critical; a bound merely asserting rapid decay for fixed parameters would not provide the required powers after multiplying over f places. The exponent 6 in $\min(1,u)^6$ is a separate damping term used for summability of frequencies after applying the product formula. The construction of these kernels and the bounds (including the identity of the top kernel with the correlation) is the formal-source input `hasBandLimits_of_isCutBessel` and the preceding real-place calculations in `New/RealPlace/BandLimits.lean` and `LocalTheorem.lean`.

A second local estimate is used for the terms that cannot be treated by the sharp main-term calculation. Let $Z_v(\gamma)$ denote the real-place frame-product matrix and let $\mathfrak{d}_v = |\det Z_v|^2 - |D_h(Z_v)|^2 > 0$. If $\rho_0 \geq b$ and the box in which all frame coordinates have norm at most ρ_0 meets its translate by γ, define

$$M_v(\rho_0,\gamma) = \left(\frac{\rho_0^{9s_0}}{\mathfrak{d}_v}\right)^{10^5 s_0^4} \frac{D_v^2}{d_v c_G (\mathfrak{d}_v/2)^{k-1}}. \tag{14.7}$$

The open-cell estimate (*Formal-source input*: `BesselSystem.open_cell` in `New/RealPlace/OpenCellBound.lean`) asserts the existence of a compact integration set, an averaging parameter in $U^0 H_0$, and a measure of total mass at most M_v, such that the averaged squared norm of the integrals against $W(\gamma ty)W(ty)^*$ is at most the integral of the squared modulus of the test function against this measure. The compact integration set contains the required translates of the box, almost surely in the averaging parameter. This is a local operator bound with an explicit discriminant pole, not a pointwise estimate for W. The normalization in (14.5) is intentionally different: in the sharp term one preserves cancellation between Fourier phases, while in (14.7) one uses upper bounds allowing losses. The large exponent on $\mathfrak{d}_v$ is compensated globally by discriminant bounds and congruence counts.

14.5 Construction and exact unfolding of the Poincaré map

Using the coordinate isomorphism of Proposition 14.4, endow C with its orthonormal coordinate basis $e_{u,i}$ for $u \in \mathcal{C}$, $i \in I$. Write $W(y)_i = \prod_v W_{v,i_v}(y_v)$ and similarly $W^\flat$ for the cut. The superscript $\flat$ now denotes the approximate function, not the auxiliary cutoff on unipotent cosets. The product Bessel functions live on the product of the normalized infinite groups, denoted G_∞. For any rational element γ, use γ_∞ and γ_f for its images in these normalized coordinates and in the finite lattice coordinates, respectively.

For $u \in \mathcal{C}$, let $U_u(g)$ be the rational unipotents with $h_u^{-1}\gamma_f h_u \in K(g)$. The image at infinity is a discrete group. Take a measurable nonnegative class cutoff β_u on G_∞ with

$$\sum_{v \in U_u(g)} \beta_u(vy) = 1. \tag{14.8}$$

Such a function may be obtained by taking a measurable fundamental domain, or by normalizing a nonnegative locally finite function that is positive somewhere on each orbit. The choice in the source is expressed as a cutoff for a discrete subgroup. The arithmetic assertion of discreteness, and the existence of a cutoff with the specified compatibility, is `exists_cutoff_classUnip` in `New/Filling/Setup.lean`. The measures are chosen so that the covolume of $U_u(g)$ in the infinite unipotents is a constant $\nu > 0$ independent of u. Explicitly, put $d_U = h(h+k)$ and $D_E = |\operatorname{disc}(E)|$. With $v_f = \operatorname{vol}(\widehat{\mathcal{O}}_E)$ in the self-dual finite measure of the source, the covolume is

$$\nu = \frac{(g^{2d_U+h}/N^k)^f}{v_f^{d_U}} \frac{D_E^{h/4}}{\mathcal{L}}. \tag{14.9}$$

This is `Filling.unipCovolume`. To see the origin of the exponent, the unipotent filtration has a middle step of size $k+1$ and two strings of sizes i, i for each $0 \le i < h$; the sum of their sizes is $\sum_i (2i+k+1) = h(h+k)$. Each such adelic coordinate has congruence density and a self-dual covolume, while the normalized archimedean change of basis supplies the Jacobian $\mathcal{L}$. The norm term N^{-kf} occurs in the middle lattice modification. The exact quotient formula (14.9), including independence of u, is the *Formal-source input* `Filling.lintegral_classCutoff_eq_unipCovolume` of `New/Filling/UnipotentCovolume.lean`; the numerical factors are those of `Setup.lean`. For an integrable function F on the unipotent quotient, the cutoff identity implies

$$\int_{G_\infty} \beta_u(y) F(y)\, dy = \nu \int_{U_\infty \backslash G_\infty} F(y)\, dy$$

provided F is invariant on unipotent orbits: integrate first over the left unipotent and then use (14.8) to replace that integral by its quotient covolume. This is the quotient calculation used repeatedly in the unfolding.

For (u, i) let $\phi_{u,i}$ be the seed given by $\beta_u(y) W_i(y)$ at infinity times the indicator that $h_u^{-1} g_f \in K(g)$, in the common coordinates. A fixed common adelic coordinate change is understood (the map called `Filling.Fits.finiteEquiv` and the inclusion of its fixed η in `New/Filling/PoincareSeries.lean`); it has no effect on the formulas for relative positions since it cancels from them. The Poincaré map $Q : C \to V$ is the periodization

$$Qz(g) = \sum_{u,i} z_{u,i} \sum_{\gamma \in G(E^+)} \phi_{u,i}(\gamma g). \tag{14.10}$$

When summing first over unipotents, the periodization replaces the cutoff β_u by 1 on the permitted finite coset, by (14.8) and the Bessel character being trivial on $U_u(g)$. This identifies (14.10) with the Poincaré series on the Bessel quotient; note that the source function called `cutBessel` in this construction contains the full W and the cutoff β_u, not the cut $W^\flat$.

To state the convergence and unfolding precisely, put

$$\operatorname{pos}_{u,u'}(\gamma) = h_u^{-1} \gamma_f h_{u'}, \qquad \mathcal{E}(u, s, u') = \{\gamma : \operatorname{pos}_{u,u'}(\gamma) \in K_f, \operatorname{red}(\operatorname{pos}_{u,u'}(\gamma)) = s\}.$$

Here K_f is the lattice stabilizer used in the filling datum, with reduction kernel $K(g)$. For a set T of rational terms define the amplitude matrices (coordinates $I \times I$)

$$A(\gamma; u, u') = \int_{G_\infty} \beta_u(y) \beta_{u'}(\gamma_\infty^{-1} y) W(y) \otimes W(\gamma_\infty^{-1} y)^*\, dy, \tag{14.11}$$

$$A_T(u, s, u') = \sum_{\gamma \in T \cap \mathcal{E}(u,s,u')} A(\gamma; u, u'). \tag{14.12}$$

Define $A_T^\flat$ with $W^\flat$ in place of W.

Proposition 14.8 (Formal-source input: convergence and reproduction)**.** *For every u, s, u' the sums in (14.12) are absolutely convergent in the finite-dimensional tensor space. For fixed u, s the sum over u' and the corresponding rational terms of the integrated norm difference of the full and cut kernels, with the product of the two cutoffs as in (14.11), is at most $\nu b^{-k^3 f}/2$. The sum of the norms in the inner periodization in (14.10) is bounded uniformly in g. The Poincaré map has image in V.*

The first assertions are `amplitudeAbs_lt_top_and_cut_le`, `summable_amplitude` in `New/Filling/AmplitudeSums.lean`; the last two are `Fits.tsum_cutBessel_le`, `Fits.poincareMap` in `New/Filling/PoincareSeries.lean`. They are not consequences of pointwise boundedness of the seed alone. The majorant in Definition 14.5(2), together with the bounded box partition at infinity and lattice estimates, is used to prove convergence before interchanging the integrals. The reproducing-kernel condition and integration on the Bessel quotient then put the Poincaré series in the space of the required archimedean cusp type and principal finite level.

Lemma 14.9 (Exact coefficient identity)**.** *In the Petersson normalization there is a constant $v_0 = \mathrm{vol}(K(g)) > 0$ such that*

$$\langle Qe_{u,i}, \rho(s) Qe_{u',j} \rangle = v_0 A_{\mathrm{all}}(u, s, u')_{i,j}. \tag{14.13}$$

Here and below the pairing uses the convention in which the order of i, j agrees with (14.11); if one takes the opposite convention, both sides are conjugated/transposed accordingly.

Proof. Unfold the first periodization in the quotient Petersson integral and expand the second into its defining series. This gives a single sum of integrals over the whole adelic group, of the product of a seed of class u with a translated conjugate seed of class u'. Absolute convergence and Fubini are justified by Proposition 14.8 and the integrability statement `Fits.integral_poincare` in `New/Filling/PoincareSeries.lean`. After left translating the finite integration variable by the class representative h_u^{-1}, its range is $K(g)$. Normality of $K(g)$ in K_f implies that the second indicator is identically 1 on this range precisely when the relative element has reduction s, and otherwise identically zero. Hence the finite integral contributes v_0 and restricts the rational terms to the set $\mathcal{E}(u, s, u')$. The archimedean part of the same integral is (14.11); the two class cutoffs remain in it. Thus the integral is precisely (14.13). This is the unfolding carried out by `Fits.petersson_poincare` in `New/Filling/PoincareSeries.lean`. □

Set $P = (v_0\nu)^{-1/2} Q$. Then $P^* \rho(s) P$ in the coordinate basis has entries $\nu^{-1} A_{\mathrm{all}}(u, s, u')_{i,j}$. Projecting by R gives the translates of the row map after restriction to the ramified space, up to their fixed nonzero scalar. It is the normalized map P that will enter the variance estimate, whereas the integral-index theorem uses the actual unscaled map of Theorem A.

Lemma 14.10 (Identity term)**.** *The amplitude of the Bessel subgroup at $s = 1$ is exactly*

$$A_{S_B}(u, 1, u')_{i,j} = \nu \delta_{u,u'} \delta_{i,j}.$$

In particular the Bessel-subgroup contribution to P^*P is the identity operator.

Proof. The elementary reduction property used here is `Filling.eq_and_mem_classUnip_of_besselS` in `New/Filling/NearTerms.lean`: a term in S_B over $(u, 1, u')$ forces the two column classes to be equal and belongs to $U_u(g)$. This follows by projecting to H_0 and using the definition of the column classes and triviality of their automorphism groups at K_c (the precise compatibility with the modified lattice and the level is the cited formal-source input). Thus only

$u = u'$ can contribute. On summing over the unipotents, covariance of W and (14.8) eliminate the second cutoff. The remaining integral is ν times the integral of $W_i\overline{W_j}$ on the unipotent quotient. Use the product over v of the orthonormality relation in Lemma 14.7. This proves the claimed identity (`amplitudeSum_besselS_one` in `New/Filling/AmplitudeSums.lean`). □

14.6 Energy, pairs of congruent terms, and removal of lower cells

The following notation allows the estimates for the moment to be stated with their actual normalizations. For a subset $\Omega \subseteq \mathfrak{G}$ define

$$\mathcal{E}(A_T;\Omega) = \frac{1}{\nu^2}\sum_{u,u'\in\mathcal{C}}\sum_{s\in\Omega}\|A_T(u,s,u')\|_{\mathrm{HS}}^2. \tag{14.14}$$

The norm here is the Euclidean norm on the $I \times I$ tensor. Put

$$D = \frac{c^2}{d_\infty \operatorname{mass}_r(E)}, \quad D_1 = |\bar{S}_B|c, \quad \tau = b^{kf/4}, \quad \mathfrak{c} = 2\tau^2(b/g)^{kf},$$

where $\bar{S}_B$ is the reduction of the integral Bessel subgroup. These are `mainEnergy`, `besselEnergy`, `tailScale`, and `codimFactor`, respectively; the formulas and inequalities in the following proof use this normalization. In particular, the admissibility assumptions give

$$\tau^{44}D_1 \le D, \qquad \tau^{18}\mathfrak{c}D \le D_1, \qquad \mathfrak{c} \le 1. \tag{14.15}$$

This pair of bounds is central: the Bessel-subgroup energy, despite its large number of translations, is extraordinarily small on the scale of the full expected moment. The first bound means it is inexpensive to use coarser estimates on the lower cells; the second expresses the strong saving from nontrivial finite congruence codimension. They are part of `Filling.IsAdmissible` and are verified for the data in `New/Filling/FillingAdmissible.lean`; the verification is a *Formal-source input* using the mass and unipotent measure formulas and degree bounds. The first equality in (14.2) shows $D/|\mathfrak{G}| = c^2/n_0$.

For two compatible collections of amplitudes A, B, any $\delta > 0$ gives

$$\mathcal{E}(A+B;\Omega) \le (1+\delta)\mathcal{E}(A;\Omega) + (1+\delta^{-1})\mathcal{E}(B;\Omega). \tag{14.16}$$

This follows from $2\Re\langle a,b\rangle \le \delta\|a\|^2 + \delta^{-1}\|b\|^2$ in each summand. By absolute convergence the amplitude for a set decomposes into amplitudes for disjoint subsets. The global truncation error follows from Proposition 14.8 by applying the triangle inequality to each tensor, squaring, and summing:

$$\mathcal{E}(A_T - A_T^{\flat};\Omega) \le |\mathcal{C}|\,|\mathfrak{G}|\,b^{-2k^3f}/4 \le b^{-30}D_1.$$

For the last inequality use $|\mathfrak{G}| \le (g_0g_1)^{r^2f} \le b^{22r^2f}$, $c \ge |\mathcal{C}|$, $|\bar{S}_B| \ge 1$, and k much larger than h, hence $2k^3f - 22r^2f \ge 30$. This is `energy_sub_cut_le` in `New/Filling/SchurTest.lean`. Thus there is substantial freedom to cut off the Bessel functions. No comparable freedom is allowed in the constant of the main term.

The expansion of the squared amplitude in (14.14) gives pairs of rational elements *with equal finite reductions*: for each (u,s,u'), their relative positions belong to the same coset of $K(g)$ in K_f. These pairs are stratified by degeneracy of their flags and by the links between their middle spaces. We describe the stratification and record its precise resulting estimate.

Character orthogonality. Call a set of terms stable if it is invariant under left and right multiplication by rational unipotents. Write ψ_U for the additive Bessel character on rational unipotents as read in finite adelic coordinates; equivalently it is the inverse product of the infinite-place characters. For a term γ, a necessary condition to contribute is that

$$\psi_U(u) = \psi_U(\gamma^{-1}u\gamma) \quad (u \in U(E^+) \cap \gamma U(E^+)\gamma^{-1}). \tag{14.17}$$

Indeed the product $W(y) \otimes W(\gamma_\infty^{-1}y)^*$ transforms, on replacing y by $u_\infty y$, by the quotient of these characters. To see why cutoffs cause no obstruction, group the sum over $T \cap \mathcal{E}(u, s, u')$ into $U_u(g)$–$U_{u'}(g)$ double cosets. On a double coset, the product of the cutoffs can be replaced by any measurable weight of total one over the double action: this follows by Fubini from summing the two cutoff identities (14.8). The resulting integral is an orbit integral, and the mismatch character on the stabilizer makes it vanish. For a nontrivial character choose an element at which it is not one; the integral equals that nontrivial scalar times itself. In the formal proof the rational stabilizers relevant to the finite cosets and the full unipotent intersection give the same vanishing condition by the approximation theorem for the unipotent group, and the exact weight-change identity (*Formal-source inputs*: `hasSum_amplitudeSum_change_of_weight` in `New/Filling/AmplitudeSums.lean`, `amplitudeSum_eq_zero_of_not_charsAgree` in `New/Filling/CrudeEnergy.lean`, and `exists_rat_approx_unipotentChar_eq` in `NearTerms.lean`). The character must be retained until this step; replacing it by its modulus would lose the exact vanishing.

Normal forms and links. A term with degenerate discriminant $\mathfrak{d}(\gamma) = 0$ but compatible characters has a shared vector in the frame flags. The proof uses a descending family of subgroups G_j fixing the first j hyperbolic pairs ($0 \le j \le h$). For such a shared vector, compatibility of characters forces equality of the Hermitian norms of the two leading coefficients and propagates the shared-vector relation to subsequent levels. To see the mechanism, let v be a vector orthogonal to the relevant preceding flag, and w an isotropic vector of the flag lying in that orthogonal complement. The unitary Eichler transformation is

$$x \longmapsto x + w\langle v, x\rangle - (v + \eta\langle v, v\rangle w)\langle w, x\rangle, \qquad \eta + \bar{\eta} = 1.$$

Expanding the two pairings using $\langle w, w\rangle = \langle w, v\rangle = 0$ and the relation on η shows that this transformation preserves the Hermitian form. Its simple-root coordinates against the flag are linear in the coefficients of w and in the pairings of v with the successive flag vectors. The Bessel character of the transformation tests this linear combination. Applying the transformation with the corresponding common vector on each side of the conjugation, equality of characters for all rational scalings, and nondegeneracy of the additive pairing, forces equality of the linear functionals. Elementary linear algebra then propagates the shared-vector relation up the flag. In the source these steps are `exists_mem_besselU_eq_eichlerMap` in `EichlerFlag.lean`, `hasCommonFlagVector_of_le` and `CharsAgree.exists_hasLevel` in `CrudeEnergy.lean`. The latter supplies the exact assertion used below: a term satisfying (14.17) with discriminant zero has a level $j \le h$ in the sense of `Filling.HasLevel`. We use this implication as a *Formal-source input*, retaining its character hypothesis.

The link of a term is a linear subspace of $X_0 \oplus X_0$. In the split basis, let $M(\gamma)$ have as columns $e_0, \dots, e_{h-1}, \gamma e_0, \dots, \gamma e_{h-1}$ and let p_{X_0} denote projection to the rank-k positive middle coordinates. The link is the image of the kernel of $M(\gamma)^* J_Y$ under the map

$$z \longmapsto (p_{X_0} z, p_{X_0}\gamma^{-1} z).$$

It describes the relation between the two positive middle coordinates after imposing orthogonality to both isotropic frames. This definition is `link` and `linkMatrix` in `New/Filling/`

`EichlerFlag.lean`. When two terms have different links, their common part has smaller dimension; the resulting loss is reflected in the congruence codimension factor $\mathfrak{c}$. For the top cell (nonzero discriminant) the primary surviving energy is the sum of squared norms of blocks for a fixed link, not the sum of squares of individual terms. For a set $\mathcal{L}$ of links put

$$\mathcal{E}_{\mathrm{link}}(T,\Omega,\mathcal{L}) = \sum_{u,u',s\in\Omega}\sum_{\Lambda\in\mathcal{L}} \|A^{\flat}_{T\cap\{\mathrm{link}(\gamma)=\Lambda\}}(u,s,u')\|^2.$$

Unlike (14.14), this definition does not divide by ν^2; we keep the source normalization.

To estimate the discarded terms one first replaces the double cutoff by a partition in frame boxes of size $b^{2s_0^2}$, using a cutoff modulo the frame stabilizer U^0. Within such boxes all frame-product entries have bounded archimedean norm. Their finite congruence constraints and the product formula give lower bounds on the nonzero discriminants and on relevant lattice indices. The proof further selects near representatives of the U^0 double cosets (`IsNearRep` of `BoxPartition.lean`). For each pair of near representatives with matching finite reductions and level j, one bounds its contribution by products of explicit weights, introducing the lattice of common link vectors. The weights are inverse powers of: the discriminant of the lattice of link vectors, the index of the frame-value lattice, the saturation index of the paired frame lattice, and the factor $g^{2f\min(j,h)}$. In particular the pair invariant of `pairInvariant` is the product of the discriminant, frame-value index, and saturation index. The weight `termWeight` includes the powers $-(k+1), -k, -k$ of the three corresponding invariants (with a count normalization depending on the codimension).

The point of these inverse powers is that the total weight can be bounded by estimates for sublattices rather than a direct count of pairs. The geometry-of-numbers input (*Formal-source input*: `Filling.tsum_sublatticeDisc_inv_le` in `NearTerms.lean`) bounds sums of inverse discriminant powers using short generators obtained by a relative Siegel lemma. Bounds on the congruence indices and lattice intersections followed by these sums give the explicit theorem `sum_pairWeight_le` in `BlockBound.lean`: the total pair weight is at most $2\tau^2$. The class amplitude size prior to summation is the function `classSize` of `SchurTest.lean`. With M_0 the normalization `countWeight0` and z the relevant index parameter it has the exact form

$$a_j(z) = \tau^2\left(\frac{D_j}{M_0 g^{2\min(h,j)kf} z^{k+1}}\right)^{1/2}, \quad D_j = \begin{cases} \tau^{13-32j}D_1 & j \le h, \\ D & j = h+1. \end{cases}$$

At the last level the estimate for $a_j(z)$ is obtained from (14.7) and the archimedean and finite discriminants. At earlier levels the weighted Schur estimate bounds the retained common-vector blocks. The exact comparisons with class sizes and the total weights are `IsNear.le_classSize_and_cellMass_le` in `SchurTest.lean`, and `enorm_amplitudeSum_le_classSize` and `enorm_amplitudeSum_le_classSize_top` in `BlockBound.lean`. We take these as formal-source inputs. These estimates are uniform in the rank and in the field; qualitative finiteness of the double cosets or a fixed-rank Siegel lemma estimate alone would not give the stated bound.

Proposition 14.11 (Formal-source input: reduction to coherent top-cell pairs)**.** *Retain the notation D, D_1, τ. Suppose T is stable, Ω is a set of reductions with the common fixed-vector norm bound $X \ge 1$ of `Filling.FixesNormAtLeast`, and $C_0 \ge D_1$ is a constant satisfying the following extra restrictions. If a term in T over Ω is in normal form zero, require $C_0 \ge \tau^{20}D_1$; if an omitted link exists, require $C_0 \ge \tau^{12}D/X^k$. For $w \ge 0$, suppose*

$$\mathcal{E}_{\mathrm{link}}(T\cap\{\mathfrak{d}\neq 0\},\Omega,\mathcal{L}) \le w\nu^2.$$

Then

$$\mathcal{E}(A_T;\Omega) \le (1+b^{-4})^2 w + b^{-12}C_0. \tag{14.18}$$

This is `Filling.energy_le_of_linkEnergy_le` in `New/Filling/CrudeEnergy.lean`; the two subordinate assertions are `energy_le_of_level_pairs` and `energy_le_of_level_lt`. Their quantitative errors are $C_0/(b^{30}4^{h+1-j})$ at level j and the accumulated error below level n' is at most $C_0/(b^{30}4^{h+1-n'})$. To explain the accumulation, split off the next level by (14.16) (taking $\delta = 1$ here); its own error and the preceding one together cost at most four times the preceding-level budget. Thus a geometric factor 4 suffices from one level to the next. If a term has discriminant zero, the character obstruction or a lower normal form just discussed applies, so it contributes only to this error. At the last level, separate the coherent equal-link blocks, of cost $w\nu^2$, from the unequal-link blocks estimated by the common link and its codimension. Returning to the full amplitudes with (14.16), now using $\delta = b^{-4}$ and the extremely small truncation error, gives (14.18). Both the precise coefficient of w and the estimate on the omitted terms matter: choosing a fixed loss at the final step instead of $1 + b^{-4}$ would not yield a nearly spanning family. The nontrivial block estimates above, rather than their algebraic assembly, are the imported part of this proposition.

A first consequence concerns translations that locally lie in the image of the Bessel subgroup. For $i = 0, 1$, set

$$\Omega_i = \{s : \text{there exists } t \in \bar{S}_B \text{ with } s_i = t_i\}.$$

Proposition 14.12 (Formal-source input: Bessel-coset error)**.** *For each i one has*

$$\mathcal{E}(A_{G(E^+)\backslash S_B}; \Omega_i) \le b^{-12} D_1.$$

This is `energy_compl_besselS_le` in `New/Filling/Moments.lean`. For these reductions the frames satisfy the exact Bessel constraints on an entire prime's chosen factors. The finite congruences and the character condition imply that a remaining normal-form-zero term would actually be in the Bessel subgroup, so such a term does not contribute to the complement. The common constraints supply the norm bound X for applying Proposition 14.11 with $\mathcal{L} = \varnothing$, $w = 0$, $C_0 = D_1$; the scale inequality (14.15) guarantees the requisite domination of the full-cell error. The exact group-theoretic constraint and norm calculation are part of the cited source proposition. This restrictive error estimate isolates the diagonal term and is used again to bound the non-Bessel part of a central projection.

14.7 The sharp coherent main term

The preceding estimates deliberately allow coarse constants, because they are used only for errors. For the coherent open-cell pairs a different argument must retain enough positivity to give the leading constant one. We record the form of this argument and isolate its three substantive inputs: local orbital kernels, rational-to-local block comparison, and the global sieve.

Kernel systems and their products. Consider a measured commutative ring A with an additive character $\psi : A \to \mathbb{C}^\times$ of modulus one, and let $A^\times$ be its group of units. For a general semilocal ring put $u(z) = z$ if z is a unit and $u(z) = 1$ otherwise. Let p, q denote sequences of units; $p[i : c]$ means the sequence obtained by replacing entry i by $u(c)$. A sieve system is a family of forms $J_i(p, q)$, $0 \le i \le h + 1$, with the following relations:

$$\int_A J_{i+1}(p[i : c], q[i : c])\, dc = J_i(p, q), \qquad |J_0(p, q)| \le \mathfrak{e},$$

$$\left| \int_A \psi(\xi c) J_{i+1}(p[i : c], q[i : c])\, dc \right| \le \beta(\xi)\Lambda_i(p)\Lambda_i(q). \tag{14.19}$$

The Fourier inequality is required for unit frequencies, and all the integrals in these relations are absolutely convergent. For $i < h$ one also requires the unequal-entry bound

$$|J_{i+1}(p,q)| \le \beta(\iota(p_i) - \iota(q_i))\Lambda_{i+1}(p)\Lambda_{i+1}(q)$$

when the difference of inverse-head coordinates in this expression is a unit; the map ι is inverse of a unit, with the additional factor $1/g$ at the finite places. At the top the kernel is continuous and every matrix $\big(J_{h+1}(p_a,p_b)\big)_{a,b}$ with common last entry is positive semidefinite. This is an abstraction of `SieveSystem` in `New/Filling/EichlerFlag.lean`; we have omitted the infinite tails of the sequences from the notation.

Here is a useful consequence. If p_a is any finite family of heads, for $i \le h$ define the function

$$F_i(c) = \sum_{a,b} J_{i+1}(p_a[i:c], p_b[i:c]).$$

This is nonnegative, because the top forms are positive semidefinite and lower forms are obtained by integrating a common coordinate, successively. It is lower semicontinuous wherever the head maps are continuous, by Fatou's lemma after integration over subsequent entries. Fubini and (14.19) imply

$$\int F_i = \sum_{a,b} J_i(p_a,p_b), \qquad |\widehat{F_i}(\xi)| \le \beta(\xi)\left(\sum_a \Lambda_i(p_a)\right)^2.$$

The same scheme applies to the product of systems at finitely many places; the top positive matrices multiply entrywise and remain positive semidefinite by the Schur product theorem, the integration compatibility follows from Fubini, and the constants and majorants multiply. This is the elementary content of `SieveSystem.pi`, `SieveSystem.step_sum`, and `SieveSystem.nonneg_and_lowerSemicontinuous`. Positivity is retained in the sum over both heads; the individual summands need not be nonnegative, and no absolute value may be inserted before the sum without changing the argument.

Finite orbital kernels. Let Σ be a set of chosen places above g_0 and fix, for each $v \in \Sigma$, a nonzero frame vector at the place or its conjugate. Define $\Omega^\circ(\Sigma)$ to consist of the level reductions fixing the specified vectors, and for which, unless Σ is the full set, there exists at least one selected place above g_0 at which all required leading frame minors are units. This is the set denoted `fixingSet` in the formal proof. This is the exact source condition; imposing simultaneous unit-minor tests at all places gives a smaller subset and is sufficient for the later positivity lower bound. At an arbitrary finite conjugation-stable set of places containing the level primes and all denominators of a finite initial selection of rational terms, the argument works over the product of the corresponding local fields. Choose an intermediate compact open subgroup between $K(g)$ and K_f whose local factor at each place contains the appropriate fixed-vector subgroup. The ratio of the bottom orbital integral to the group measure is diminished by a factor b^{-6k} at each specified place. The precise construction is `exists_intermediateGroup_one_prime` and the product version `exists_intermediateGroup_lines` in `New/Filling/FixedLines.lean`; it uses counting on split matrices and a bound on the volume of nonunits (at all the level places this is at most b^{-9}).

The finite system incorporates the Bessel character of conductor given by $1/g$ and the different. If $p = (p_i)$ is a head over the local product of fields, let $[D_j]$ denote the generalized ring index associated to the partial product $D_j = p_0 \cdots p_j$. Write $\mathrm{den}(\xi)$ for its denominator index. The finite head majorant is

$$\Lambda_i^f(p) = \prod_{j<i} [D_j]^{s_0^6/2} \quad \text{if all } D_j \text{ are integral, and 0 otherwise.}$$

The finite frequency majorant is

$$\beta_f(\xi) = \frac{[\xi]^{(k-s_0^7)/2}}{\mathrm{den}(\xi)^6} b^{s_0^9 f},$$

where $[\xi]$ is the generalized index, not necessarily an integer. These definitions are `finHeadWeight` and `frequencyGain` in `EichlerFlag.lean` and `GlobalSieve.lean`. After dividing the top orbital kernel by the volume of a compact open additive neighborhood of zero and by the prescribed mass ratio, the bottom constant of the finite sieve system is at most

$$(1 + b^{-6}) b^{-6k|\Sigma|}. \tag{14.20}$$

These are the conclusions of the *Formal-source input* `Filling.IsFillingLevel.exists_finite_sieveSystem` in `New/Filling/FiniteSystem.lean`. We emphasize that the normalization of the orbital kernel and the bottom constant include the measure of the fixed-vector subgroup, rather than only a bound on its order. The proof passes down the tower of groups fixing successive hyperbolic pairs, integrating one entry of the frame Gram matrix at each step. Its key step is a local integral comparison (`Level.integral_succ_floor_eq` and `Gauge.integral_floorForm_succ`, `New/RealPlace/LocalTheorem.lean`). In the nonzero-frequency bound one uses the local character along a simple-root direction of an Eichler transformation, and averaging gives a power of the inverse module of the frequency. A parallel unequal-head calculation uses frequency $(p_i^{-1} - q_i^{-1})/g$; this is the origin of the unequal-entry bound in the definition of a sieve system. These local comparisons are substantive statements for indefinite unitary groups of growing rank, not applications of Plancherel alone.

The archimedean and finite systems together. At the infinite places use the systems of Section 14.4, and at the finite ones the systems of (14.20). The diagonal embedding of the rational frame entry into the local products scales it at each v by l_{vi}. Its covolume factor in Poisson summation is

$$c_i = \frac{\mathrm{vol}(\widehat{\mathcal{O}}_E)}{\prod_{v\in\Phi} l_{vi}^2} \qquad (0 \le i \le h).$$

The square in the denominator is the real Jacobian of multiplication by l_{vi} on a complex coordinate. Let A_0 be a finite set of rational units integral away from the auxiliary finite set, along with their inverses, and let $V_0 \subseteq A_0$ index the last rational entry. With the product system write $p[t]$ for the $h+1$ entries obtained from $p \in A_0^h$ and $t \in V_0$, read diagonally in the local coordinates. The following inequality is the output of the global sieve.

Proposition 14.13 (Formal-source input: the global sieve)**.** *If a finite sieve system has bottom bound $\mathfrak{e}_f$, head and frequency weights as above, then*

$$\Re \sum_{t\in V_0} \sum_{p,q\in A_0^h} J_{h+1}^f(p[t], q[t]) \prod_{v\in\Phi} J_{h+1}^v(p[t], q[t]) \le \prod_{i=0}^{h} c_i \left(\mathfrak{e}_f \prod_{v\in\Phi} \frac{1 + b^{-10}}{d_v} + \frac{b^{-6kf-10}}{d_\infty} \right).$$

The exact theorem is `sieve` in `New/Filling/GlobalSieve.lean`. The proof uses `one_sided_poisson` in that file; we give its mechanism because it explains the small additive error. For the function F_i above, one-sided Poisson summation on the additive adele lattice gives a bound for the sum over a finite set of rational points by c_i times its zero-frequency integral plus c_i times an error estimated by the sum of its nonzero frequency coefficients. A one-sided version is needed for lower semicontinuous kernels; it follows by comparison with integrable smoothing approximants to a nonnegative function and the elementary finite-subset inequality. The Fourier bounds and the assumption $(\sum_p \Lambda_i(p))^2 \le \Gamma \prod_{a<i} c_a$ give an error at

most $\Gamma\varepsilon\prod_{a\le i}c_a$ at this step if the sum of nonzero-frequency majorants is at most ε. When undoing a common head coordinate to descend another floor, the unequal head contributions give a second error of this size, bounded using the off-diagonal condition. Consequently induction from $i=0$ gives a bound

$$\prod_{a\le i}c_a\bigl(\mathfrak{e}+(2i+1)\Gamma\varepsilon\bigr).$$

The source's `one_sided_poisson` states this explicitly, with rational frequencies that are integral away from the finite set. Products of the weights β_f and β_v are summable uniformly in the finite set: by the product formula the large powers of the norms cancel, leaving the damping power 6 and denominator factors. The needed elementary sum with exponent k' is the precise inequality `sum_frequencies_le`: over a finite set of nonzero T-integers y, if $|c_v|^{-1}\le|x_v|\le B_1$, $B_1\ge1$, then

$$\sum_y\frac{[y]_T^{k'/2}}{\mathrm{den}(y)^6}\prod_v\lambda_{k'}(|c_vv(y)|^{-1})\le(b\,\mathrm{rd}_E)^{3f}\prod_v\bigl(|x_v|^{k'}(bB_1^2)^3\bigr).$$

This is proved in the source using a sum of inverse discriminants of rank-one sublattices of a rank-two lattice (the vector $(1,y)$), including its denominator lattice. It is a formal-source input, together with the one-sided Poisson lemma in the stated level normalization. Taking $k'=k-s_0^7$ and the actual archimedean scaling the strong inequalities of `Filling.IsAdmissible` give $(2h+1)\Gamma\varepsilon\le b^{-6kf-10}/d_\infty$. This is substantially better than a sieve bound with a factor b^η, and is the reason the second moment has the requisite leading constant.

Comparison of blocks with the orbital kernel. On the nondegenerate cell the terms are put in Weyl normal form $w_p\sigma$ with common residual middle element σ, common top entry $t=\langle z,\sigma z\rangle$, and $|v(t)|>N$ at every infinite place. The terms supported on identical finite reductions and identical link are sorted into such packets. Finite sets of initial representatives suffice by the absolute convergence of the amplitudes; subsequently enlarge the finite set of places to contain all denominators and take limits of local compact-open approximations. The exact comparison is `sum_blocks_le_topKernel` in `New/Filling/FiniteSystem.lean`; here is its normalization. Let $\mathcal{U}=|\bar U|=g^{fh(2k+2h+1)}$ be the order of the reduced integral unipotent subgroup, let

$$c_*=(\pi^k/((k-1)!N^k))^f\,D_E^{1/4}\,v_f^k,\qquad A_*=|\bar S_B|\mathcal{U}|\mathcal{C}|c_*.$$

For a finite packet of distinct quadruples (u,u',s,σ) with the given t, the squared norm of the corresponding sum of coset amplitudes, multiplied by the common normalization, is bounded by

$$(1+b^{-6})A_*\,\Re\sum_{p,q}R_{pq}\frac{J^{\mathrm{top}}(p,q)}{\mathrm{vol}(L_0)\ \mathrm{massRatio}},$$

for sufficiently small compact open additive L_0. Here R is the Gram matrix of the product of normalized archimedean correlation tensors. The exact factors in R are those of (14.5), and the finite module factors in the packet comparison cancel against these factors. This is a *Formal-source input*: the proof uses a count of integral points on a quadric for the last coordinate (`siegel_sum_le_lintegral_quadric`), and a local block comparison with the positive phase kernel, then the global-to-local limit identity `exists_hasSum_cosetAmplitude` in `New/Filling/GlobalToLocal.lean`. The normalization A_* incorporates the unipotent order and the same-class condition. It should not be replaced by a crude number of rational points. The factor c_* is the constant of the rank-k definite sphere sum. The factor $(|v(t)|^2-N^2)^{k-1}$ in

this count is the one that cancels against the open-cell Jacobian in the normalized real-place correlations.

The size comparison used in the compact proof is the following hypothesis of the filling datum (with our notation):

$$A_* \le (1+b^{-9})D\, d_\infty c_G^f\, v_f^{-s_0} \left(\frac{\nu\mathcal{L}}{|I|}\right)^2.$$

This formula is `packetConst_mul_mass_le` in `Filling.IsAdmissible`. It accounts for the fact that the second moment is compared with a definite genus mass, not just a count of the rows in the particular target representation. Combining the block comparison, Proposition 14.13, the scale factors, and the crude stratification bound yields the following sharp estimate. In contrast with a fixed-group spectral equidistribution statement, the constants here are uniform for the auxiliary r, k, f of this proof.

Proposition 14.14 (Formal-source input: sharp energy on the open congruence set)**.** *For every set Σ of chosen places above g_0, with specified nonzero fixed vectors, one has*

$$\mathcal{E}(A_{\mathrm{all}}; \Omega^\circ(\Sigma)) \le (1-b^{-4})(1+9b^{-4})D\, b^{-6k|\Sigma|}.$$

This is `energy_fixingSet_le` in `New/Filling/Moments.lean`. To obtain exactly this factor the coherent bound is taken with slack $(1+b^{-4})Db^{-6k|\Sigma|}$ before applying the two small error steps of Proposition 14.11. The resulting expression is at most $(1-b^{-4})(1+9b^{-4})Db^{-6k|\Sigma|}$ by the size bounds and $\tau^{20}D_1 \le Db^{-6kf}$, which follows from (14.15). The main analytical content of this proposition is precisely the uniform packet and local sieve inputs specified above. We have not replaced those inputs by a claimed classical trace formula theorem without quantitative estimates.

14.8 A finite-field positivity argument for singular minors

The open-set estimate of Proposition 14.14 does not yet imply a total energy estimate, since it leaves out the locus of singular minors. The next finite group argument shows that a nonnegative function of positive type cannot concentrate too heavily on such a locus unless it concentrates on fixed-vector subgroups. The distinction between the energy itself and a general positive-type function is useful because multiplying by a fixed-vector indicator preserves positive type, allowing induction.

Let G be a finite group, $\pi : G \to \mathrm{GL}_R(\mathbb{F}_q)$ a representation, and let A, A' be $R \times j$ matrices. Let $\Xi : G \to \mathbb{R}_{\ge 0}$ be of positive type; choose vectors u_g in a complex Hilbert space with $\langle u_a, u_b\rangle = \Xi(a^{-1}b)$ (the Gram representation follows directly from positive semidefiniteness). Write $S_\Xi = \sum_g \Xi(g)$. For any nonzero $x, y \in \mathbb{F}_q^j$, suppose

$$\sum_{g:,\pi(g)Ax=Ax} \Xi(g) \le Y, \qquad \sum_{g:,\pi(g)^{-T}A'y=A'y} \Xi(g) \le Y.$$

Lemma 14.15 (Finite-field determinant estimate)**.** *Under these assumptions,*

$$q \sum_{g:,\det(A'^T\pi(g)A)=0} \Xi(g) \le 2S_\Xi + q^{j+R/2}Y.$$

Proof. Fix a nontrivial additive character ψ of $\mathbb{F}_q$. For nonzero y, x consider the weighted exponential sum

$$S(y,x) = \sum_g \psi(y^T A'^T \pi(g) A x)\Xi(g).$$

In $\ell^2(\mathbb{F}_q^R)$ let e_z be the delta basis and f_z its normalized Fourier basis with entries $q^{-R/2}\psi(z^T z')$. These are orthonormal bases by additive character orthogonality. For the tensors $U_1 = \sum_a f_{\pi(a)^{-T}A'y}\otimes u_a$ and $U_2 = \sum_b e_{\pi(b)Ax}\otimes u_b$, the squared norms are at most $|G|Y$: on expanding the inner products of the tensor sums, the basis inner products give exactly the indicators of the associated fixed-vector subgroups. Their mixed inner product is $|G|q^{-R/2}S(y,x)$, or its conjugate according to the Fourier convention, by reindexing $a^{-1}b$. Cauchy–Schwarz thus gives $|S(y,x)| \le q^{R/2}Y$. If at least one of y, x vanishes then $S(y,x) = S_\Xi$.

Put $B_g = A'^T\pi(g)A$. Additive orthogonality gives $\sum_y q^{-j}\psi(y^T B_g x) = 1_{B_g x=0}$. Summing in x and weighting by $\Xi(g)$ shows

$$\sum_g \Xi(g)|\ker B_g| = q^{-j}\sum_{y,x} S(y,x) \le 2S_\Xi + q^{j+R/2}Y.$$

On the singular locus, $|\ker B_g| \ge q$, and all summands are nonnegative. This proves the inequality. It reproduces `card_mul_sum_det_le` in `New/Heights/OneDeterminant.lean`. □

In the application $R = r$ and j is at most the size s_0 of the frame. A vector fixer indicator is of positive type, because it is the matrix coefficient of the permutation representation on the orbit of that vector. Multiplying it by Ξ gives another nonnegative function of positive type, by taking a tensor product of Gram representations. Thus one can successively impose new fixed-vector conditions. This is exactly the feature needed to derive the following result from Proposition 14.14.

Proposition 14.16 (Removing the unit-minor condition)**.** *Let $\Xi \ge 0$ be of positive type on $\mathfrak{G}$ and satisfy*

$$\sum_{s\in\Omega}\Xi(s) \le \mathcal{E}(A_{\mathrm{all}};\Omega)$$

for every set Ω of translations. Suppose its support consists of reductions fixing specified nonzero frame vectors over every place in a set Σ above g_0. Then

$$\sum_{s\in\mathfrak{G}}\Xi(s) \le (1+9b^{-4})D\,b^{-6k|\Sigma|}.$$

Proof. Use descending induction in $|\Sigma|$. If all places above g_0 belong to Σ, there is no additional minor condition, and the assertion follows from Proposition 14.14. In general suppose the assertion fails for Ξ at Σ. At each new place v outside Σ, intersect its support with the prescribed fixer of any nonzero frame vector x in the appropriate standard or dual frame. The restricted function remains of positive type, is dominated by the energy, and is supported on the extended set of fixers. By the induction hypothesis its total mass is at most $(1+9b^{-4})Db^{-6k(|\Sigma|+1)} < b^{-6k}S_\Xi$. Lemma 14.15 bounds the mass of each singular-minor locus by

$$\frac{2+q^{j+r/2}b^{-6k}}{q}S_\Xi \le 3b^{-10}S_\Xi,$$

where $q = g_0$, and the last inequality follows from $q \le b^{11}$, $q \ge b^{10}$, and $j + r/2 \le k/2 + 2h + 2$, which with the scale bounds gives $11(j+r/2) < 6k$ (in particular $11(2h+2) < k/2$). The exact singular-minor test of `fixingSet` involves the standard and contragredient frame matrices, with one test for each of the leading minors at a chosen place; this identification is the finite matrix calculation in `sum_le_of_positive_type` of `New/Filling/Moments.lean`. Even bounding the number of tests by $2f(h+1)$, the excluded mass is at most $6f(h+1)b^{-10}S_\Xi \le b^{-4}S_\Xi$, by $f \le b/1000$, $h+1 \le b$. Consequently at least $(1-b^{-4})S_\Xi$ lies in $\Omega^\circ(\Sigma)$. Domination by energy and Proposition 14.14 imply $(1-b^{-4})S_\Xi \le (1-b^{-4})(1+9b^{-4})Db^{-6k|\Sigma|}$, contradicting the supposition. This proves the claimed induction. □

Here we have used only the imported sharp estimate on the open set and finite group positivity; we have not assumed that the geometric energy is automatically of positive type. In fact the function

$$\Xi(s) = \|P^*\rho(s)P\|_{\mathrm{HS}}^2$$

is of positive type. To see this, with $T_0 = PP^*$ put $u_s = \rho(s)T_0\rho(s)^*$ in the Hilbert space of Hilbert–Schmidt operators. Then $\langle u_a, u_b\rangle = \mathrm{Tr}(T_0\rho(a^{-1}b)T_0\rho(b^{-1}a)) = \Xi(a^{-1}b)$. It equals the energy per reduction by Lemma 14.9. With $\Sigma = \varnothing$ Proposition 14.16 therefore gives

$$\frac{1}{|\mathfrak{G}|}\sum_s \|P^*\rho(s)P\|_{\mathrm{HS}}^2 \le (1+9b^{-4})\frac{c^2}{n_0}. \tag{14.21}$$

This is where both local and global sharpness are indispensable.

14.9 First moment and bad finite-type loss

The estimates needed in addition to (14.21) are not general consequences of that second moment. They use the diagonal and central projections at the two different level primes. In fact

$$\sum_{i=1}^{c}\|(P^*P-I)e_i\| \le b^{-4}c, \qquad \|(I-R)P\|_{\mathrm{HS}}^2 \le 4b^{-4}c. \tag{14.22}$$

We explain the derivation and specify the further inputs.

For a central self-adjoint idempotent $z \in \mathbb{C}[\mathfrak{G}_i]$ and the inclusion $\jmath_i : \mathfrak{G}_i \to \mathfrak{G}$, let $M_z(s)$ be the $c \times c$ matrix with block from u to u' equal to

$$\nu^{-1}\sum_{a\in\mathfrak{G}_i} z(a)A_{G(E^+)\setminus S_B}(u, \jmath_i(a)^{-1}s, u').$$

The complement-coset energy controls this matrix at 1 after summing its row norms. Indeed convolution by the coefficients of z is an orthogonal projection on the regular representation and hence contracts the squared sum over s. To restrict to the relevant set Ω_{1-i}, note that it is invariant under left translation by the embedded factor $\mathfrak{G}_i$: the other component is unchanged. It contains the Bessel reductions. The convolution estimate and the covariance under the Bessel subgroup give

$$|\bar{S}_B|\|M_z(1)\|_{\mathrm{HS}}^2 = \sum_{s\in\bar{S}_B}\|M_z(s)\|_{\mathrm{HS}}^2 \le \mathcal{E}(A_{G(E^+)\setminus S_B}; \Omega_{1-i}).$$

The equality uses the following exact covariance calculation, a *Formal-source input* in the proof of `sum_norm_amplitudeSum_le` in `New/Filling/Moments.lean`. Each $s \in \bar{S}_B$ induces a permutation of the column classes by the rational Bessel element over (u, s, u'). The left and right covariance of the Bessel functions replaces the corresponding block by a unitary matrix acting on its coordinate space, times a scalar of modulus one. Norms are unchanged because the finite-dimensional compact types are unitary. This is derived from `amplitudeSum_move` in `New/Filling/AmplitudeSums.lean` and `exists_besselS_mem_termsOver` in `NearTerms.lean`. The convolution estimate is applied after keeping the terms whose factor-$(1-i)$ reduction lies in the image of the Bessel subgroup; the exact group-module contraction, together with this covariance, is the assertion `sum_norm_amplitudeSum_le` cited above. In particular the resulting sum of row norms is at most $b^{-4}c$ (using $\sqrt{c}\|M_z(1)\|_{\mathrm{HS}}$ and Proposition 14.12). Taking $z = 1$ and using the identity term gives the first inequality of (14.22); bounding row or column norms is equivalent for the self-adjoint $P^*P - I$.

The remaining statement is the trace of a small central projection on the Bessel-subgroup packet. If z is central and $\Re(z(1)) \le b^{-8}/2$, then

$$\Re \sum_{a\in\mathfrak{G}_i} z(a) \sum_{u,j} A_{S_B}(u, \jmath_i(a)^{-1}, u)_{j,j}/\nu \le b^{-4}c. \tag{14.23}$$

This is the *Formal-source input* `re_trace_besselS_le` in `Moments.lean`. It is not a trace formula on arbitrary cusp forms. Its proof realizes the class sum on the finite orbital integrals over the Bessel subgroup, using invariance of z under conjugation and the local character. The ell^2 norm of the coefficients of z is exactly $z(1)$ since $z = z^* = z^2$ in the regular group algebra. The remaining orbital coefficient has squared norm at most $2c^2$ after the measure normalizations (the lower-floor estimate in `New/Filling/FixedLines.lean`); hence Cauchy–Schwarz bounds the trace by $c\sqrt{2z(1)} \le b^{-4}c$. The crucial nontrivial part is the bound on the orbital coefficient, incorporated in (14.23). Applying this with $z = 1 - e_i$ and Lemma 14.2, the trace of $(1-\rho(e_i))PP^*$ is at most $2b^{-4}c$, since the non-Bessel part contributes at most the same bound $b^{-4}c$ from the previous paragraph. The two projections commute, so $1 - R \le (1-\rho(e_0)) + (1-\rho(e_1))$ as positive operators. Taking their traces against PP^* proves the second inequality in (14.22).

The notation of the formal proof for these two inputs is deliberately separate: `re_trace_besselS_le` controls the Bessel term, whereas `sum_norm_amplitudeSum_le` controls the complement for any central projection, including the identity. Both are needed even though the discarded relative Plancherel mass itself is tiny.

14.10 From the moment bounds to the covering theorem

We now give the final argument without any further automorphic estimate; it is a general Hilbert-space inequality. First, W is nonzero. If $R = 0$ then $\|P\|_{\mathrm{HS}}^2 \le 4b^{-4}c$ by (14.22), but the first bound in (14.22) gives $\mathrm{Tr}(P^*P) \ge (1-b^{-4})c$. These two inequalities contradict $c > 0$ and $b \ge 1000$. Thus $n > 0$.

For $s \in \mathfrak{G}$ put $T_s = \rho(s)PP^*\rho(s)^*$ and $\overline{T} = |\mathfrak{G}|^{-1}\sum_s T_s$. Unitarity and the trace identity show

$$\begin{aligned}\|\overline{T}\|_{\mathrm{HS}}^2 &= |\mathfrak{G}|^{-2}\sum_{s,t}\mathrm{Tr}(T_sT_t)\\ &= |\mathfrak{G}|^{-1}\sum_s \|P^*\rho(s)P\|_{\mathrm{HS}}^2.\end{aligned}$$

If $J \subseteq W$ and p_J is its orthogonal projection in V, set $U_s = (R\rho(s)P)^{-1}J$. For each unit vector $u \in U_s$, the two orthogonal components of $\rho(s)Pu$ lie respectively in J and $\ker R$. Summing over an orthonormal basis of U_s and enlarging the sums for these components gives

$$\mathrm{Tr}(p_{U_s}P^*P) \le \mathrm{Tr}(p_JT_s) + \|(I-R)P\|_{\mathrm{HS}}^2.$$

In this computation p_{U_s} is the orthogonal projection within C and R commutes with $\rho(s)$. For any orthogonal projection p_U in C, one has

$$|\,\mathrm{Tr}(p_U(P^*P - I))| \le \sum_i \|(P^*P - I)e_i\|,$$

by expanding in the coordinate basis, applying Cauchy–Schwarz to each inner product, and using $\|p_Ue_i\| \le 1$. Applying this with $U = U_s$ gives

$$\dim U_s \le 5b^{-4}c + \mathrm{Tr}(p_JT_s).$$

Averaging and using Hilbert–Schmidt Cauchy–Schwarz, $\|p_J\|_{\mathrm{HS}} = \sqrt{\dim J}$, and (14.21), we obtain

$$\begin{aligned}\frac{1}{|\mathfrak{G}|}\sum_s \dim U_s &\le 5b^{-4}c + \sqrt{\dim J}\sqrt{(1+9b^{-4})c^2/n_0}\\ &\le \left(5b^{-4} + \sqrt{(1+9b^{-4})\alpha}\right)c\\ &\le \tfrac{1}{2}(1+\alpha+19b^{-4})c,\end{aligned}$$

where $\alpha = \dim J/n$. The second line uses $n \le n_0$, and the third uses $2\sqrt{uv} \le u+v$. Thus there exists s attaining the indicated bound. This is `Fits.exists_translate_finrank_comap_le` in `New/Filling/Moments.lean`, with the underlying Hilbert–Schmidt variance argument `finrank_sub_le_re_inner_dyadSum` in `New/Heights/OneDeterminant.lean`. The constant 19 comes from the projected and identity estimates $(2 \cdot 5)$ and from the second moment (9).

Finally, ramify the columns by applying $D_N = \prod_{w|\pi}(1 - z_w)$, with z_w the generator of the finite central-unit action at a selected place above N. Its image is C^{ram}, since in a cyclic finite group the image of $1 - z_w$ is the kernel of the average over the group, and these operators at different places commute. In the source the estimate `ramifyingOp_finrank_comap_and_short` in `New/Columns/ChainValuation.lean` shows that for any subspace $U \subset C$ the dimension of its inverse image under D_N is at most $\dim U + fc/(N-1)$: the central-unit action permutes classes freely in cycles of $N-1$, so the kernel has relative dimension at most $1/(N-1)$ at each place. The freeness and compatibility with the column type are the formal-source input used here; the dimension inequality also follows immediately from rank-nullity for D_N from its kernel bound. Admissibility gives $19b^{-4}/2 + f/(N-1) \le b^{-2}$, since $N - 1 \ge b^9$ and $f \le b/1000$. Inserting this loss and removing the common nonzero normalization scalar proves (14.1). This completes the derivation of the spread property appearing in `Data.ramifiedColumns_fill_rows` in `New/Filling/FillingTheorem.lean`.

It remains to convert spread into a small number of translates. Put $\eta = b^{-2}$ and start with $J = 0$. If the current codimension $q_* = n - \dim J$ satisfies $q_* > 4\eta n$, choose s by (14.1). By rank-nullity applied to the quotient map from C to W/J, the increase Δ in dimension obtained by adjoining $P_s(C^{\mathrm{ram}})$ satisfies

$$\Delta \ge c\left(\frac{q_*}{2n} - \eta\right) \ge \frac{cq_*}{4n} > 0.$$

In particular this is a genuinely new translate. Define a potential on the remaining codimension by

$$\Psi(q_*) = 5n \log\left(1 + \frac{q_*}{\eta n}\right).$$

When $q_* > 4\eta n$, adjoining the new image decreases the potential by at least

$$\Psi(q_*) - \Psi(q_* - \Delta) = -5n \log\left(1 - \frac{\Delta}{q_* + \eta n}\right) \ge \frac{5n\Delta}{q_* + \eta n} \ge c,$$

since $q_* + \eta n \le 5q_*/4$. The process terminates because dimension increases by a positive integer at every step. The potential remains nonnegative; hence the number of selected translates times c is at most $\Psi(n) = 5n\log(1+b^2)$. The final codimension is at most $4b^{-2}n$. This proves Theorem 14.1. The greedy covering is the result `exists_few_translates` of `New/Heights/OneDeterminant.lean` (the preceding derivation is a complete mathematical proof of that deduction).

For the application, take a spanning set of short integral columns as in Theorem A, and for the above finite set T select independent vectors from its projected images after the rational

rescaling. The span has the dimension supplied by Theorem 14.1, and the accumulated logarithmic index bound is at most $5n\log(1+b^2)B$. Apply the individual squared-length estimates to these independent vectors and Hadamard's inequality to their lifted Gram matrix. The loss from the small missing fraction of dimensions is controlled in the final comparison by the crude bound for the heights. No further spanning theorem is needed. The substantive assumptions for B have thus been reduced from the theorem itself to the precise uniform Bessel system, lattice block, trace transfer and local orbital/Poisson estimates; the finite-field counting, second-moment positivity, variance, and greedy selection are explicit deductions from those estimates.

Chapter 15

The analytic and arithmetic mechanism of the Bessel sieve

In Chapter 14 we isolated the quantitative covering property of the projected Poincaré map and its second-moment proof. We now give a more detailed account of the estimates behind that property. It is useful to organize the material by the three changes of measure that occur. First one goes from Haar measure on a real unitary group to a Bessel quotient. Next one introduces a tower of measures on successive open cells, replacing one pivot coordinate at a time by a Fourier integral. Finally one compares the rational orbit count with the finite-place measure and uses Poisson summation. The last two steps must preserve the leading constant. A conventional argument that merely bounded the integrals up to a factor depending on the ranks or the fields would not produce a sufficiently large subspace of rows.

We retain the notation $h = m/2 - 1$, $s_0 = h + 1$, $r = k + 2h + 1$ of Section 14.3. The letter d_v is a real formal degree (not the quadratic prime d). Unless the measure is explicitly specified otherwise, matrix norms in this section are operator norms. We use $\|\cdot\|_{\mathrm{HS}}$ for the Hilbert–Schmidt norm and write J_* for a Hermitian matrix with $J_*^2 = 1$. Thus $U(J_*)$ is embedded in $\mathrm{GL}_r(\mathbb{C})$ and

$$g^{-1} = J_* g^* J_*, \qquad \|g^{-1}\| = \|g\|, \qquad \|g\| \geq 1.$$

The symbols $\mathbb{A}$ and $\mathcal{O}$ refer to the adeles and integers of the CM field E. When using an assertion with "formal-source input" below, we mean the precise result at the indicated Lean declaration. In several instances we derive its consequences, or a substantial part of its proof, but the remaining uniform coordinate estimate is explicitly retained as an input. The purpose is to separate those inputs from elementary arguments of harmonic analysis and linear algebra, and to make their quantitative roles intelligible; we do not claim to have independently established every assertion.

15.1 Norm balls in the local groups

A useful feature of coefficient estimates on $U(J_*)$ is that a fixed inverse power of the operator norm is sufficient to dominate all integrals on the Bessel subgroup, and the exponent required depends only polynomially on the rank. We give the argument because it is repeatedly used both to justify rearrangement of integrals and to prevent the normalization of a Bessel function from destroying the tail bound.

Lemma 15.1 (Norm-ball covering)**.** *Let S be a locally compact group, $j : S \longrightarrow U(J_*)$ a continuous homomorphism, and ν_S a right Haar measure on S. Suppose that $D, R \geq 1$,*

$y \in U(J_*)$, and $K \geq 3DR^2$. Then

$$\nu_S\{s : \|j(s)y\| \leq R\} \leq K^{2r^2}\nu_S\{s : \|j(s) - 1\| \leq D^{-1}\}.$$

If j is proper, the set on the right has finite measure; if the Haar measure is nonzero it has positive measure.

Proof. In a real normed vector space of dimension $2r^2$, a set of points in the radius-R ball with mutual distances at least $\delta > 0$ has cardinality at most $(1 + 2R/\delta)^{2r^2}$. Indeed disjoint open balls of radius $\delta/2$ about them lie in the radius-$R + \delta/2$ ball; comparison of Lebesgue volumes proves the bound. Take $\delta = (DR)^{-1}$ and choose a maximal such separated family among the matrices $j(s)y$ with norm at most R. There are at most $(1 + 2DR^2)^{2r^2} \leq K^{2r^2}$ of them, with lifts $s_i \in S$. Every s in the set under consideration satisfies $\|j(s)y - j(s_i)y\| \leq \delta$ for some i, and

$$\|j(ss_i^{-1}) - 1\| \leq \|j(s)y - j(s_i)y\|\|(j(s_i)y)^{-1}\| \leq 1/D.$$

By right invariance of ν_S, counting the covering sets proves the assertion. Properness gives compactness of the inverse image of the indicated closed ball. Positivity follows since it contains the inverse image of an open neighborhood of the identity. □

This is `measure_ball_le_covering` in `Lib/DiscreteSeries/Haar/RegularRepL2.lean`. In the proper case, put $Q = \nu_S\{s : \|j(s) - 1\| \leq 1\}$. For $c > 8r^2$ the covering lemma implies

$$\int_S \left(\frac{2}{\|j(s)y\|}\right)^c d\nu_S(s) \leq 16^c Q. \tag{15.1}$$

For clarity, use the elementary layer-cake identity, for a function $H \geq 1$,

$$\int_S H^{-c} d\nu_S = c\int_1^\infty t^{-c-1}\nu_S\{H \leq t\}\, dt.$$

Indeed $H^{-c} = \int_H^\infty ct^{-c-1}dt$ pointwise. The lemma with $D = 1$, $R = t$ bounds the last integrand by $3^{2r^2}Qct^{4r^2-c-1}$, and $c/(c - 4r^2) \leq 2$. Multiplying by 2^c and weakening the constant proves (15.1). The corresponding formal result is `integral_decay_le`. The argument works on a subgroup, not only on the full group; the properness assumption is what ensures finite measure in the initial ball.

We also need a version of (15.1) with the real formal degree included. Let $l = 2\lambda$ be the doubled Harish–Chandra parameter, let ψ_l denote the matrix coefficient of the unit vector defining the discrete-series kernel, and let $d_l = \mathrm{formalDegree}(l)$. We use the following representation-theoretic facts.

Proposition 15.2 (Formal-source inputs: lowest-type coefficient estimates)**.** *For the large-gap parameter l of rank r with $|l_i| \leq L_0$, one has the normalized orthogonality integral $d_l \int |\psi_l(g)|^2 dg = 1$ and $|\psi_l(g)| \leq 1$. For each integer c at most the gap bound, one has*

$$|\psi_l(g)|^2 |\det a_g|^c \leq 1, \qquad |\psi_l(g)| \leq (2/\|g\|)^c, \tag{15.2}$$

where a_g is the block diagonal part relative to the signs of J_*. Also

$$|\psi_l(xg) - \psi_l(x)| \leq 100(r + L_0 + 1)^5\|g - 1\|. \tag{15.3}$$

These are `norm_dsCoeff_sq_mul_det_pow_le_one`, `norm_dsCoeff_le_pow` in `Lib/DiscreteSeries/FlenstedJensen/PsiDecay.lean` and `norm_dsCoeff_sub_le`, `half_le_norm_dsCoeff` in `Lib/DiscreteSeries/FlenstedJensen/TailBounds.lean`, using the orthogonality results of `Lib/DiscreteSeries/CuspKernel/SchurOrthogonality.lean`. We will use the exact estimates, not an unspecified asymptotic of ψ_l. For the first inequality in (15.2) the source uses polar decomposition and Fischer's inequality for positive definite block matrices; the second follows from the bound $\|g\|^2/4 \le |\det a_g|$. The Lipschitz bound comes from differentiating along a short Cayley-coordinate path and bounding the infinitesimal action on the lowest vector. We describe that differentiation and the relevant operator estimate below. The full coefficient formula in the variable signature and normalization is kept as a representation-theoretic input.

Lemma 15.3 (Degree times tail mass)**.** *With $M_0 = r + L_0 + 1$, if $c > 8r^2$ then*

$$d_l \int_{U(J_*)} (2/\|g\|)^c \, dg \le 8(2000M_0^6)^{2r^2} 2^c.$$

Proof. Put $D = 200M_0^5$. Inequality (15.3) implies $|\psi_l(g)| \ge 1/2$ on the ball $\|g - 1\| \le D^{-1}$. Hence the measure of this ball is at most $4\int |\psi_l|^2 = 4d_l^{-1}$. Lemma 15.1 gives $\nu(\|g\| \le R) \le 4d_l^{-1}(2000M_0^6R^2)^{2r^2}$. Integrating the same layer-cake identity as before and using $c/(c - 4r^2) \le 2$ proves the result, including finiteness. □

Thus the normalization by the formal degree is harmless at the level of an explicit polynomial in the maximal weight. This explains why `integral_tail_le` can be used even when d_l grows far faster than any fixed power of the scale x.

15.2 A rank-one convolution estimate and its differentiated form

Here is a direct argument for the Schur bound needed on the Bessel quotient. It is useful to formulate it with a general reproducing coefficient rather than choose a preferred Hilbert model of the real representation. On a unimodular locally compact group write the usual right convolution convention as

$$(F * x)(y) = \int x(a^{-1})F(ya)\, da.$$

Translation of an S-equivariant F to the right preserves the norm on $S\backslash G$; for our application $S = U$ is the Bessel unipotent, and its character is unitary, so the scalar products of equivariant vectors descend. Let ψ be a Hermitian integrable coefficient with $\psi(1) = d > 0$. The normalization of the idempotent used by the source consists precisely in

$$\int_G \psi(a^{-1})\psi(gay)\, da = \frac{\psi(g)\psi(y)}{d}. \tag{15.4}$$

Taking $y = 1$ gives $\psi * \psi = \psi$. This is the rank-one convolution identity, for $\psi = d\psi_l$ by Schur orthogonality. We record an abstract lemma that also explains why no unwanted dimension factor appears in the local Fourier estimate.

Lemma 15.4 (Schur orthogonality for reproduced vectors)**.** *Suppose that F is bounded, measurable and integrable modulo the left unitary character, and $F * \psi = F$. Suppose x is*

integrable, $\psi * x = x$, and $\phi \in L^1(G) \cap L^2(G)$. With the above quotient measure,

$$\|F * x\|_2^2 = \frac{\|x\|_2^2}{d}\|F\|_2^2, \tag{15.5}$$

$$\left\| \int_G \phi(a)(F * x)(\cdot\, a)\, da \right\|_2^2 \le \frac{\|x\|_2^2}{d^2}\|F\|_2^2\|\phi\|_2^2. \tag{15.6}$$

The inequality, interpreted with right side $+\infty$, also holds for integrable ϕ.

Proof. Denote by $x^*(z) = \overline{x(z^{-1})}$ the convolution involution. The adjoint of the right convolution operator on quotient L^2 is convolution with x^*; this follows by the change of variables $y \mapsto ya^{-1}$ and Fubini. The hypotheses imply $x^* * \psi = x^*$. For every integrable w, (15.4), after substitution and Fubini, gives

$$\psi * w * \psi = \frac{(\psi * w)(1)}{d}\psi.$$

Consequently, applying this with $w = x * x^*$ and using $\psi * x = x$, one obtains

$$x * x^* = \psi * (x * x^*) * \psi = \frac{(x * x^*)(1)}{d}\psi = \frac{\|x\|_2^2}{d}\psi.$$

The integral on the right is finite: $x = \psi * x$ is bounded since ψ is bounded, and x is integrable. Pairing $F * (x * x^*)$ with F proves (15.5). For $z(a) = \phi(a^{-1})$, put $w = x * z$; then $\psi * w = w$. By the same argument $w * w^* = d^{-1}\|w\|_2^2\psi$. Since

$$\|w\|_2^2/d = (w * w^*)(1)/d$$

and $w * w^* = (\|w\|_2^2/d)\psi$, we can use Cauchy–Schwarz in the identity

$$\|w\|_2^2 = (x * z * z^* * x^*)(1) = \int z(a)\,(z^* * x^* * x)(a^{-1})\, da.$$

Left convolution by $x^* * x$ on $L^2(G)$ has norm at most $\|x\|_2^2/d$: it is a positive operator with its square equal to $\|x\|_2^2/d$ times itself, by $x * x^* = (\|x\|_2^2/d)\psi$ and $\psi * x = x$. Equivalently this follows from the L^2 norm of the convolution operator of x, since its product with its adjoint is $\|x\|_2^2/d$ times a projection. Thus $\|w\|_2^2 \le (\|x\|_2^2/d)\|z\|_2^2$. Applying (15.5) to w proves (15.6). All identities may first be taken for bounded compactly supported inputs and then extended by Young's inequality and approximation; boundedness and L^1 of x, ψ, F justify the stated Fubini manipulations. □

In the formal source these are `IsBoundedMeasurable.lintegral_sq_convolve` and `lintegral_sq_convolve_convolve_le` in `Lib/DiscreteSeries/CuspKernel/BesselSchurBound.lean`. In ordinary notation one can recognize them as the fact that the right orbit of any reproduced vector contains an isometric copy of the discrete series, up to the single formal degree. The identity uses the product of the convolutions in their indicated order and the rank-one identity; ordinary L^1 convolution estimates alone would not produce d^{-1}.

To differentiate the Bessel function, it is not sufficient to know that F is smooth. The bound has to be uniform in the representation and in the maximal order of differentiation. We use smooth coefficient vectors in the discrete-series representation. If u is the normalized lowest vector, and u_w is obtained by differentiating it in the coordinate directions of a word w of length j, the needed estimate is

$$\|u_w\| \le \big((r+1)^4(L_0 + r + 1 + j)\big)^j. \tag{15.7}$$

Formal-source input: this is the filtration-wise infinitesimal action bound used in `IsBesselFunction.hasSchurBound_of_convolve_eq` in `New/RealPlace/BesselCut.lean`, from `isKHighest_dsGenerator`, the unitary Casimir inequality and `norm_apply_le_of_mem_filtration`. It applies in a unitary normalized sign frame; conjugation of the coordinate differential directions by that frame does not increase their operator norms. To obtain (15.7) from the one-step bound on filtration of depth at most j, iterate from the unit vector, always bounding the previous depth by j. This retains the powers of $(r+1)$ and cannot introduce a factorial in L_0.

For completeness, given such coefficient vectors one may recover the corresponding right derivatives of F by differentiating the convolution reproduction formula. In the source this is done by first smearing the smooth lowest coefficient by an integrable bump operator β satisfying $\mathrm{smear}(\psi_l, \beta) = \psi_l$, and differentiating the bump instead of assuming differentiability of F. Differentiation under the integral is thereby justified by the smooth compactly supported test function. The word w produces an integrable kernel x_w which is d times the matrix coefficient of u_w and the original unit vector; it has $\psi * x_w = x_w$. Orthogonality gives $\|x_w\|_2^2/d = \|u_w\|^2$. Apply Lemma 15.4 to x_w to conclude that for every integrable ϕ,

$$\left\| \int \phi(a) F_w(\cdot\, a)\, da \right\|_{U\backslash G,2}^2 \le P^{2j} \frac{\|F\|_{U\backslash G,2}^2}{d} \|\phi\|_2^2, \qquad P \ge (r+1)^4 (L_0 + r + 1 + 2s_0^2), \quad j \le 2s_0^2. \tag{15.8}$$

This is exactly the constant in (14.4). The continuity of all differentiated convolutions follows from the continuity in $L^1(G)$ of translates of an integrable kernel, and the basic boundedness of F. The differential relation between successive words follows by testing along paths of right multiplication and using the ordinary chain rule on the bump. In particular no differentiation of the compact cutoff is used in (15.8); the estimate concerns the *full* Bessel function. This distinction is used later in the energy comparison between full and cut functions.

15.3 Threshold projectors and a lowest-type coefficient

The construction of a peak starts with a coefficient, not with a compactly supported test function. We give the reason that the alternating signs and steepness of the weights in that construction matter. In this subsection all matrix norms without an explicit subscript are Hilbert–Schmidt norms, and

$$J = \mathrm{diag}(\epsilon_1, \dots, \epsilon_n), \quad K = \{g : g^* g = 1,\ gJ = Jg\}.$$

The two blocks of K are compact; they must not be confused with the two unitary groups in the pair construction below.

For real integral weights $w = (w_i)$ define

$$Q_w(X) = \sum_{i,j} |w_i - w_j| |X_{ij}|^2.$$

If P_t is projection onto the coordinates at which $w_i \ge t$, then

$$\sum_t \|[P_t, X]\|^2 = Q_w(X), \qquad \sum_t \mathrm{tr}(P_t X) = \sum_i (w_i + a) X_{ii}, \tag{15.9}$$

where the thresholds run from $-a+1$ to a, with $a = \sum_i |w_i|$. Indeed the ij entry of the commutator has absolute value $|X_{ij}|$ if precisely one of w_i, w_j exceeds the threshold, and is zero otherwise. There are exactly $|w_i - w_j|$ such thresholds. Each diagonal entry is counted $w_i + a$ times. The same identity holds with negative or shifted weights. In particular repeated weights give flat directions, unrelated to the size of the absolute values of the weights.

Lemma 15.5. *For a compact unitary matrix u one has*

$$\prod_t |\det(P_t u P_t|_{\operatorname{im} P_t})|^2 \le \exp(-Q_w(u)/2).$$

Proof. For a single orthogonal projection P put $A = PuP|_{\operatorname{im} P}$ and $B = (1-P)uP$. Then $A^*A = 1 - B^*B$, whose eigenvalues lie in $[0,1]$. The inequality $1 - x \le e^{-x}$, including at $x = 1$, gives

$$|\det A|^2 \le e^{-\operatorname{tr}(B^*B)} = e^{-\|B\|^2}.$$

Unitarity gives $\|Pu(1-P)\| = \|B\|$: the trace of A^*A equals the trace of AA^* and subtracting both traces from $\operatorname{rank} P$ proves the equality. Consequently $\|[P,u]\|^2 = 2\|B\|^2$. Multiplication and (15.9) conclude. □

This argument is `norm_weightMinors_sq_le_exp` in `New/DiscreteSeries/GaussianShape.lean`. It explains why integrating the coefficient over K does not destroy the concentration: weighted minors of a unitary matrix themselves act as a concentration density around the commutant of the weights.

Here are the precise smaller-scale statements required to evaluate the phase. They are recorded as formal-source inputs because a proof of the full Flensted–Jensen formula with explicit remainder would take us beyond the scope of this account. Write ψ_l for the normalized highest-vector coefficient, $\Lambda(l)$ for the diagonal of the lowest K-type (in the notation of the source, `lowestKType`), and $L_0 \ge \max|l_i|$.

1. In `New/DiscreteSeries/MinorExpansion.lean`, `exists_det_compress_eq_exp_logOneAdd`, if P, Q are orthogonal projections, $\|X\| \le \delta \le 1/2$, $\operatorname{tr}([P,Q]Y) = 0$, $\|[P,X] - Y\| \le \delta c$, and $\|Q - P\|, \|Y\| \le c$, then

   $$\det(1 + PX) = \exp\big(\operatorname{tr}(Q\log(1+X)) - \tfrac{1}{4}\operatorname{tr}(Y^2) + e\big), \quad |e| \le 100\delta c^2.$$

   The noncommuting term costs a norm *squared*; a remainder simply bounded by $O(\|X\|^3)$ would not suffice when some of the weights coincide.

2. In the same file, `exists_principalMinor_eq_exp` applies to the projected minors of the Flensted–Jensen integrand at $g = 1 + X$, $\|X\| < 10^{-3}$. For the corresponding threshold projection P and compact u it writes the minor as

   $$\exp\big(\operatorname{tr}(P\log g) + \tfrac{1}{4}\|[P,X]\|^2 + e_P\big), \qquad |e_P| \le 200000\|X\|(\|[P,X]\|^2 + \|[P,u]\|^2).$$

   The integrand is evaluated at the matrix y related to g by the triangular shifts in the Flensted–Jensen formula; this is not an asserted expansion of $\det(Pu^{-1}gu)$ at arbitrary g, u.

3. The resulting `dsCoeff_gaussian_shape` in `New/DiscreteSeries/GaussianShape.lean` is the uniform statement: take the minimum required noncompact weight separation at least $500n^3$ (in addition to the given gap parameter); let $T \ge \max(1, n^2, n^2\log(n + L_0 + 1))$, $\delta T \le 10^{-12}$ and $\|g - 1\| \le \delta$. Then there exists $0 \le \kappa \le 1$ such that

   $$\left|\frac{\psi_l(g)}{\exp(\operatorname{tr}(\Lambda(l)\log g))} - \kappa\right| \le \kappa/10 + e^{-100T}. \tag{15.10}$$

   The proof in the source takes κ to be the Gaussian with exponent minus one quarter of the weighted commutator quadratic form associated with the distinct shifted weights `-paramAddRho`; these weights must be distinguished from the lowest K-type in the exponential phase.

We explain the quantitative mechanism linking these statements. The Flensted–Jensen integral is a normalized compact integral of a product of minors times a lowest-type weight. Apply (15.9) to the first term and the quadratic term of the logarithmic expansion in each minor, using the appropriate weight sign. The phase term is independent of the integration parameter u and can be removed from the integral. The sum of the remainders is then bounded in the source by

$$10^6\delta(Q(X)+Q(u)/2+n^2),$$

with the quadratic forms determined by the shifted weights; the n^2 accommodates the shifts by the half-sum of roots. By Lemma 15.5, the compact density includes the factor $e^{-Q(u)/2}$. Thus, on the set where $Q(X)+Q(u)/2$ is at most a large multiple of T, the exponential of the remainder differs from 1 by a small relative error. On the complement, the Gaussian dominates the exponential of the remainder. The normalization of the compact integral cannot cost more than a power of the dimension of the weight representation, which is at most exponential in $n^2\log(n+L_0+1)$. This accounts for the condition on T. The separation between these two regimes is essential: requiring a relative error uniformly, including where κ is extremely small, would be false. Equation (15.10) supplies precisely the relative error *plus* exponentially small absolute error that can be combined over two factors.

For illustration, a general 2-by-2 matrix direction with $w_1 > w_2$ has

$$Q_w(X) = (w_1-w_2)(|X_{12}|^2+|X_{21}|^2).$$

The Gaussian width in these coordinates therefore scales as $(w_1-w_2)^{-1/2}$, not $L_0^{-1/2}$ if the two weights move together. A rank-one compact direction with $w_1 = w_2$ is completely invisible to the Gaussian. The common compact group in the next construction is exactly what allows these invisible directions to be removed before using (15.10).

15.4 Two-factor Laplace estimate

The restriction integral is not an integral of a positive coefficient: the additive character must cancel the first-order phase. Let $j_i : S \to G_i$ be a unitary pair of homomorphisms, of ordered ranks n_i. Write $u_i(s) = j_i(s) - 1$ and let Λ_i be the two lowest-K-type weight matrices. The norm on differences of matrices here is Hilbert–Schmidt for the phase and compact-correction arguments, and the norm of a group element itself is the operator norm in the decay bound. The relevant additive phase is

$$\varphi(s) = \operatorname{tr}(\Lambda_1 u_1(s)) - \operatorname{tr}(\Lambda_2 u_2(s)).$$

In the actual pair it is an imaginary-valued homomorphism. It follows that for $\|u_i(s)\| \le 1/2$,

$$\operatorname{tr}(\Lambda_1 \log j_1(s)) - \operatorname{tr}(\Lambda_2 \log j_2(s)) = \varphi(s). \tag{15.11}$$

Here is a useful elementary reason. Put

$$A_k = \operatorname{tr}(\Lambda_1 u_1(s)^k) - \operatorname{tr}(\Lambda_2 u_2(s)^k).$$

Additivity of the phase on powers s^q and the binomial expansion give $\sum_{k=1}^{q}\binom{q}{k}A_k = qA_1$. At $q=2$ this gives $A_2 = 0$; induction gives $A_q = 0$ for all $q \ge 2$. The logarithmic series is absolutely convergent, giving (15.11). Thus the cancellation is exact, not merely a first-order approximation. The source statement is `phase_logOneAdd` in `New/DiscreteSeries/TwoFactorSetup.lean`.

There is another elementary matrix identity behind the bound on the real part of the remaining phase. For $gJg^* = J$, the ii row gives

$$\mathrm{Re}(g_{ii} - 1) \cdot 2 + \sum_j |g_{ij} - \delta_{ij}|^2 = 2 \sum_{\epsilon_j \neq \epsilon_i} |g_{ij}|^2.$$

In fact the left side is the squared ordinary row norm minus 1, while the signed row norm is ϵ_i. Multiplying by nonnegative row weights w_i, bounding them by C on rows with noncompact directions, and summing, gives

$$\sum_i w_i \sum_j |g_{ij} - \delta_{ij}|^2 + 2\,\mathrm{Re}\,\mathrm{tr}(\mathrm{diag}(w)(g-1)) \leq 2C \sum_{\epsilon_j \neq \epsilon_i} |g_{ij}|^2.$$

One can also replace w_i in the first term by any smaller nonnegative quantity, as in `norm_sq_add_re_trace_le` of the same file. This accounts for the noncompact quadratic errors when comparing the two phases.

If g_{diag} is the part block diagonal with respect to signs, the decay estimate gives

$$|\psi_l(g)|^4 \leq (1 + \|g - g_{\mathrm{diag}}\|^2)^{-\mathrm{gap}}.$$

Indeed the two diagonal square determinants come from $1 + BB^*$ and its companion with B the off-diagonal block; the product inequality $\det(1 + BB^*) \geq 1 + \mathrm{tr}(BB^*)$ follows by expanding the product of its nonnegative eigenvalues. Combine this with the bound of the coefficient by the inverse power of the diagonal determinant recorded in the previous subsection. The precise simultaneous two-block expression is `norm_dsCoeff_pow_four_le` in `New/DiscreteSeries/TwoFactorSetup.lean`.

The remaining group-theoretic ingredient is an important distinction from a decay estimate in a given representation. The setup `IsTwoFactor` in `New/DiscreteSeries/TwoFactorSetup.lean` includes that the pair has matching zero-weight compact blocks, that it contains their common compact group, and that the steep weights away from this block have alternating signs. Under these conditions `exists_near_one` asserts that if the squared norms of the two off-block directions sum to at most ρ^2, $0 < \rho \leq 1/100$, one can multiply by an element of the common compact subgroup without changing the integrand so that

$$\|j_1(s) - 1\|, \|j_2(s) - 1\| \leq 10\rho.$$

This is a formal-source input (a specific matrix estimate on the pair); it does not follow just from j_i being representations. The two images of the compact subgroup have to come from the very same s.

Proposition 15.6 (Restriction integral)**.** *With the two-factor setup let $N = (n_1 + n_2)^2$ and $\beta \geq 10^{12}$, with $n_1 + n_2 + L_0 \leq \beta$ and weight gap at least $10^{30}(N \log \beta)^3$. Assume that j_1 is continuous and proper, j_2 is continuous and $\|j_2(s)\| \leq \beta + 1$ whenever $\|j_1(s) - 1\| \leq 1$. Use the integrand*

$$\theta(s) = \psi_{l_1}(j_1(s))\overline{\psi_{l_2}(j_2(s))}e^{-\varphi(s)}.$$

For Haar measure ds on S,

$$\mu\{s : \|j_1(s) - 1\| \leq 1\} \leq \beta^{20N} \left| \int_S \theta(s)\, ds \right|. \tag{15.12}$$

We give the integration argument of the source, separating from it the precise pointwise input used. In `New/DiscreteSeries/RestrictionLemma.lean`, `IsTwoFactor.laplace_method`,

the Gaussian estimate (15.10), the off-diagonal decay and `exists_near_one` give two auxiliary bounds:

$$|\theta(s)| \le (2/\|j_1(s)\|)^a \qquad (0 \le a \le 18N), \tag{15.13}$$

$$4|\theta(s)| \le 8\left(\operatorname{Re}\theta(s) + 16\beta^{-50N}(2/\|j_1(s)\|)^{8N}\right). \tag{15.14}$$

The numerical pointwise estimates are formal-source inputs in this proof. To see their structure, use the Gaussian estimate with T proportional to $N\log\beta$ after the common compact correction. The product of the two ratios to their exponential phase factors is within a small relative error of the product of the two nonnegative Gaussians, except for the exponentially small absolute errors. The exponential phase factors cancel by (15.11). The row norm estimate prevents the absolute errors from being magnified by an unbounded real exponential. Away from the applicable neighborhood the gap forces decay, and (15.13) permits the residual error to be measured against an integrable weighted function, rather than just the uniform bound β^{-50N}. The cubic power in the gap is responsible for having both a sufficiently small logarithmic neighborhood and sufficiently drastic decay beyond it.

Proof of Proposition 15.6 from the pointwise estimates. Put $A = \{\|j_1(s) - 1\| \le 1\}$ and $A' = \{\|j_i(s) - 1\| \le (200\beta^5)^{-1},\ i = 1, 2\}$. The two-factor norm ball covering lemma gives

$$\mu(A) \le \beta^{16N}\mu(A').$$

To apply it use the proper block direct sum of the two homomorphisms (properness follows from j_1), the norm bound by $\beta + 1$ for the second group on A, and the small radius $1/(200\beta^5)$. The elementary net cardinality is a power of β with exponent $16N$ (real matrix dimension at most $2N$). Properness ensures the sets have finite Haar measure. The near-identity coefficient inequality, with its weight bound by β, gives $|\psi_{l_i}(j_i(s))| \ge 1/2$ on A'. The phase has absolute value one. So

$$\mu(A') \le 4\int_S |\theta| \le 8\left(\left|\int_S \theta\right| + 16\beta^{-50N}\int_S (2/\|j_1(s)\|)^{8N}\,ds\right).$$

By the norm decay integral bound in the opening subsection the last integral is at most $16^{8N}\mu(A)$, because $8N > 8n_1^2$. Multiplying by β^{16N}, the coefficient of $\mu(A)$ on the right is at most $128 \cdot 16^{8N}\beta^{-34N} < 1/2$. Absorb it, then use $16\beta^{16N} \le \beta^{20N}$. This gives (15.12). □

Several features of the argument are worth noting. It never invokes positivity of the original coefficient or replaces the oscillatory integral by an integral of its modulus. It is the near-positive cone estimate (15.14) which prevents cancellation. Rescaling the Haar measure multiplies both sides of (15.12) by the same factor. Hence the estimate is compatible with taking Haar probability on the compact factor of the Bessel subgroup, while leaving Haar measure on the unipotent factor fixed by the global construction. Finally the resulting lower bound depends on a volume near the identity, not on an arbitrarily normalized coefficient at some moving point. This is what permits the peak comparison to be transported to the quotient where the Bessel kernel is used.

15.5 Interpolation, isotropic frames, and the Bessel peak

The role of the strict order of the weights is not merely to ensure membership in regular discrete series. It also permits a constructive real embedding whose restriction is guaranteed to be nonzero by Proposition 15.6. We reproduce the algebra of this embedding.

Lemma 15.7 (Alternating interpolation weights)**.** *Let $a_0 > \cdots > a_{d-1}$ be distinct integers, and set*

$$c_i = \Big((-1)^i \prod_{j \neq i} (a_i - a_j)\Big)^{1/2}.$$

Then $c_i \geq 1$ and, for $\deg p < d$, one has

$$\sum_i \frac{(-1)^i p(a_i)}{c_i^2} = [t^{d-1}]p(t).$$

Proof. There are exactly i negative factors in the defining product, so its signed value is a positive integer. The polynomial

$$L_i(t) = \frac{\prod_{j \neq i}(t - a_j)}{\prod_{j \neq i}(a_i - a_j)}$$

equals 1 at a_i and 0 at all other a_j. The degree $< d$ polynomial $p - \sum_i p(a_i)L_i$ has at least d roots and hence is zero. Compare leading coefficients. □

This is `exists_lagrange_weights` in `Lib/General/LinearAlgebra/FramedCoefficient.lean`. We use it in a space with $2h + 2$ noncompact coordinates, of alternating signs $+, -, +, -, \ldots, +, -$. The other coordinates form the common positive compact block. Let the exceptional weights in descending order be λ_i, $0 \leq i \leq 2h + 1$, and let a be the distinguished weight in the compact factor. The steepness guarantees that $t_i = \lambda_i - a$ are strictly descending and negative. Apply the lemma to

$$(a_0, a_1, \ldots, a_{2h+2}) = (0, t_0, \ldots, t_{2h+1}).$$

Form columns in the exceptional coordinates by

$$(C_0)_i = c_0/c_{i+1}, \qquad (C_j)_i = t_i^j/c_{i+1} \quad (j \geq 1).$$

Multiplying the $i + 1$ sign in the lemma by -1 is the essential sign change: the ambient Hermitian sign is $(-1)^i$. For $j + l \leq 2h + 2$ the raw identity is

$$\sum_i \frac{(-1)^i t_i^{j+l}}{c_{i+1}^2} = \frac{\delta_{j+l,0}}{c_0^2} - \delta_{j+l,2h+2}.$$

Consequently

$$\begin{gathered} H(C_0, C_0) = 1, \\ H(C_j, C_l) = 0 \quad (0 < j + l < 2h + 2), \\ H(C_j, C_l) = -1 \quad (j, l > 0,\ j + l = 2h + 2). \end{gathered}$$

In particular $C_1, \ldots, C_h$ form a totally isotropic frame and C_{h+1} is a negative unit vector orthogonal to this frame. No assumption about unknown zeros of an L-function enters here; the relation is just interpolation.

For a diagonal weight operator $D = \mathrm{diag}(\lambda_i)$ on the exceptional coordinates, one has $DC_j = aC_j + C_{j+1}$ for $j \geq 1$ and $DC_0 = aC_0 + c_0C_1$. These equalities show how the shifted weight projects onto a nilpotent chain. The trace phase of the source's pair is designed using that chain. If one simply chose an arbitrary isotropic flag of the same dimension, its Bessel character need not agree with the coefficient's differential.

As a check, take $h = 0$, $t_0 = -1$, $t_1 = -2$. Then $(c_0, c_1, c_2) = (\sqrt{2}, 1, \sqrt{2})$ and the two columns are

$$C_0 = (\sqrt{2}, 1)^t, \qquad C_1 = (-1, -\sqrt{2})^t.$$

Relative to the sign matrix $\mathrm{diag}(1, -1)$ these have squared lengths $1, -1$ respectively, and they are orthogonal. The construction generalizes the same cancellation to successively many isotropic columns.

It is not enough to prove that these columns give the desired Hermitian Gram matrix. We must bound the size of a conjugation identifying the resulting flag with the standard flag. The relevant elementary estimate behind that step is that the matrix of minors determining the columns has denominators $c_i \geq 1$. If $|t_i| \leq T$ then a raw column of degree j has norm at most $\sqrt{2h+2}\,T^j$. The initial column has the additional factor c_0, bounded by T^{h+1}. Nondegeneracy in the reverse direction follows from the negative anti-diagonal: a vector with pairing of magnitude 1 against a vector of norm M has norm at least M^{-1}. This explains why passing to orthogonal normalized frames here costs only a polynomial in the base parameter with exponent linear in h, rather than in the number of all entries of a k-by-k matrix.

The formal source makes the preceding construction precise in the following two assertions, which we use as formal-source inputs.

- `besselSize_le_of_dual`, `New/RealPlace/FrameNorm.lean`: if $V_0, \dots, V_{2h}$ are vectors of norm at most M, with V_j equal to the Bessel frame columns for $j \leq h$, and $\langle V_i, V_j \rangle = 0$ for $i + j < 2h$ and $= -1$ for $i + j = 2h$, then the Bessel size of the frame is at most $8(h+1)M$ (in the normalized sign frame).
- `IsSteep.exists_isTwoFactor_besselPairHom`, `New/DiscreteSeries/BesselExistence.lean`: for the steep sign and weight pattern on the two groups, there exists a unitary realization of the preceding columns, carrying the standard Bessel unipotent with its character to a pair of homomorphisms satisfying `IsTwoFactor`. The compact part is identified with the common zero-weight block. The linear trace difference is imaginary and additive and exponentiates to the given Bessel character on the unipotent. In the parameter ranges of the local construction this realization has $\mathrm{bs}(x) \leq b^{5s_0}$ and both homomorphisms meet the small-ball hypotheses of Proposition 15.6 for $\beta = b^3$.

The second assertion is stronger than Witt extension: the phase must match, and the quantitative bound for the dual array must also survive the identification. The anti-diagonal identity above exhibits the input to the verification of both properties.

Let S denote this Bessel subgroup, with Haar measure the unipotent Haar measure times Haar probability on its compact factor. In local coordinates its embedding into the ambient group is $s \mapsto s_1$; its action on the compact type space is s_2. For a unit vector v of the compact representation ρ with coefficient ψ_{l_2}, define

$$Y(y) = \int_S \psi_{l_1}(x^{-1} s_1 y)\, e^{-\varphi(s)} \rho(s_2^{-1}) v \, ds. \tag{15.15}$$

Conjugating the two group maps by their chosen frames, Proposition 15.6 gives

$$\|Y(x)\| \geq |\langle Y(x), v \rangle| = \left| \int_S \theta \right| > 0.$$

The inner product uses the normalized unit compact vector, so bounding the integral by the norm has no dimension factor. The norm of the integrand in (15.15) is the absolute value of the coefficient. The norm-ball decay theorem then gives absolute convergence; the same domination is uniform on compact sets of y by submultiplicativity and properness. Bessel

equivariance follows by changing the integration variable s by right multiplication, noting that the character is multiplicative and that Haar measure on the semidirect product is unimodular. Fubini and the coefficient idempotent ψ show $Y * \psi = Y$ with the normalized coefficient $\psi = d\psi_{l_1}$. Thus the whole orbit of the peak, and not just its value, lies in the reproduced subspace to which the Schur estimates apply.

For clarity, the useful property `HasPeakAt` of `New/DiscreteSeries/BesselExistence.lean` comprises both

$$\|Y(y)\| \le \int_S |\psi_{l_1}(x^{-1}s_1y)|\,ds, \qquad \int_S (2/\|x^{-1}s_1y\|)^{8r^2+1}\,ds \le b^n\|Y(x)\| \tag{15.16}$$

for every y in the normalized frame, with the specified exponent n, and $Y(x) \neq 0$. The second inequality follows from Proposition 15.6 by the norm-ball integral bound with the moving coset y; its use does not presume that y belongs to the image of S. The source declaration `exists_isBesselFunction` supplies Y with these properties under the above pair assumptions. The exact admissible n is chosen in that declaration so that in the global application $n + 99r^2 \le r^3$. This margin will pay for the conversion of the pointwise peak into an L^2-normalized cutoff.

15.6 A peak-to-cut inequality with all mass factors

The lower bound at the peak is of little global use if most of the resulting function is concentrated far from a bounded Bessel-size region. We now prove the cutoff argument in a form that displays why the formal degree, Haar measure, and peak height cannot be normalized independently.

Put

$$M = 8r^2 + 1, \quad C = 16^M, \quad R = b^{6s_0+2}, \quad \omega_a(g) = (2/\|g\|)^a, \quad I = \int_G \omega_{2M}(g)\,dg.$$

In what follows the norm is the source's submultiplicative matrix size, symmetrized under inversion, in the normalized frame. Let B be the group norm ball of radius 1 about the identity (measured by the norm of the difference); it has Haar measure $v_B > 0$. The Bessel size of a coset satisfies

$$\mathrm{bs}(yg) \le \mathrm{bs}(y)\|g\|.$$

This follows from its definition as the infimum of $1+$ the norm of a representative under the Bessel subgroup and the norm inequalities. Under the source's bounds $\mathrm{bs}(x) \le b^{5s_0}$ and $b \ge 2000$, every translate $t^{-1}B$, where $t = x^{-1}s_1y$, lies in $\{g : \mathrm{bs}(yg) \le R\}$. Indeed for $g = t^{-1}u$ with $u \in B$ one has $s_1yg = xu$ and $\mathrm{bs}(xu) \le 16b^{5s_0} \le R$.

Define the positive majorant and the positive finite measure

$$\Phi(y) = \int_S \omega_{2M}(x^{-1}s_1y)\,ds, \qquad d\nu(g) = \frac{C}{v_B} b^n\|Y(x)\|\omega_M(g)\,dg.$$

The degree-tail bounds in the first subsection imply $I < \infty$ and $\int \omega_M \le Cv_B$, so $\nu(G) \le C^2b^n\|Y(x)\|$. Writing $t = x^{-1}s_1y$, for $g \in t^{-1}B$ the norm of g is at most $16\|t\|$. Thus

$$\nu\{g : \mathrm{bs}(yg) \le R\} \ge b^n\|Y(x)\|\omega_M(t).$$

Multiply by $\omega_M(t)$ and integrate over s. The second inequality of (15.16) cancels $b^n\|Y(x)\|$ and proves

$$\Phi(y) \le \nu\{g : \mathrm{bs}(yg) \le R\}. \tag{15.17}$$

This argument does not require the family of translates $t^{-1}B$ to be disjoint.

The quotient integration convention is important. With Haar probability on the compact part of S and the given unipotent Haar, Weil integration gives

$$\int_{U\backslash G} \Phi(y)\,dy = I.$$

Indeed unfold first along the unipotent, then use left invariance under the compact part and x. The decay exponent can be chosen so that the first inequality in (15.16) implies $\|Y(y)\| \le \Phi(y)$ everywhere. If $\mathrm{bs}(y) > R-1$, the estimate $\mathrm{bs}(y) \le \mathrm{bs}(x)\|t\|$ forces $\|t\| > 2b$ for every s. An extra decay exponent A therefore improves the same bound to

$$\|Y(y)\| \le b^{-A}\Phi(y) \quad (\mathrm{bs}(y) > R-1).$$

In particular Y is square integrable, since $\Phi \le \nu(G)$ and Φ has finite quotient integral. Its squared norm is strictly positive by continuity and $Y(x) \neq 0$.

It remains to estimate how much of the majorant is needed when Y is normalized. The coefficient reproduction formula at x, combined with the integral construction of Y and Bessel equivariance, gives the bound

$$\|Y(x)\| \le d \int_{U\backslash G} \Phi(y)\|Y(y)\|\,dy.$$

Apply Cauchy–Schwarz with weight Φ and use (15.17):

$$\|Y(x)\|^2 \le d^2 I \int_{U\backslash G} \Phi(y)\|Y(y)\|^2\,dy \le d^2 I\nu(G)\|Y\|_2^2. \tag{15.18}$$

Set $c^2 = D_v/\|Y\|_2^2$, where D_v is the index cardinality of the local type. Since $Y(x) \neq 0$ we can cancel its norm in (15.18), using the bound for $\nu(G)$, to obtain

$$c^2\nu(G)I \le D_v(dIC^2b^n)^2.$$

The degree-tail estimate gives $dIC^2 \le b^{99r^2}$ for the weight and rank ranges $r + L_0 + 1 \le b^4$ in use. Thus $n + 99r^2 \le r^3$ implies

$$c^2\nu(G)I \le D_v b^{2r^3}.$$

These explicit numerical inequalities are precisely the budget in `HasPeakAt.exists_cut_isCutPair`, `New/RealPlace/BesselCut.lean`; they make the following tail estimate independent of the original normalization of Y or of the Haar measure.

Use the continuous cutoff $\chi(t) = \max(0, \min(1, R-t))$. Write $F = cY$, $V = \chi(\mathrm{bs})F$. Then $\|F\|_2^2 = D_v$ and V is supported on Bessel size at most R. Multiplicative properties of bs and continuity of the norm give continuity of bs, so V is continuous. More explicitly if $g \to 1$ both $\|g\|$ and $\|g^{-1}\|$ in the multiplicative size tend to one; sandwich the two values of bs using the above inequality. The tail and mass estimates prove

$$\begin{aligned}\|V-F\|_2^2 &\le b^{-2A}c^2 \int_{U\backslash G} \Phi(y)^2\,dy \\ &\le b^{-2A}c^2\nu(G)I \le D_v b^{-10r^2}.\end{aligned} \tag{15.19}$$

The last step uses the source margin $A \ge 3r^3 f$, with f the number of infinite places, and the given ranges of the ranks. Moreover $c\Phi$ majorizes both norms, and (15.17) gives a majorant formed by translating a bounded-size region against a measure. These are stronger properties than a small L^2 tail; they will be needed on cell charts.

One must not apply the differentiated convolution estimate to V. Its cutoff is only continuous and piecewise linear, and it is not reproduced by the idempotent. The later bandlimit argument uses the full function F and inserts cutoff regions at the level of integration. The tail estimate (15.19) is exactly the price for passing between the two, in a setting where no derivatives of the cutoff are required.

15.7 From an L^2 reproducing estimate to the bandlimit tower

We detail the architecture of the still necessary local orbital calculation. The norm and tail estimates just proved do not imply arbitrarily strong Fourier decay on the orbit. The Fourier exponent depends on the dimension of the positive middle space, whereas the differentiated Schur estimate uses derivatives only of order at most $2s_0^2$. These are different sources of regularity.

The local calculation in `New/RealPlace/LocalTheorem.lean` has a descending tower of groups G_j (fixing an initial portion of the isotropic flag and its dual) with Haar measures, upper quotients and measures on floor coordinates. At a level j the notation md stands for the module of multiplication on the local additive field, so at a complex place $\mathrm{md}(z) = |z|^2$. Write the two heads as p, q. The multiplicative prefactor

$$D_j(p,q) = \prod_{a<j} \mathrm{md}(p_a q_a)^{d_a/2}$$

is the source's `Level.headDensity`, with d_a the successive floor dimensions. Under a change of variable by multiplication by p_a in d_a complex additive coordinates the real Jacobian is $|p_a|^{2d_a}$; writing the geometric mean of the two head Jacobians gives precisely the product above. It is especially important not to discard the square of the usual complex absolute value here: it would change the exponent of the final lattice-covolume factor.

The local normalized form is denoted by `floorForm`. It integrates a translate pairing of two vector-valued Bessel functions against the upper group and its descending convolution of floor measures. The corresponding absolute integrand is `formModulus` and its dominating integral is `formMajorant`. Descending convolution of measures means simply the pushforward of the product under the ordered group multiplication; it is not commutative. For bounded compactly supported functions, expanding one factor of this product and applying Fubini is the identity `upperPairing_eq_integral_succ`. For the actual forms, integrability on all the intermediate levels is obtained by the majorant bound below. Thus the equality of the zero-frequency integral with the next floor is an equality, not a rough inequality: the sharp constant in the final sieve depends on it.

The majorant is controlled by a product of two self-norms, denoted $\mathcal{D}_j(x)$ here and `Gauge.selfNormProduct` in the source. To illustrate the use of such a product, for a positive measure λ on the upper quotient and two vector-valued functions A, B, the elementary inequality is

$$\left| \int \langle A(z), B(z) \rangle \, d\lambda(z) \right|^2 \le \left(\int \|A(z)\|^2 \, d\lambda(z) \right) \left(\int \|B(z)\|^2 \, d\lambda(z) \right).$$

If the integration is also restricted to a measurable region of measure s^e, a further Cauchy–Schwarz with that region's indicator supplies a factor s^e on the right. This is the source of powers of a small floor volume, in contrast with powers coming from derivatives. In the actual gauge, there are changes of variables in the two factors and the self-norms live on different floors. The following are the formal-source inputs, with their quantitative conclusions; all are in `New/RealPlace/LocalTheorem.lean`:

1. `Gauge.lintegral_selfNorm_le`: for a pairing kernel at level $j \le s_0$, if the product of floor constants times its upper-quotient squared integral is at most c, the integral of its self-norm over the gauge's bounded positions is at most $s_0 c$.

2. `Gauge.formMajorant_sq_le`: if the floor mass is at most V, $V \ge 1$, with the specified near-region conditions and floor-value volume at most $s < 1$ on nonzero pairing integrands, then
$$\mathcal{I}_j^+(x)^2 \le s^{\mathrm{volumeExponent}} V^2 \mathcal{D}_j(x).$$
Here $\mathcal{I}_j^+$ is the form majorant. The hypotheses include nearness of both arguments of the pairing after their Weyl translations, as well as nearness of the support of the second Bessel function.

3. `Gauge.formModulus_sq_le_of_char`: for the character of frequency ξ on the diagonal frame Gram coordinate, assuming the gauge's multiplicative size bound for the character variant, the squared modulus integral is at most
$$(P.c/\mathrm{md}(\xi))^{|\iota|}\mathcal{D}_j(x).$$
The exponent $|\iota|$ is the cardinality of the complex middle-coordinate index in that gauge. The same right side occurs in `Gauge.formModulus_sq_le_of_heads` for the ordinary pairing at level $j = i + 1$ if $\xi = g^{-1}(p_i^{-1} - q_i^{-1})$. The first bound comes from averaging a variable character; the second comes from the same character on two unequal Weyl heads. Neither is just the assertion that an L^2 Fourier transform is square integrable.

4. `Gauge.lintegral_positions_le`: with the product of pairing norms equal to Θ^2, the function $a(x)$ on a bounded position set satisfying $a(x)^2 \le e^2 \mathcal{D}_j(x)$ obeys
$$D_j(p,q) \int a(x)\,dx \le e\,\Theta \left(\prod_{l<s_0} \mathrm{floorModulus}(l) \right) (P.c)^{|\iota| s_0} \Lambda^{1500000 s_0^3 + 1} (\Lambda')^{50000 s_0^2}.$$
Here Λ, Λ' are the sizes of the two gauges on the two position sets. This is a bound uniform in the Weyl heads; the head-density factor on the left is part of the assertion.

5. `Level.integral_succ_floor_eq` and `Gauge.integral_floorForm_succ`: integration with one additive character against the floor-value measure is the corresponding integral of the translate pairing with the character on the diagonal frame Gram variable, with the next upper group measure and stack measure. At zero character it gives the preceding floor. The hypotheses include the finite integral of the majorant. These statements implement the ordered convolution change of variables described above.

The distinction between the last two Fourier estimates can be seen abstractly. Suppose a pairing along a torus orbit contains two exponential factors $e^{2\pi i \,\mathrm{Tr}(p^{-1}z)}$ and $e^{2\pi i \,\mathrm{Tr}(q^{-1}z)}$ in its two vectors. Their inner product contains their ratio, so its character is that of $(p^{-1} - q^{-1})z$, independent of the common phase. This explains why the same Fourier gain controls both a nonzero additive frequency at the integration step and an unequal-head contribution at the step between floors. The scale by g^{-1} comes from the Bessel character's root-coordinate normalization, not a change of the frequency decay exponent.

One can therefore assemble the local tower as follows. At $j < s_0$ set J_j equal to the normalized floor form; at $j = s_0$ set $J_{s_0}(p,q) = \langle v(p), v(q) \rangle$, where v is the vector of normalized Bessel correlations. This last form is a Gram kernel when the last entries agree. The step identity gives J_j for zero frequency, and the majorant and frequency estimates give head weights times the common damped-power function at nonzero frequency or unequal heads.

The top floor identity uses the change of variables between the diagonal Gram entry and the common last vector. The exact geometric normalization also uses that compact double cosets of the last scalar coordinate with the same module have the same pairing. The source theorem implementing this assembly is

hasBandLimits_of_isCutBessel in New/RealPlace/BandLimits.lean.

It states that for a local cut/full Bessel pair of an admissible rank, positive degree d, Schur bound $\mathrm{SchurLe}(F, d^{-1}D_v)$ and $d \cdot \mathrm{bandConst} \le b^{4r^2}$, there exists a kernel tower satisfying `HasBandLimits`. Its bottom bound $(1 + b^{-10})/d$, the head weights and the nonzero-frequency majorant are those in Section 14.4. We retain the local coordinate and numerical inequalities just stated as formal-source inputs. The purpose of separating them is that the analytic normalization $d^{-1}D_v$, which was derived in Section 15.6 and the Schur argument, now becomes the precise number entering this tower, whereas the large Fourier exponent $k - s_0^7$ results from averaging in the middle coordinates.

15.8 The cell L^2 estimate and its limited differentiation cost

The open-cell estimates serve a separate purpose from the bandlimit tower: crude and unequal-link bounds need an integral over the frame Gram section. The local lemma which transfers quotient norm estimates to bounded sets in the group is worth spelling out. The residual stabilizer U^0 fixes the entire chosen frame, while the quotient of U by that stabilizer is described by upper triangular frame entries. Its dimension in the complex real-place coordinates is at most $2s_0^2$.

Suppose w is nonnegative, measurable and left Bessel-invariant, and a is a nonnegative measurable cutoff satisfying

$$\sup_y \int_{U^0} a(uy)\,du \le K.$$

Suppose whenever $a(y), a(sy) \ne 0$ and $s \in U$ the upper frame entries of s have module (ordinary complex norm) at most R. For any U-orbit meeting the support of a, choose a single point y_0 of this intersection. All other points in the intersection are obtained by elements whose frame matrices lie in the box of side $2R$ about zero in each of the real and imaginary coordinates. Integrating across U^0 first, the contribution of the outer coordinates is at most $K(2R)^{2s_0^2}$; the use of a single basepoint on each orbit is only a pointwise bound of that integral, and requires no measurable choice. Weil integration and the invariance of w give

$$\int_G a(y)w(y)\,dy \le K(2R)^{2s_0^2} \int_{U\backslash G} w(y)\,dy.$$

This is the substance of `lintegral_mul_le_of_frameMatrix_le`, `New/RealPlace/SobolevCell.lean`, using the source's coordinates with the product additive measure. The dimension that appears here is *not* the real dimension of U, which includes the large positive-middle block.

Here is an elementary Sobolev estimate explaining the other appearance of $2s_0^2$. For a continuously differentiable scalar function f on an interval of length $2t$ centered at 0, one has

$$|f(0)| \le (2t)^{-1} \int_{-t}^{t} |f(x)|\,dx + \int_{-t}^{t} |f'(x)|\,dx.$$

This follows from $f(0) = f(x) - \int_0^x f'(u)\,du$ and averaging. The same inequality holds with the norm of a vector-valued function. Iterate the corresponding two operators in each of d_0

coordinates of a box of side $2t$; by Cauchy–Schwarz and the inequality $(\sum_I a_I)^2 \le 2^{d_0}\sum_I a_I^2$, obtain

$$\|f(0)\|^2 \le 2^{d_0}\sum_{I\subseteq\{1,\dots,d_0\}} (2t)^{-d_0+2|I|}\int_{[-t,t]^{d_0}} \|\partial_I f(z)\|^2\,dz.$$

Only derivatives of order at most d_0 appear, one per coordinate. Apply this to the functions obtained by smoothing F with an L^2 test function; the derivatives are controlled by (15.8). Choosing the small coordinate box comparable in size to P^{-6}, as in the source, costs a power of P of order s_0^2. The ambient rank affects derivative sizes but not the number of derivatives.

What converts these inequalities into the actual open-cell bound is the source's change of variables through the frame Gram map. Specifically `le_lintegral_of_gramSection` of the real-place cell calculation supplies the Jacobian proportional to $c_G\,\mathrm{disc}(Z)^{k-1}$, with Haar normalization c_G fixed by the Gram normalization. Together with the chart and near-coordinate estimates, `lintegral_gramSection_integral_le` in `New/RealPlace/SobolevCell.lean` is the formal-source input with the following conclusion: for a full Bessel function satisfying Schur bound with parameter $P \ge 8r$, frame Gram discriminant at least $c > 0$, and the specified chart bounds on the P^{-6} box, the integral of

$$\left\|\int \phi(y)F(\gamma_Z(uh)^{-1}y)\,dy\right\|^2$$

over the compact-factor variable h and the indicated unipotent cell variables u ($\|u\| \le P^{20}$), multiplied by $c_G c^{k-1}$, is bounded by $P^{2000s_0^2}D\int|\phi|^2$ when ϕ is supported in the P^{20} size region and D is the Schur mass bound. This is the chart-dependent inequality used by the `BesselSystem.open_cell` clause of Section 14.5. Our box estimate explains the required exponent scale and the order of differentiation; it does not on its own identify c_G or verify the boundedness of every Gram section.

15.9 A finite-subset version of adelic Poisson summation

We give a direct proof of the inequality actually used in the global sieve. Its hypotheses and conclusion differ from classical Poisson equality for a Schwartz function. For a number field E let $\mathbb{A}_E$ carry the self-dual measure relative to the usual trace character, so that the quotient $\mathbb{A}_E/E$ has measure one. The diagonal copy of E is a discrete cocompact subgroup and its annihilator is E. These structural properties are the background `WeilData` in the formal source. For a nonnegative real-valued L^1 function F on $\mathbb{A}_E$, define $\widehat{F}(\xi) = \int F(x)\overline{\psi(\xi x)}dx$.

Lemma 15.8 (One-sided Poisson)**.** *Assume F is lower semicontinuous at each member of a finite subset V of E. Then*

$$\sum_{t\in V} F(t) \le \sum_{\xi\in E} |\widehat{F}(\xi)|.$$

More precisely, if the nonzero coefficients have summable bounds $B(\xi)$, the right side may be replaced by $\int F + \sum_{\xi\neq 0} B(\xi)$.

Proof. Use the quotient Haar probability and the nonnegative periodization $P(x) = \sum_{t\in E} F(x+t)$, which is defined almost everywhere and belongs to $L^1(\mathbb{A}_E/E)$ by Tonelli. Its Fourier coefficients are $\widehat{F}(\xi)$, by Fubini and the annihilator identification. If the sum of their absolute values diverges the claim is immediate. Otherwise the Fourier series converges absolutely and uniformly to a continuous function P_0 and equals P almost everywhere. This last assertion follows from uniqueness of Fourier coefficients on a compact abelian group: characters span a uniformly dense subalgebra of the continuous functions, by conjugation

stability and separation of points, so the measure with density $P - P_0$ is zero. Given $\delta > 0$, lower semicontinuity gives a neighborhood O of zero on which $F(t + x) \geq F(t) - \delta$ for all $t \in V$. For almost every x in this neighborhood,

$$\sum_{t \in V} (F(t) - \delta) \leq P(x) = P_0(x) \leq \sum_{\xi} |\widehat{F}(\xi)|.$$

Nonempty open subsets of the quotient have positive measure. Now let $\delta \to 0$; the same argument with arbitrarily large bounds applies if some $F(t) = \infty$. □

The source supplies this inequality in `Lib/Adelic/Adeles/OneSidedPoisson.lean`. Finite subsets and lower semicontinuity permit the original orbital functions to be used without assuming differentiability at zero or at nonunit elements. For example a single upward spike at a rational point invalidates the inequality without lower semicontinuity: it contributes to the left side but is invisible to the Fourier transform. A downward spike is harmless for a one-sided bound.

We verify the Jacobian and the frequency support in the setting of the sieve. Let T be the finite auxiliary set of finite places, A_T the product of E_w at $w \in T$, and

$$U_T = \{x \in \mathbb{A}_E : x_w \in \mathcal{O}_w \ (w \notin T)\}.$$

This is open in the restricted product topology. For a function H on $\mathbb{C}^f \times A_T$ and positive or nonzero real scales l_{vi}, form

$$F(z) = \mathbf{1}_{U_T}(z) H((l_{vi} z_v)_{v|\infty}, (z_w)_{w \in T}).$$

Multiplication by l_{vi} on $\mathbb{C}$ has absolute real determinant l_{vi}^2. Off T, integrating the compact indicator introduces its Haar volume. Thus for the normalized measures on the local factors in the source,

$$\int F(z)\, dz = c_i \int H(y)\, dy, \qquad c_i = \frac{\mathrm{vol}(\widehat{\mathcal{O}})}{\prod_{v|\infty} l_{vi}^2}.$$

The finite-source convention is that the local integral volumes of the unit rings at T are one; equivalently one may absorb them into the normalization of H. In a nonzero-frequency integral the same change of variable replaces the infinite frequency by its quotient by l_{vi} (since l_{vi} is real). The off-T Fourier integral of an indicator of a compact open additive subgroup is zero outside its annihilator: translating by an element of the subgroup that has a nontrivial character multiplies the integral by a phase different from one. With the self-dual character divided by the source's different-normalizing element, the annihilator of the off-T integral lattice is the integral lattice. Thus the rational frequencies that remain are T-integers. The constants κ_v and g in the source convert the usual trace character to the local Bessel root characters; those constants occur in the frequency arguments, not in the absolute Jacobian factor. This is the measure and frequency content of `measurePreserving_scale_evalPlaces`, `extendToAdeles` and `one_sided_poisson` in `New/Filling/GlobalSieve.lean`.

15.10 An explicit induction for the sieve system

For the precise kernels defined in Section 14.13, we work with a finite alphabet A of admissible rational units (the elements and the relevant inverses integral away from T). The head maps $H_i(p, t)$, for $p \in A^i$ and next entry t, place these entries into the local product scaled by l_{va}, filling later entries with the chosen units. They are continuous on the unit locus. The bottom term is bounded by E_0; to simplify notation let $\lambda_i(p)$ be the nonnegative weight of a head of

length i, in the normalized coordinates. The scaling and filling maps are understood at all levels. We derive the finite-set induction used in `one_sided_poisson`; abstracting from those coordinate choices, it has two distinct error terms.

First, the positivity hypothesis at the top of a sieve system means that, for any finite family p_a with a common last entry, $\sum_{a,b} \overline{z}_a J_{s_0}(p_a, p_b) z_b \geq 0$. It is preserved by taking products over places: represent each matrix as the Gram matrix of a family e_a of vectors, and similarly f_a for a second matrix. The new Gram vectors $e_a \otimes f_a$ have pairings equal to the product of the original entries. Integration at zero frequency preserves the property for any fixed coefficients z_a, by linearity and integrability. It also preserves lower semicontinuity of these nonnegative sums, by Fatou's lemma. More explicitly, if the continuity hypothesis at the top is restricted to the open unit locus, a path of common head entries approaching a unit can be taken within this open set; the other integrated coordinates are units almost everywhere or replaced by the map `unitOrOne`. Iterating Fatou gives the required inequality even when the lower kernels have not been shown continuous separately. This proves the positivity and lower-semicontinuity content of `SieveSystem.pi` and `SieveSystem.nonneg_and_lowerSemicontinuous` in `New/Filling/EichlerFlag.lean`.

For a finite $V \subseteq A$ put

$$S_i(V) = \operatorname{Re} \sum_{t \in V} \sum_{p,q \in A^i} J_{i+1}(H_i(p,t), H_i(q,t)).$$

The summand after summation over heads is nonnegative, as just proved. We take the uniform nonzero-frequency assumption of the source in the form

$$\sum_{y \in X} \beta(\operatorname{ratPoint}(c_v / l_{vi}, y)) \leq \varepsilon, \tag{15.20}$$

for every finite set X of nonzero rational T-integers, for each i in the range, and for each family of admissible scale constants c_v (in `one_sided_poisson`, $|c_v|^{-1} \leq |\kappa_v|/g$). Here β is the product frequency majorant. The head sum assumption is

$$\left(\sum_{p \in A^i} \lambda_i(p) \right)^2 \leq \Gamma \prod_{a<i} c_a. \tag{15.21}$$

The zero-frequency step integral descends to J_i; at each nonzero frequency the sum over heads is at most the majorant times $(\sum \lambda_i)^2$, by triangle inequality. Fubini justifies this identity and the bound for a finite sum of head pairs: no convergence in the number of heads is being used. Lemma 15.8 and the Jacobian computation give

$$S_i(V) \leq c_i \operatorname{Re} \sum_{p,q \in A^i} J_i(p,q) + c_i \varepsilon \Gamma \prod_{a<i} c_a. \tag{15.22}$$

At $i = 0$ the sum on the right has one entry and the bound is $c_0(E_0 + \Gamma\varepsilon)$. For $i \geq 1$, split the right-hand sum according to the last entries of p and q. The equal-entry part is exactly an $S_{i-1}(A)$, not a sum of norms of individual head vectors. For the unequal-entry part the difference frequency at level i is $g^{-1}(t^{-1} - u^{-1})$. For fixed t and variable $u \neq t$, the differences are nonzero and distinct. They are rational T-integers before the stipulated normalization of the character. With

$$L_t = \sum_{p \in A^{i-1}} \lambda_i(p,t),$$

the absolute value of the unequal-entry contribution is bounded by

$$\sum_{t \neq u} \beta(g^{-1}(t^{-1} - u^{-1})) L_t L_u.$$

One can estimate this by the row and column bounds for the nonnegative frequency matrix: each row sum is at most ε by (15.20), and so is each column sum, because the differences there too are nonzero and distinct. Since $2L_tL_u \le L_t^2 + L_u^2$, the bound is $\varepsilon \sum_t L_t^2 \le \varepsilon(\sum_t L_t)^2$. Combining this with the Poisson error gives

$$S_i(V) \le c_i S_{i-1}(A) + 2\varepsilon\Gamma \prod_{a\le i} c_a.$$

Together with the base case this yields

$$S_i(V) \le \left(\prod_{a\le i} c_a\right)(E_0 + (2i+1)\Gamma\varepsilon).$$

This is the induction and the coefficient $2i+1$ in `one_sided_poisson` of `New/Filling/GlobalSieve.lean`. The complex forms $J_i(p,q)$ in the source include the head padding and unit replacements; the displayed shorthand drops those redundant arguments. The term $2i+1$ accounts for one Poisson error at every floor and a different error each time the previous floor is split by its last entries. Bounding each Gram entry in absolute value at the top would lose the exact identity between floors, which supplies the leading term E_0.

15.11 Frequency summation and the product formula

The bound (15.20) would be useless if it contained the size of A or the sizes of denominators of its members. The damping exponent 6 is designed to avoid either dependence after product-formula cancellation.

Let $y \in E^\times$ be integral outside the conjugation-stable set T. In a local field write $|y|_w$ for the module normalized by the cardinality of the residue field, and $[y]_T = \prod_{w\in T} |y|_w^{-1}$ for the generalized index. For a complex place v, $|v(y)|$ denotes the ordinary absolute value. The product formula gives

$$\prod_{v|\infty} |v(y)|^2 = \prod_{w \text{ finite}} |y|_w^{-1} \ge [y]_T.$$

The inequality uses integrality away from T, not integrality at T. In particular the potentially small or large quantity $[y]_T$ can be nonintegral. Write $L_K(u) = u^K \min(1,u)^6$ for the damped power. At scales satisfying $|c_v|^{-1} \le |x_v| \le B$, $B \ge 1$, we get

$$\prod_v L_K(|c_v v(y)|^{-1})\,[y]_T^{K/2} \le \prod_v |x_v|^K \min\left(1, \frac{B}{|v(y)|}\right)^6.$$

The remaining local factor satisfies

$$\min(1, B/t)^6 \le \frac{(1+B^2)^3}{(1+t^2)^3} \quad (t>0).$$

If $t \le B$, the right side is at least one. If $t > B$, use $B^2(1+t^2) \le (1+B^2)t^2$ and cube. Thus each local term is controlled by the inverse cube of the natural two-dimensional vector norm $1+|v(y)|^2$, with the factor B^6 separated.

The denominator index $\mathrm{den}(y)$ is the index of the ideal $\{z \in \mathcal{O}_E : zy \in \mathcal{O}_E\}$, with the finite contribution outside the stipulated places taken according to the lattice convention. The point y determines a rational line in E^2 via $(1,y)$, and the lattice on that line is formed by multiplying $(1,y)$ by the denominator lattice. If $y = a/b$ is rational in lowest terms, this is simply the primitive vector (b,a), whose length squared is b^2+a^2; indeed $b^6(1+|a/b|^2)^3 = (a^2+b^2)^3$.

The corresponding sum over distinct lines converges: in concentric integer shells of maximum coordinate n, the number of pairs is at most $8n$ and each summand is at most n^{-6}. This illustrates why the same denominator power and the same infinite-place damping are used, rather than unrelated tapering functions at the two kinds of places.

The general CM assertion also needs the lattice covolumes and ideal classes. We retain as formal-source input the rank-one sublattice discriminant sum `tsum_sublatticeDisc_inv_le` and discriminant comparison `sublatticeDisc_mul_index_sq_le` used in the proof of `sum_frequencies_le` in `New/Filling/GlobalSieve.lean`. In that declaration, under WeilData and conjugation stability of T, the precise consequence is

$$\sum_{\substack{y\in X\\ y\neq 0}} \frac{[y]_T^{K/2}}{\operatorname{den}(y)^6} \prod_v L_K(|c_v v(y)|^{-1}) \le \left(\prod_v |x_v|^K (bB^2)^3\right) (b\,\mathrm{rd}_E)^{3f}$$

for any finite subset of nonzero T-integers. The explanation above proves the product-formula and damping portions of the inequality; the full sublattice discriminant sum is more than ordinary Euclidean lattice-point counting, because one must count rational lines with nonprincipal denominator ideal.

In the source $K = k - s_0^7$, and the finite multiplier defining `frequencyGain` is exactly $[\xi]_T^{K/2} \operatorname{den}(\xi)^{-6} b^{s_0^9 f}$. Thus the large power K disappears from the summation at the cost of the explicit powers of the scales x_v. The other weight bound (15.21) comes from summing over the successive head products in `sq_sum_globalWeight_le` (same file). Those products can be used as summation coordinates because knowing all successive products and the leading unit recovers each entry by division; denominators are again accounted for by the integral-lattice condition. The numerical comparison with the parameters of `Filling.IsAdmissible`, including these head bounds and frequency sums, is the remaining formal-source input `sieve`. It gives the error $(2h+1)\Gamma\varepsilon \le b^{-6kf-10}/d_\infty$ and the leading constant of Proposition 14.13. The part now proved directly is the one-sided Poisson argument and its induction, with the correct covolume factor; it requires the uniform frequency and head sums just described, not merely absolute summability for each fixed rank.

15.12 Finite packets, their normalization, and the global comparison

We expand the finite-place mechanism in the same notation as Section 14.13. At a finite product of local fields, a compact-open level permits a locally constant vector function to be integrated by counting cosets. Here is the elementary step used in constructing the top kernel. If f is continuous and locally constant on a compact open subset D of a finite product of nonarchimedean fields, cover D by finitely many translates of compact open additive subgroups on each of which f is constant; intersect those subgroups. The result is a compact open subgroup L with D a finite union of its cosets, on each of which f is constant. Consequently

$$\int_D f(t)\,dt = \operatorname{vol}(L) \sum_{t+L\subset D} f(t).$$

For a finite family of functions the same subgroup works, by further finite intersection. If the functions are indexed by finitely many cosets of a compact packet group, the same argument applies simultaneously to their pairings. This is the basic step-function mechanism in `New/Filling/FiniteSystem.lean`. It identifies the normalized top floor as an averaged Gram matrix, not just the indicator of an algebraic equation over a residue field.

There is an independent counting problem for the mass at the bottom of the tower. In the lattice model the quadratic condition for a last-vector packet is of the form

$$Ny^{*}y = t\bar{t} - N^{2}.$$

This follows by pairing the last frame vector with itself in the changed basis and substituting its coordinates in the positive middle space. Thus the number of choices for the k-component vector in the packet is subject to a Hermitian norm condition. For a finite field of square cardinality q^2, the equation $\sum_{j=1}^{k} x_j\overline{x_j} = a \neq 0$ has $q^{2k-1} - (-1)^k q^{k-1}$ solutions: for $k = 1$, each nonzero a is the norm of exactly $q + 1$ elements. If Z_k and N_k denote the counts at zero and at a fixed nonzero value, respectively, then $Z_k + (q-1)N_k = q^{2k}$ and adding one coordinate gives

$$N_{k+1} = N_k + (q+1)(Z_k + (q-2)N_k) = (q+1)q^{2k} - qN_k.$$

This recursion proves the formula by induction, and subtraction gives the corresponding zero count. It demonstrates the high-dimensional equidistribution in the packet counting step, but it does not prove the particular local mass comparison when congruences, ramified factors and higher floors are imposed. The latter uses the weighted quadric integral and the passage between successive floors in the source.

Precisely, `Filling.IsFillingLevel.exists_finite_sieveSystem` of `New/Filling/FiniteSystem.lean` constructs an intermediate packet group Z containing all level stabilizer elements whose reductions belong to the prescribed fixing set, and a compact open additive subgroup L of the finite-local product. Its kernels have the form

$$J_j^{\mathrm{fin}}(p,q) = \begin{cases} \mathrm{floorForm}_j(p,q;Z)/\mathrm{massRatio}(Z), & j < s_0, \\ \mathrm{topKernel}(L,Z;p,q)/(\mathrm{vol}(L)\mathrm{massRatio}(Z)), & j = s_0. \end{cases}$$

The upper kernel is a Gram kernel at a common top entry. The bottom bound in that declaration is

$$(1 + b^{-6})/(b^{6k})^{|S_g|}.$$

The head weights are `finHeadWeight`, and the nonzero frequency majorant is `frequencyGain` in Section 15.11. The exact floor and mass comparison, including `sum_blocks_le_topKernel`, is a formal-source input. The factor $1/\mathrm{vol}(L)$ compensates for the additive coset volume in the integration identity. The factor $1/\mathrm{massRatio}(Z)$ accounts for the measure of the fixed-vector subgroup; an arbitrary upper bound for its cardinality could not be used instead, because the subsequent product with the archimedean leading terms must give the stated main-energy normalization.

To compare the local orbital integrals with the rational amplitude, one first removes the compact unipotent class cutoffs by unfolding, then collects all rational terms lying in the prescribed finite cosets. For a convergent sum, partitioning the rational orbits into those cosets is just reindexing the sum. The height restrictions and quotient cutoff estimates ensure convergence for the actual amplitude. The identity `IsFillingLevel.correlation_smul_eq_cosetAmplitude` in `New/Filling/GlobalToLocal.lean` gives the exact conversion. For a term s at the last floor, local head q obtained by reading a rational head p at T, and packet cosets a, b satisfying the level-congruence relation, it equates

$$\mathrm{correlation}_{Z,q}(a,b) \otimes \mathrm{besselCorrelation}_{\infty}(\gamma) = D_{s_0}(q,1)\ \mathrm{vol}(K(g))\,\mathrm{cosetAmplitude}_T(\gamma).$$

The left arguments in the source include the inverse of the local term position of s, and γ denotes the rational Weyl-head translate of s. This is a formal-source input: proving it

requires the changes of variables in the exact integral normalization at all places. It specifies which mass factors survive the global comparison. Taking the product sieve system and applying Section 15.10 to the common-cell entries of that amplitude is therefore a comparison of normalized Gram matrices with the prescribed bottom constant, rather than a pointwise comparison of positive scalar functions.

15.13 Degenerate terms and the accumulation of energy errors

For completeness we supply two small calculations that govern the nonmain cell estimates. The first is character compatibility. Consider an unfolded integral on a double orbit with left and right Bessel characters. Translating within their common stabilizer by s leaves the cell and measure unchanged, but changes the integrand by the quotient $\eta(s)$ of the two characters. If $\eta(s) \neq 1$, the integral I must satisfy $I = \eta(s)I$, hence $I = 0$. For rational orbit sums with unipotent class cutoffs the same argument applies after averaging over the stabilizer, because the sum of the translates of the cutoff equals one. The precise source identity, using stable sets of rational terms and the hypotheses of admissibility, is `Filling.amplitudeSum_eq_zero_of_not_charsAgree` in `New/Filling/CrudeEnergy.lean`. We use the cutoff-invariance statement there; pointwise invariance of an arbitrarily chosen cutoff is not assumed.

To check the Hermitian geometry behind the compatible degenerate terms, use the convention that H is conjugate-linear in its first variable. For $H(x,x) = H(x,v) = 0$ and $\alpha + \overline{\alpha} = 1$ the transformation

$$T(y) = y + xH(v,y) - (v + \alpha H(v,v)x)H(x,y)$$

preserves H. Indeed write $l(y) = H(x,y)$, $m(y) = H(v,y)$ and $a = H(v,v)$ real. Then $T(y) = y + x(\overline{m(y)} - \alpha a\overline{l(y)}) - vl(y)$. Expanding $H(T(y),T(z)) - H(y,z)$, the cross terms $\overline{l(y)}m(z)$ and $\overline{m(y)}l(z)$ cancel, leaving $a(1-\alpha-\overline{\alpha})\overline{l(y)}l(z) = 0$. Its difference from the identity is nilpotent, since its image is contained in the span of x, v, sends x to zero and sends v to a multiple of x. It is thus an invertible unitary transformation. Choosing x in the appropriate shared flag gives the source's elementary Eichler map. The precise placement in the Bessel group is `exists_mem_besselU_eq_eichlerMap` in `New/Filling/EichlerFlag.lean`; the paired-coordinate propagation with compatible characters is `hasCommonFlagVector_of_le` in `New/Filling/CrudeEnergy.lean`. Their flag-dependent consequence `Filling.CharsAgree.exists_hasLevel` is a formal-source input: a rational term with compatible characters and discriminant zero has a level below the top open cell in the sense of `Filling.HasLevel`. The explicit calculation shows why preservation of the indefinite form in those test transformations does not itself impose any restriction on their Bessel characters. It is the additional character compatibility that supplies the relation between the two flags.

The second calculation is the handling of many error strata. The amplitude spaces in the energy definition are Hilbert spaces (with the common positive scalar unipotent-covolume normalization). For $\delta > 0$,

$$\|A + B\|^2 \leq (1+\delta)\|A\|^2 + (1+\delta^{-1})\|B\|^2. \tag{15.23}$$

This is obtained from $2|\langle A, B\rangle| \leq \delta\|A\|^2 + \delta^{-1}\|B\|^2$. It applies to integrable families by taking their Hilbert space norms; it does not require orthogonality of the strata.

We recall the actual pair-error bound from the source. Let C satisfy the prescribed size inequalities (at least the Bessel energy, with stronger bounds τ^{20} times that energy at level zero and τ^{12} times the main energy divided by X^k if the retained set of links omits one). The

input `Filling.energy_le_of_level_pairs` in `New/Filling/CrudeEnergy.lean` says that for a stable set of a common level j, the cut energy is at most

$$w + \frac{C}{b^{30}4^{j^{\mathrm{rev}}}},$$

if the retained coherent link energy divided by the square of the unipotent covolume is at most w. The reversed level index is the distance from the top of the finite tower. This inequality uses the summable pair weights, the exact assertion `sum_pairWeight_le` in `New/Filling/BlockBound.lean` being their total mass at most $2\tau^2$. Incoherent links contribute the arithmetic error rather than being discarded as orthogonal. Those lattice and pair estimates are retained as formal-source inputs.

To sum all levels strictly below a floor, begin with zero energy below the first level (the incompatible terms vanish by the character identity). If the error budget at one floor is E, the previous floor budget is $E/4$. The new single-level budget is likewise $E/4$ by the pair bound with no retained coherent link energies. Apply (15.23) with $\delta = 1$ to get $2(E/4)+2(E/4) = E$. This proves the induction of `Filling.energy_le_of_level_lt` with the geometric factor $4^{j^{\mathrm{rev}}}$; there is no factor proportional to the number of levels.

One still needs a bound for replacing the locally cut functions by the full ones after global unfolding. The formal-source estimate `energy_sub_cut_le` in `New/Filling/SchurTest.lean` bounds this difference energy by b^{-30} times the Bessel energy. It uses the quotient tail and transported majorants, not just the local inequality (15.19); sums of translates at distinct rational points would prevent the latter alone from proving the former. Now the cut top-cell energy is at most $w + Cb^{-30}$ and the lower-cell cut energy at most Cb^{-30} by the preceding induction. A first use of (15.23) with $\delta = b^{-4}$ bounds the total cut energy by

$$(1 + b^{-4})w + (2 + b^4 + b^{-4})Cb^{-30}.$$

A second use with the same δ, for full minus cut, shows that the full energy is at most

$$(1 + b^{-4})^2 w + \Big((1 + b^{-4})(2 + b^4 + b^{-4}) + 1 + b^4\Big)\, Cb^{-30}.$$

Since $b \geq 1000$, the error coefficient is less than b^{-12}. This gives precisely the combination in `Filling.energy_le_of_linkEnergy_le` of `New/Filling/CrudeEnergy.lean`. The retained pair estimate, cutoff transfer, flag-level theorem and link/coherent comparison are substantial formal-source inputs. The geometric summation and energy combination are elementary deductions. In particular the local Fourier sieve supplies the only non-negligible w; the other cells are controlled separately with large exponents. Substituting the coherent bound into this deduction gives the sharp energy bound of Section 14.11 and the subsequent `energy_fixingSet_le` calculation, with the arithmetic numerical comparisons of the specialist chapter left unchanged.

Chapter 16

Quantitative Bessel construction and the arithmetic of the filling datum

The filling argument is unusually sensitive to the normalization of local integrals and the exact exponent in their bounds. The proof of the first and second moments in Chapter 14 and the local mechanism in Chapter 15 left a number of these statements as formal-source inputs. In this chapter we reconstruct further parts of those arguments: the explicit Gaussian approximation of the discrete-series coefficient, the simultaneous removal of compact directions and the phase-cancelling realization of the Bessel subgroup, the arithmetic level and its unipotent covolume, the elimination of vanishing-discriminant terms by a flag argument, the height summation in the global sieve, and the finite-group covariance required for the moments. We maintain all quantitative bounds needed by those arguments. There remain local Fourier-floor estimates, some uniform arithmetic pair-counting estimates, and the exact global packet comparison; these are separately identified rather than inferred from the results proved here. Standard representation-theoretic results used here include the identification of the normalized lowest-type coefficient with the Flensted–Jensen integral, its square-integrable representation, the Schur orthogonality and Harish-Chandra degree formulas. These are logically different from the particular quantitative estimates proved below.

Throughout the chapter the hyperbolic rank is $h = \varrho$, $n = h + 1$, the dimension of the definite part is k, the full rank is $r = k + 2h + 1$, the number of complex places is f, $g = q_0 q_1$ is the product of the two split rational level primes, and N is the separate split prime used for the marked lattice. In the discrete-series calculations, where we use matrices of varying sizes, we use d for a matrix rank, not for the formal degree, which we denote by d_λ or d_∞ for the product of degrees. For a matrix A let A^* denote its conjugate transpose, $\|A\|_2$ its Frobenius norm, and $\|A\|_{\mathrm{op}}$ the Euclidean operator norm. This distinction is used throughout the estimates.

16.1 The compact integral and the change of splitting

Let $J = \mathrm{diag}(\epsilon_i)$, where each $\epsilon_i = \pm 1$, and let $G_J = \{g : g^* J g = J\}$ and $K = G_J \cap U(d)$. The doubled Harish-Chandra parameter ℓ has distinct entries, with $\ell_i + d - 1$ even. Put $\lambda = \ell/2$, let ρ be the half sum of roots positive for the order of λ, and ρ_c the half sum of the compact positive roots (the roots $e_i - e_j$ with $\epsilon_i = \epsilon_j$). Thus

$$\mu = \lambda + \rho, \qquad \Lambda = \mu - 2\rho_c$$

are integral. In particular μ has the same strict order as ℓ. The weight Λ is dominant in each compact block for this order. The parity and distinctness here are essential: the compact type with highest weight Λ is an honest representation, not merely a character on a cover.

For an integral weight w and square matrix A write $m_w(A)$ for the product of principal minors used in the source:

$$m_w(A) = \det(A)^{-a} \prod_{t=-a+1}^{a} \det A_{\{i:w_i \geq t\}}, \qquad a = \sum_i |w_i|.$$

Empty principal minors have determinant 1. This is independent of enlarging a. If all the principal minors in the product are invertible it obeys

$$m_{-w}(A) = m_w(A^{-1}), \qquad m_{w+c}(A) = m_w(A)\det(A)^c. \tag{16.1}$$

For the first identity apply Jacobi's complementary-minor identity at each threshold, pairing threshold t with the complementary threshold $1-t$; the determinant factors compensate because $(A^{-1})_S$ has determinant $\det A_{S^c}/\det A$. The second follows by shifting the thresholds.

Let a be invertible block diagonal for J and z have strictly positive Hermitian part, so that every principal minor of u^*zu is nonzero for $u \in K$. Put

$$F_\ell(a,z) = \frac{1}{Z_\Lambda}\int_K \frac{m_\Lambda(au)\overline{m_\Lambda(u)}}{m_{-\mu}(u^*zu)}\,du, \qquad Z_\Lambda = \int_K |m_\Lambda(u)|^2\,du.$$

Here the Haar measure on K is probability. This is `HCParam.flenstedJensen` in `Lib/DiscreteSeries/FlenstedJensen/RepresentationDefs.lean`. If $g \in G_J$, write a_g for its block diagonal part. The representation-theoretic identification used in the source defines its coefficient by

$$\psi_\lambda(g) = F_\ell(a_g, a_g^{-1}g).$$

We rely on the lowest-type coefficient theorem in this identification, not on an analytic continuation assertion about arbitrary highest-weight functions.

We prove directly the splitting property of the integral. This proof is useful because the pointwise Gaussian calculation uses a different splitting in which the large weights occur only in the denominator.

Lemma 16.1 (Change of splitting)**.** *Let g be a fixed complex invertible matrix. On the set of block diagonal matrices a such that $g^{-1}a$ has strictly positive Hermitian part, the value $F_\ell(a, a^{-1}g)$ is constant.*

Proof. The domain is convex: positivity of the Hermitian part of $g^{-1}a$ is a convex condition and implies that a is invertible. Positivity is equivalent for an invertible matrix and its inverse, by conjugating the Hermitian part with the inverse. Hence both arguments of the integrand are valid on the domain. Rewrite it, by (16.1), as

$$\Phi(a,u) = \frac{m_\Lambda(au)m_\Lambda(u^{-1})}{m_\mu(u^{-1}g^{-1}au)}.$$

The integral of this function is locally holomorphic in the block diagonal matrix a: all denominators are nonzero, and on the compact set K boundedness of the derivatives on a sufficiently small neighborhood permits differentiation under the integral. For the same reason the integrand extends holomorphically in u to a neighborhood of K in its complexified group.

Let X be a block diagonal variation of a and write $Y = u^{-1}a^{-1}Xu$. At fixed u the derivative in direction X of Φ is the sum of the derivatives in directions $auE_{ij}u^{-1}$ multiplied

by Y_{ij}, over pairs with $\epsilon_i = \epsilon_j$. If $\ell_i \geq \ell_j$ then right multiplication by $1 + tE_{ij}$ is an upper triangular elementary matrix (ordering the weights in descending order), and the principal-minor transformation law gives the derivative

$$D_{auE_{ij}u^{-1}}\Phi(a,u) = \delta_{ij}(\Lambda_i - \mu_i)\Phi(a,u).$$

Indeed for an upper triangular elementary right factor each high-weight principal submatrix is multiplied on the right by its restricted triangular factor; the diagonal case multiplies by $(1+t)^{\Lambda_i-\mu_i}$.

Suppose instead that $\ell_i < \ell_j$. A lower triangular unipotent left factor leaves the high-weight principal minors unchanged, since the principal submatrix of each initial segment depends only on the corresponding rows of the left factor. Under $u \mapsto u(1+tE_{ij})$, the left factor $(1+tE_{ij})^{-1}$ therefore drops out of both $m_\Lambda(u^{-1})$ and the denominator. The derivative of Φ in that u-direction equals $D_{auE_{ij}u^{-1}}\Phi$. In the same variation, Y is conjugated by $1+tE_{ij}$, so the derivative of its ij-entry is $Y_{ii} - Y_{jj}$. Haar right invariance implies

$$\int_K \frac{d}{dt}\Big|_{t=0} \big(\Phi(a, u(1+tE_{ij}))Y_{ij}(u(1+tE_{ij}))\big)\, du = 0.$$

This follows first for real Lie algebra directions and hence for E_{ij} by complex linearity of the holomorphic derivative (the Lie algebra of each complexified compact block is the complexification of that of the compact block). We conclude

$$\int_K Y_{ij} D_{auE_{ij}u^{-1}}\Phi\, du = -\int_K \Phi(Y_{ii} - Y_{jj})\, du \quad (\ell_i < \ell_j).$$

Sum over the indicated opposite triangular directions. For each i the coefficient of Y_{ii} in $\sum_{\ell_i<\ell_j,\epsilon_i=\epsilon_j}(Y_{ii} - Y_{jj})$ is $-2(\rho_c)_i$. Thus their contribution to the derivative is $+\int_K \Phi \sum_i 2(\rho_c)_i Y_{ii}\, du$. The diagonal contribution from the first formula is $\int_K \Phi \sum_i (\Lambda_i - \mu_i) Y_{ii}\, du$, precisely its negative. The total derivative vanishes along every variation X. Constancy follows from convexity. □

This is the argument of `flenstedJensen_eq_of_accretive` and `dsCoeff_eq_flenstedJensen` in `Lib/DiscreteSeries/FlenstedJensen/FlenstedJensenInvariance.lean`. In particular the half sum of the compact positive roots accounts exactly for the derivative of the conjugating matrix in the compact integral. Replacing Λ with λ would not give splitting invariance.

The following useful consequence isolates the triangular shifts. If L is strictly lower triangular in the order of descending ℓ and block diagonal for J, with X and $(1-L)X$ accretive, then

$$F_\ell(1,(1-L)X) = F_\ell(1,X). \tag{16.2}$$

Indeed apply Lemma 16.1 to $(1-L)X$, using the splitting $a = 1-L$, and use $m_\Lambda((1-L)u) = m_\Lambda(u)$. Taking conjugate transpose shows the analogous invariance under a right factor $1-L^*$, since $F_\ell(1,X^*) = \overline{F_\ell(1,X)}$. The adjoint statement is `flenstedJensen_shift_eq_conj` in `Lib/General/LinearAlgebra/FramedCoefficient.lean`.

16.2 Logarithmic principal-minor estimates by resolvents

The delicate feature of the next calculation is that the error must be proportional to the *squared commutator* rather than to a uniform cubic matrix error. This is what permits repeated or nearly equal compact weights. We give a proof using the derivative of the compressed logarithm.

Lemma 16.2 (Compressed logarithm with an absolute error)**.** *Let P, Q be orthogonal projections, $\|X\|_2 \le \delta \le 1/2$, and suppose*

$$\operatorname{tr}([P,Q]Y) = 0, \qquad \|[P,X] - Y\|_2 \le \delta c, \qquad \|Y\|_2, \|Q-P\|_2 \le c.$$

Then

$$\log\det(1+PX) = \operatorname{tr}(Q\log(1+X)) - \tfrac{1}{4}\operatorname{tr}(Y^2) + e, \qquad |e| \le 100\delta c^2,$$

where the logarithm of the determinant denotes the branch obtained by continuation along $1+tPX$, $0 \le t \le 1$.

Proof. The trace logarithm is valid because the operator norm of X is at most δ. Write X in blocks relative to P: its blocks are A, B, C, D. For $0 \le t \le 1$, put $U = (1+tA)^{-1}$ and $W = (1+tX)^{-1}$. The PP-block of the inverse satisfies

$$W_{PP} - U = t^2 U B W_{P^\perp P^\perp} C U.$$

This follows by multiplying the inverse equations, or twice using the equations for the off-diagonal inverse blocks. Differentiate the trace logarithms of $1+tPX$ and $1+tX$, compressed by P in the second case. The difference of their derivatives is $t^{-1}\operatorname{tr}(W_{PP} - U)$ (at $t = 0$ use the continuous extension). Hence

$$\operatorname{tr}\log(1+PX) - \operatorname{tr}(P\log(1+X)) = \int_0^1 t\,\operatorname{tr}(UBW_{P^\perp P^\perp}CU)\,dt. \tag{16.3}$$

The inverses have operator norm at most 2; their difference from the identity has operator norm at most $2t\delta$. Replace the three inverse factors in the integrand successively by 1 and use $|\operatorname{tr}(B'C')| \le \|B'\|_2\|C'\|_2$. The error is at most $14t\delta\|B\|_2\|C\|_2$ before multiplication by t. Thus the right side of (16.3) differs from $\frac{1}{2}\operatorname{tr}(BC) = -\frac{1}{4}\operatorname{tr}([P,X]^2)$ by at most

$$\tfrac{14}{3}\delta\|B\|_2\|C\|_2 \le 3\delta\|[P,X]\|_2^2.$$

We next estimate the trace incurred by replacing P with Q. Let $L = Q - P$, and separate its diagonal and off-diagonal blocks relative to P. The diagonal blocks are respectively negative and positive semidefinite. The sum of their trace norms is

$$\operatorname{tr}(P - PQP) + \operatorname{tr}(P^\perp Q P^\perp) = \|P - Q\|_2^2.$$

Their trace against $\log(1+X)$ has absolute value at most $2\delta c^2$. For the off-diagonal block part L_{off} one has

$$\operatorname{tr}(L_{\mathrm{off}}X) = -\operatorname{tr}([P,Q][P,X]).$$

The hypothesis on Y and $\|[P,Q]\|_2 \le \|L\|_2$ give an absolute bound δc^2. The logarithm series also shows

$$\|[P, \log(1+X) - X]\|_2 \le \sum_{m\ge 2} \delta^{m-1}\|[P,X]\|_2 \le 2\delta\|[P,X]\|_2.$$

Here the commutator of X^m is a sum of m terms with a single commutator of X. The off-diagonal part of a matrix has the same Frobenius norm as its commutator with P, so the remaining trace is at most $2\delta c\|[P,X]\|_2$.

Finally $\|[P,X]\|_2 \le (1+\delta)c \le 3c/2$. Replacing its trace square by $\operatorname{tr}(Y^2)$ costs at most $\|[P,X] - Y\|_2(\|[P,X]\|_2 + \|Y\|_2) \le (5/2)\delta c^2$. Adding these estimates, the factor 100 is more than sufficient. □

In the source this is `exists_det_compress_eq_exp_logOneAdd`, `New/DiscreteSeries/MinorExpansion.lean`. In the proof the resolvent identity avoids a factor exponential in the matrix rank. All estimates use a Hilbert–Schmidt product only twice; other factors are measured in operator norm.

Proposition 16.3 (Triangularly shifted principal minors)**.** *Let $g = 1 + X \in G_J$ with $\delta = \|X\|_2 < 10^{-3}$. There is an accretive y such that $\psi_\lambda(g) = F_\ell(1, y)$ and, for each projection P onto an initial segment of the* ascending *ℓ-order, each $u \in K$,*

$$\det((u^* y u)|_{\operatorname{im} P}) = \exp\left(\operatorname{tr}(P \log g) + \tfrac{1}{4}\|[P, X]\|_2^2 + e_P\right),$$

with

$$|e_P| \le 200000\delta(\|[P, X]\|_2^2 + \|[P, u]\|_2^2).$$

Here the compressed matrix in the determinant is $Pu^* y u P|_{\operatorname{im} P}$.

Proof. Order the indices by descending ℓ. Let L and R be the strictly lower and upper parts, respectively, of the same-sign block part of X. Put

$$y = (1 - L)(1 + X)(1 - R).$$

Since $\|L\|_2, \|R\|_2 \le \delta$, $\|y - 1\|_2 \le 4\delta$, and both intermediate shifted matrices are accretive. Lemma 16.1 first yields $\psi_\lambda(g) = F_\ell(1, g)$, and (16.2) and its right-hand version give $F_\ell(1, y) = F_\ell(1, g)$.

Write $A = [P, X]$, $B = [P, L]$, $C = [P, R]$ and $A_o = A - B - C$. Thus A_o is the opposite-sign part of A because P is diagonal. The product rule for the commutator shows

$$\|[P, y - 1] - A_o\|_2 \le 8\delta\|A\|_2.$$

Explicitly, $[P, y] = -B(1 + X)(1 - R) + (1 - L)A(1 - R) - (1 - L)(1 + X)C$; comparing the three terms to $-B, A, -C$ incurs respectively at most $(2\delta + \delta^2)\|B\|_2$, $(2\delta + \delta^2)\|A\|_2$, and $(2\delta + \delta^2)\|C\|_2$. Each commutator of a strictly triangular mask is a submatrix mask of A. The same estimate holds with P replaced by $1 - P$ and A_o negated.

Put $P' = uPu^*$. The norm $\|P' - P\|_2 = \|[P, u]\|_2$. Consequently

$$\|[P', y - 1] - A_o\|_2 \le 8\delta(\|A\|_2 + \|[P, u]\|_2). \tag{16.4}$$

Moreover P', P are same-sign block diagonal, whereas A_o is opposite-sign block off diagonal, so $\operatorname{tr}([P', P]A_o) = 0$. Lemma 16.2 applies with first projection P', second P, matrix $y - 1$, radius 4δ, and $c = 4(\|A\|_2 + \|[P, u]\|_2)$. It yields

$$\log\det(1 + P'(y - 1)) = \operatorname{tr}(P \log y) - \tfrac{1}{4}\operatorname{tr}(A_o^2) + e_0,$$

with $|e_0| \le 12800\delta(\|A\|_2^2 + \|[P, u]\|_2^2)$.

The complementary projection $P^\perp$ corresponds to a high-weight initial segment in the descending order. Its principal determinant is not changed by the lower triangular left or upper triangular right unit factors. The determinant of the full matrix is likewise unchanged. These equalities hold for the corresponding analytic logarithms: use the path $(1 - tL)(1 + tX)(1 - tR)$, for $0 \le t \le 1$, to fix the branches. Apply the same compression lemma twice with both projections $P^\perp$, first to X with $Y = -A$ and then to $y - 1$ with $Y = -A_o$. With error bounds $100\delta\|A\|_2^2$ and $6400\delta\|A\|_2^2$ respectively this gives

$$\operatorname{tr}(P \log y) = \operatorname{tr}(P \log g) + \tfrac{1}{4}\operatorname{tr}(A^2 - A_o^2) + e_1, \qquad |e_1| \le 6500\delta\|A\|_2^2.$$

To evaluate the quadratic term, put $W = (1+X)^{-1}$. The relation $g^* = Jg^{-1}J$ and the commutator rule for inverses imply $A^* = JWAWJ$. Since $\|W-1\|_{\mathrm{op}} \le \delta/(1-\delta)$, this gives

$$\|A^* - JAJ\|_2 \le 3\delta\|A\|_2.$$

On the same-sign part $JAJ = A$ and on the opposite-sign part it equals $-A$. Taking traces against A and A_o respectively therefore proves

$$\operatorname{tr}(A^2) - 2\operatorname{tr}(A_o^2) = \|A\|_2^2 + \varepsilon, \qquad |\varepsilon| \le 9\delta\|A\|_2^2.$$

Combining the three formulas and using $\det(1+P'(y-1)) = \det(Pu^*yuP|_{\operatorname{im} P})$ proves the assertion with the stated constant. □

This is `exists_principalMinor_eq_exp` in `New/DiscreteSeries/MinorExpansion.lean`. The exact sign of the quadratic term was determined by the signed adjoint identity. Without the triangular shifts, it would not in general have this sign.

16.3 Gaussian shape of the coefficient

For a real weight vector w set $Q_w(A) = \sum_{ij} |w_i - w_j||A_{ij}|^2$. As shown in Section 15.3, summing commutator squares over threshold projections gives Q_w. We now apply Proposition 16.3 to the weight $w = -\mu$, whose upper threshold sets are the required lower ℓ-sets.

Theorem 16.4 (Uniform near-identity Gaussian)**.** *Let $|\ell_i| \le L$, and let*

$$T \ge \max(1, d^2, d^2\log(d+L+1)), \quad \delta T \le 10^{-12}, \quad \|g-1\|_2 \le \delta.$$

Under the discrete-series hypotheses,

$$\left| \frac{\psi_\lambda(g)}{\exp(\operatorname{tr}(\operatorname{diag}(\Lambda)\log g))} - \exp(-Q_{-\mu}(g-1)/4) \right| \le \tfrac{1}{10}\exp(-Q_{-\mu}(g-1)/4) + e^{-100T}.$$

Proof. Write $X = g-1$, $q = Q_{-\mu}(X)/4$ and $p(u) = Q_\Lambda(u)/2$. Sum the logarithmic estimates for the shifted principal minors, using the determinant correction in $m_{-\mu}$. The trace of $\log y$ equals that of $\log g$. By the threshold identity and the formula for Λ,

$$\frac{\psi_\lambda(g)}{\exp(\operatorname{tr}(\operatorname{diag}(\Lambda)\log g))} = \frac{e^{-q}}{Z_\Lambda}\int_K |m_\Lambda(u)|^2 e^{-\mathcal{R}(u)}\,du, \tag{16.5}$$

where $|\mathcal{R}(u)| \le 10^6\delta(q + p(u) + d^2)$. To verify the constant, the sum of the errors in Proposition 16.3 is bounded by $200000\delta(Q_{-\mu}(X) + Q_{-\mu}(u))$. Now $|\mu_i - \Lambda_i| \le d$, which implies $Q_{-\mu}(u) \le Q_\Lambda(u) + 2d\|u\|_2^2 = 2p(u) + 2d^2$. The trace incurred by replacing μ by Λ is at most $2d^2\delta$, using $\|\log g\|_2 \le 2\delta$. These inequalities give the displayed bound on $\mathcal{R}$.

On the compact group the minor bound of Lemma 15.5 gives $|m_\Lambda(u)|^2 \le e^{-p(u)}$. Schur orthogonality for the compact highest-weight coefficient gives $Z_\Lambda^{-1} = \dim V_\Lambda$, the dimension of the compact irreducible of highest weight Λ. For completeness, the function $m_\Lambda(u)$ is precisely the unit highest-vector coefficient: write the dominant weight as a sum of successive fundamental weights and a determinant weight on each sign block; take the tensor product of the corresponding unit exterior highest vectors. The tensor highest vector generates the desired irreducible, and its coefficient is the product of the principal minors. The Weyl dimension formula, $\max|\Lambda_i| \le (L+3d)/2$, and the fact that there are at most $d(d-1)/2$ compact positive roots give

$$\dim V_\Lambda \le (L+4d+1)^{d(d-1)/2} \le e^{2T}.$$

If $|\mathcal{R}(u)| \le 1/20$, then $|e^{-\mathcal{R}(u)} - 1| \le e^{1/20}/20 < 1/10$. Its contribution to the difference between (16.5) and e^{-q} is at most $e^{-q}/10$, using total probability of the normalized density. Otherwise the bound on $\mathcal{R}$ gives $q + p(u) > 49999T$. On this complementary set the absolute value of the integrand in the difference, including $Z_\Lambda^{-1} e^{-q}$, is at most

$$2e^{2T} \exp\big(-(1 - 10^{-6})(q + p(u)) + 10^{-6}\big) \le e^{-100T}.$$

Integrate against Haar probability. This proves the claimed estimate. □

This proves `dsCoeff_gaussian_shape` of `New/DiscreteSeries/GaussianShape.lean` using the Flensted–Jensen identification and ordinary compact representation theory. The absolute error is independent of the size of the Gaussian; this is what is needed in the restriction integral. The estimates concern the varying rank and weights, not merely the limit of a fixed representation at the identity.

16.4 Determinant decay and a small-ball bound

We give a further direct quantitative argument for the bounding functions used in the restriction calculation. Let γ be an integer such that $0 \le \gamma \le h_{\mathrm{gap}}$, where h_{gap} is the minimum allowed separation of λ_i, λ_j for indices of opposite sign. Thus any interval of the shifted weights μ of length γ, half open at one end, contains only indices of one sign. Let a_z denote the same-sign block part of a matrix z.

Lemma 16.5 (Decay from complementary and Fischer minors)**.** *For all $g \in G_J$,*

$$|\psi_\lambda(g)|^2 |\det a_g|^\gamma \le 1, \qquad |\psi_\lambda(g)| \le (2/\|g\|_{\mathrm{op}})^\gamma, \qquad |\psi_\lambda(g)|^4 \le (1 + \|g - a_g\|_2^2)^{-\gamma}.$$

Proof. First let $z \in G_J$ be positive Hermitian. Then $\det z = 1$ and $z^{-1} = JzJ$, so Jacobi's identity gives $\det z_S = \det z_{S^c}$ for every principal minor. Define $S = \{i : \mu_i + \gamma \ge t\}$ and $R = \{i : \mu_i \ge t\}$ for an integral threshold t. The difference $S \setminus R$ lies in one sign block. In that block Fischer's inequality gives

$$\det(a_z)_S \le \det(a_z)_R \det z_{S \setminus R}.$$

By the complementary-minor identity and then Fischer's inequality on the complementary set,

$$\det z_{S \setminus R} = \det z_{R \cup S^c} \le \det z_R \det z_{S^c} = \det z_R \det z_S.$$

Multiply these inequalities over a range of thresholds containing all weights, and include the determinant corrections. Since $\det z = 1$ and $z^{-1} = JzJ$, the result is

$$(\det a_z)^\gamma = \frac{m_{\mu+\gamma}(a_z)}{m_\mu(a_z)} \le m_\mu(z) m_{\mu+\gamma}(z^{-1}) = m_{-\mu}(z)^2.$$

For arbitrary g, the ordinary polar decomposition is $g = uz$ with $u \in K$ and z positive Hermitian in G_J; indeed $J(g^*g)J = (g^*g)^{-1}$ implies the same identity for its positive square root. Lemma 16.1 gives the splitting (u, z). Apply the preceding bound to every conjugate v^*zv with $v \in K$ (the block determinant is unchanged). Cauchy–Schwarz and Haar invariance show

$$|\psi_\lambda(g)| \le (\det a_z)^{-\gamma/2} Z_\Lambda^{-1} \int_K |m_\Lambda(uv)|\, |m_\Lambda(v)|\, dv \le |\det a_g|^{-\gamma/2}.$$

For the other two bounds write $g = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$ by signs. The two unitary identities yield

$$A^*A = 1 + C^*C, \quad AA^* = 1 + BB^*, \quad D^*D = 1 + B^*B, \quad DD^* = 1 + CC^*.$$

In particular the nonzero singular values of B, C agree. If one sign block is empty, the asserted inequalities follow immediately from $|\psi_\lambda| \le 1$ and $\|g\|_{\mathrm{op}} = 1$. Otherwise if $t = \|A\|_{\mathrm{op}} = \|D\|_{\mathrm{op}}$, then $\|B\|_{\mathrm{op}}, \|C\|_{\mathrm{op}} \le t$. All singular values of A, D are at least one, hence $|\det a_g| \ge t^2 \ge \|g\|_{\mathrm{op}}^2/4$. Furthermore

$$|\det a_g| = \det(1 + BB^*) \ge 1 + \|B\|_2^2, \quad |\det a_g|^2 \ge 1 + 2\|B\|_2^2 = 1 + \|g - a_g\|_2^2.$$

Substitution proves the last two bounds. □

This derives the bounds `norm_dsCoeff_sq_mul_det_pow_le_one`, `norm_dsCoeff_le_pow` of `Lib/DiscreteSeries/FlenstedJensen/PsiDecay.lean` and `norm_dsCoeff_pow_four_le` of `New/DiscreteSeries/TwoFactorSetup.lean`. Fischer's inequality follows, for example, by taking the Schur complement of one block of a positive definite matrix; no spectral asymptotic enters this step.

Lemma 16.6 (Very small balls)**.** *For $|\ell_i| \le L$ and $d + L + 1 \le \beta$, if $\|g - 1\|_{\mathrm{op}} \le (200\beta^5)^{-1}$, then $|\psi_\lambda(g)| \ge 1/2$.*

Proof. Set $\delta = (200\beta^5)^{-1}$. The matrix g is accretive, so use the splitting $(1, g)$ in Lemma 16.1. All principal compressed logarithms of v^*gv have absolute trace at most $2d\delta$, because the operator norm of their difference from the identity is at most δ. The number of nonconstant factors in the expression for $m_{-\mu}$ including the determinant power is at most $3\sum_i |\mu_i| \le 3d(L + d + 1)$. Thus $|\log m_{-\mu}(v^*gv)| \le 6\beta^3\delta \le 3/100$. Taking the reciprocal, the deviation from one is bounded by $e^{3/100} - 1 < 1/2$. Integrate with the positive normalized compact highest-type density to obtain $|\psi_\lambda(g) - 1| < 1/2$. □

The source proves this last assertion via the infinitesimal bound `norm_dsCoeff_sub_le`, which is also useful for the differentiated Schur inequalities. Only the small-ball assertion `half_le_norm_dsCoeff` of `Lib/DiscreteSeries/FlenstedJensen/TailBounds.lean` is needed here. The argument just given does not require an infinitesimal representation estimate.

16.5 Removal of the common compact subgroup

We retain the hypotheses of the two-factor setup of Section 15.4. Put $g_i = j_i(s)$, with ranks d_1, d_2, sign matrices J_i, and integral weight matrices Λ_1, Λ_2. Write $g_{i,\mathrm{diag}}$ for the same-sign block part of g_i. In the first group, $\Lambda_1 = 0$ on the common positive subspace X_0 of this setup (not the full definite subspace of the filling datum), and $-3B \le \Lambda_{1,j} \le -B$ elsewhere. In the second group Λ_2 is nonnegative, is bounded by $3B$ on all rows having any opposite-sign direction, and vanishes on the corresponding common positive subspace. The pair has equal traces and imaginary additive phase

$$\varphi(s) = \mathrm{tr}(\Lambda_1(g_1 - 1)) - \mathrm{tr}(\Lambda_2(g_2 - 1)).$$

It contains the simultaneous copy of the unitary group on X_0. Recall

$$\theta(s) = \psi_{\lambda_1}(g_1)\overline{\psi_{\lambda_2}(g_2)}e^{-\varphi(s)}.$$

Lemma 16.7 (Simultaneous compact correction)**.** *Suppose $B \ge 1$ and $0 < \rho \le 1/100$. If*

$$\|g_1 - g_{1,\mathrm{diag}}\|_2^2 + \|g_2 - g_{2,\mathrm{diag}}\|_2^2 \le \rho^2$$

then there is t obtained from s by left multiplication by an element of the common compact group, with $\theta(t) = \theta(s)$ and

$$\|j_1(t) - 1\|_2, \ \|j_2(t) - 1\|_2 \le 10\rho.$$

Proof. Put $H_i = g_i - 1$. The signed row identity from Section 15.4 gives, for a diagonal nonnegative weight w,

$$\sum_{a,b} w_a |(H_i)_{ab}|^2 + 2\Re\,\mathrm{tr}(wH_i) = 2 \sum_{\epsilon_a \neq \epsilon_b} w_a |(g_i)_{ab}|^2.$$

Use $w = -\Lambda_1$ and $w = \Lambda_2$ in the two groups and add. The trace contributions cancel because $\Re\varphi = 0$. If P is projection onto the common positive subspace in the first group, $Q = 1 - P$, it follows that $\|QH_1\|_2^2 \le 6\rho^2$.

In blocks P, Q, write $g_1 = \begin{pmatrix} a & b_1 \\ c & d \end{pmatrix}$. The first block lies in a positive part of the sign form. Unitarity gives

$$a^*a + c^*J_Q c = 1, \qquad a^*b_1 + c^*J_Q d = 0.$$

Thus $\|a^*a - 1\|_2 \le \|c\|_2^2 \le 6\rho^2 < 1$, so a is invertible. Write $a = u|a|$ with u unitary on the common compact subspace. Since $|\sqrt{x} - 1| \le |x - 1|$ for $x \ge 0$, $\||a| - 1\|_2 \le 6\rho^2$. The second equation gives

$$\|b_1\|_2 \le (1 - 6\rho^2)^{-1/2}(1 + \sqrt{6}\,\rho)\|c\|_2 \le \tfrac{11}{10}\|c\|_2.$$

Choose the simultaneous compact element from the hypothesis which acts by u^{-1} on both copies of the common compact subspace and fixes their complements. For the corrected first matrix $\widetilde{g}_1$, the block a becomes $|a|$, the row Q is unchanged, and b_1 is multiplied by a unitary on the left. Consequently

$$\|\widetilde{g}_1 - 1\|_2^2 \le 36\rho^4 + (1 + (11/10)^2)6\rho^2 < 16\rho^2.$$

This compact multiplication leaves the opposite-sign norms of each factor unchanged. In the two signatures apply the signed row identity with all weights equal to one. Equality of the traces of the two corrected factors implies

$$\|\widetilde{g}_2 - 1\|_2^2 \le \|\widetilde{g}_1 - 1\|_2^2 + 2\|\widetilde{g}_2 - \widetilde{g}_{2,\mathrm{diag}}\|_2^2 \le 18\rho^2.$$

Finally the lowest compact type has weight zero on the whole common compact block. Its highest vector is fixed by the unitary group of that block: the restriction has highest weight all zero, so its orbit there is the trivial irreducible. Equivalently the product of weight minors, restricted to the compact block in question, is unchanged by this compact rotation (the weight is constant zero throughout it). The normalized coefficient of this vector is unchanged by left multiplication by the corresponding element in each group. Also the linear phase is zero on the common compact subgroup, and is additive on S. Thus θ is unchanged. This proves the bounds and invariance. □

This is `IsTwoFactor.exists_near_one` in `New/DiscreteSeries/TwoFactorSetup.lean`. The polar correction here replaces a matrix inverse construction in the formal proof but gives the same error constant. In particular it uses the equal trace condition, not a separate estimate of the second group's compact block.

Proposition 16.8 (Cone estimate for the oscillatory integral)**.** *Let $D = (d_1 + d_2)^2$, $\beta \ge 10^{12}$, $d_1 + d_2 + L \le \beta$, $T = D \log \beta$, and suppose the weight separation in the two groups is at least $10^{30}T^3$ in the notation of the large-gap hypothesis. Write $G(s) = 2/\|j_1(s)\|_{\mathrm{op}}$. Then*

$$|\theta(s)| \le G(s)^a \quad (0 \le a \le 18D), \qquad 4|\theta(s)| \le 8\big(\Re\theta(s) + 16\beta^{-50D}G(s)^{8D}\big).$$

In the first inequality it suffices to take integer a; real exponents in this interval follow by interpolation with $a = 0$.

Proof. The first bound uses Lemma 16.5 for the first factor and the unit-vector coefficient bound for the second. For the second put $\rho = 10^{-13}/T$. If the sum of the squared opposite-sign norms is at most ρ^2, use Lemma 16.7 to make both corrected factors $10^{-12}/T$-close to the identity in Frobenius norm. Their exponential phases cancel exactly by the logarithmic phase identity (15.11); its proof used only additivity of φ. Theorem 16.4 with $T = D\log\beta$ gives θ as a product of two complex numbers whose differences from $a_1, a_2 \in [0,1]$ are at most $a_i/10 + \zeta$, where $\zeta = e^{-100T} = \beta^{-100D}$. (Conjugating the second factor leaves this estimate unchanged.) In particular

$$|\theta - a_1 a_2| \le 0.21 a_1 a_2 + 3\zeta.$$

Since $\Re\theta \ge a_1 a_2 - |\theta - a_1 a_2|$, it follows that

$$|\theta| - 2\Re\theta \le 3|\theta - a_1 a_2| - a_1 a_2 \le 9\zeta.$$

The operator norm of j_1 is unchanged by the compact correction. It is now at most $1 + 10^{-12}/T < 2$, so $G(s) > 1$, and the required inequality follows.

Otherwise the sum of the squared opposite-sign norms exceeds ρ^2. Multiplication of the fourth-power decay bounds gives

$$|\theta|^4 \le (1+\rho^2)^{-\gamma} \qquad \text{with } \gamma \ge 10^{30} T^3.$$

Indeed the product $(1 + \|g_1 - a_{g_1}\|_2^2)(1 + \|g_2 - a_{g_2}\|_2^2)$ is at least one plus the sum. Since $\log(1+\rho^2) \ge \rho^2/2$ for $\rho^2 \le 1$, $\frac{1}{4}\gamma\log(1+\rho^2) \ge 1250T$. Hence $|\theta| \le \zeta$. Use also the decay bound $|\theta| \le G^{16D}$ and take the geometric mean to get $|\theta| \le \beta^{-50D} G^{8D}$. Then $|\theta| - 2\Re\theta \le 3|\theta|$ proves the second inequality in this case. □

If the gap input is given with an additional integer threshold $500 d_i^3$, as in the source, the loss from taking an integer exponent in Lemma 16.5 is harmless (the argument can take the specified large integer gap). The constants here prove the pointwise inequalities (15.13)–(15.14) needed by the integration argument of Proposition 15.6. The small-ball assertion in that proof is supplied by Lemma 16.6. Thus the specialized oscillatory part of `IsTwoFactor.laplace_method` of `New/DiscreteSeries/RestrictionLemma.lean` has now been reconstructed. It is conditional on the representation-theoretic coefficient identification, rather than on an additional asymptotic estimate for discrete series.

16.6 A phase-cancelling realization of the Bessel group

We now give the geometric realization required for the Bessel peak in Section 15.5. The standard complex Hermitian form is

$$J_1 = \begin{pmatrix} 0 & 0 & 1_h \\ 0 & T & 0 \\ 1_h & 0 & 0 \end{pmatrix}, \qquad T = \operatorname{diag}(1_k, -1),$$

with standard frame $F = (e_0, \dots, e_{h-1}, z)$, isotropic duals $e_0^*, \dots, e_{h-1}^*$, and definite part with basis $x_1, \dots, x_k$. The Bessel group $S_B = UH$ consists of transformations acting by an upper unitriangular matrix on F, where $H = U(k)$ fixes the hyperbolic planes and z, and U acts trivially on the middle quotient. On S_B put

$$\sigma(s) = \sum_{i=0}^{h-1} s_{f_i, f_{i+1}}, \qquad f_i = e_i \ (i < h), \quad f_h = z.$$

Since the frame matrix is unitriangular, σ is additive on products. The real character in the source is $\mathrm{wave}(\sigma) = \exp(\sigma - \overline{\sigma})$. We first prove a useful norm estimate for its realization.

Lemma 16.9 (Controlling a representative of a Bessel frame)**.** *Suppose vectors V_j for $0 \le j \le 2h$ satisfy $\|V_j\| \le M$, $M \ge 1$, and*

$$V_a = x^{-1} f_a \ (a \le h), \qquad \langle V_i, J_1 V_j \rangle = \begin{cases} 0 & i+j < 2h, \\ -1 & i+j = 2h. \end{cases}$$

Then $\mathrm{bs}(x) = \inf_{s \in S_B}(\|sx\|_{\mathrm{op}} + 1) \le 8(h+1)M$.

Proof. Apply unnormalized ordinary Euclidean Gram–Schmidt to the ordered frame $(V_0, \dots, V_h)$, writing F_a for its orthogonalized columns. They are obtained by an upper unitriangular change of columns. Such a change lifts to an element of U: if the change on the isotropic block is A and the extra column of the last frame vector is v, the lift on (e, z, e^*) is

$$\begin{pmatrix} A & v & \frac{1}{2} v v^* (A^*)^{-1} \\ 0 & 1 & v^* (A^*)^{-1} \\ 0 & 0 & (A^*)^{-1} \end{pmatrix},$$

and it fixes the definite part. Direct multiplication against the signed form verifies this formula. Hence some Bessel translate has inverse frame precisely (F_a). Each $\|F_a\| \le M$. The anti-diagonal hypothesis gives $\langle F_a, J_1 V_{2h-a} \rangle = -1$, because the earlier columns have zero pairings with V_{2h-a}. Thus $\|F_a\| \ge M^{-1}$.

For $a < h$ put $D_a = J_1 F_a / \|F_a\|^2$. Orthogonality of the Euclidean Gram–Schmidt frame gives $\langle F_b, J_1 D_a \rangle = \delta_{ab}$ (including $b = h$), and $\langle D_a, J_1 D_b \rangle = \langle F_a, J_1 F_b \rangle / (\|F_a\|^2 \|F_b\|^2) = 0$. The vectors D_a are bounded by M and form an isotropic frame dual to $(F_0, \dots, F_{h-1})$. They can be made the inverse images of the standard dual frame by an additional translation in U^0, the subgroup fixing the whole original frame. To see this, in coordinates of the already adjusted element any such dual array has columns $(C; A; 1_h)$ in the $(e, X_{\mathrm{definite}}, e^*)$ decomposition, with zero z-part. Its isotropy is $C + C^* + A^* A = 0$. Writing $C = B - \frac{1}{2} A^* A$ with B skew-Hermitian gives the usual unipotent element fixing the frame,

$$\begin{pmatrix} 1 & -A^* & B - \frac{1}{2} A^* A \\ 0 & 1 & A \\ 0 & 0 & 1 \end{pmatrix}.$$

It has exactly these dual columns. It fixes z and is unitary for the hyperbolic and definite form by the displayed identity.

Let M_0 be the inverse of the resulting Bessel translate, with frame and dual columns denoted by F_a, D_a, and negative column F_h. Using $M_0 J_1 M_0^* = J_1$ and subtracting it from $M_0 M_0^*$ gives the exact identity

$$M_0 M_0^* = J_1 + \sum_{a<h} (F_a - D_a)(F_a - D_a)^* + 2 F_h F_h^*.$$

Therefore $\|M_0\|_{\mathrm{op}}^2 \le 1 + (4h+2)M^2$. Also the inverse has the same operator norm since J_1 is Hermitian unitary. In particular the norm plus one is bounded by $8(h+1)M$, as required. □

This proves `besselSize_le_of_dual` in `New/RealPlace/FrameNorm.lean` using the triangular-coordinate description of U^0.

For the next construction the ambient weight on the r-dimensional sign space is zero on a common positive subspace X of dimension $k-1$ and has strictly decreasing exceptional entries $\Lambda_i \in [-3B, -B]$, $0 \le i \le 2h+1$, with signs $(-1)^i$. The compact weight on the k-dimensional definite group is zero on X and equals $a \ge 0$ on its remaining distinguished coordinate. The source imposes $B \ge 1$, $3B + a \le 2b^2$, and $b \ge 8(h+1)$.

Proposition 16.10 (The interpolation realization)**.** *There is $x \in G_{J_1}$ with $\mathrm{bs}(x) \le b^{5(h+1)}$ such that the Bessel group graph $s = (g_s, h_s) \in S_B \times H$, with $g_s \in UH$ and compact factor h_s, is a two-factor subgroup on conjugating the first factor by x and an orthonormal sign frame. The phase is exactly*

$$\varphi(s) = \sigma(g_s) - \overline{\sigma(g_s)}.$$

Both factors contain the indicated common unitary subgroup diagonally, and their traces agree after subtracting their respective identities.

Proof. Put $t_i = \Lambda_i - a$; these are strictly decreasing negative integers. Apply the Lagrange interpolation Lemma 15.7 to the descending sequence $(0, t_0, \dots, t_{2h+1})$ and write its positive weights as $c_0, c_1, \dots, c_{2h+2}$. In the exceptional coordinates form

$$(C_0)_i = c_0/c_{i+1}, \qquad (C_j)_i = t_i^j/c_{i+1} \quad (j \ge 1).$$

The signed interpolation identity gives $H(C_0, C_0) = 1$, $H(C_0, C_j) = 0$ for $1 \le j \le 2h+1$, and $H(C_i, C_j) = 0$ for $i, j > 0$, $i + j < 2h + 2$, equal to -1 when $i + j = 2h + 2$. Here H denotes the signed form on the exceptional coordinates. Consequently the common positive X, followed by $C_0, C_1, \dots, C_{h+1}$, forms the same partially degenerate Hermitian frame as the standard X, its distinguished last definite vector x_0, and $(e_0, \dots, e_{h-1}, z)$. Witt extension for nonsingular Hermitian spaces over $\mathbb{C}/\mathbb{R}$ extends this isometry of subspaces to an isometry of the whole space. Equivalently choose duals for the isotropic part and subtract half their Gram matrix to make them isotropic, then complete the two orthogonal complements by orthonormal bases. Let p be an orthonormal sign frame for J_1. Choose x so that xp maps the above columns to the standard ones.

The columns pC_j, $1 \le j \le 2h+1$, are the dual array in Lemma 16.9, with $V_i = pC_{i+1}$. Since each $c_{i+1} \ge 1$ and $|t_i| \le 2b^2 \le b^{5/2}$,

$$\max_{1 \le j \le 2h+1} \|pC_j\| \le \sqrt{2h+2}\, b^{(5/2)(2h+1)}.$$

Lemma 16.9 gives $\mathrm{bs}(x) \le b^{5(h+1)}$, since $8(h+1)\sqrt{2h+2} \le b^{3/2}$.

We verify the phase rather than infer it from Witt extension. Let

$$D = xp\,\mathrm{diag}(\Lambda_i)(xp)^{-1}.$$

It is self-adjoint for J_1. In the above columns, multiplication by a weight is multiplication by $t_i + a$, hence

$$Df_j = af_j + f_{j+1} \quad (0 \le j < h), \qquad Dx_0 = ax_0 + c_0 f_0, \qquad D|_X = 0.$$

Order the blocks of G_{J_1} as the isotropic lines $e_0, \dots, e_{h-1}$, the middle definite block together with z, and the dual lines in opposite order. Matrices of U are strictly block upper triangular off the identity. To contribute to $\mathrm{tr}\, D(g_s - h_s)$ an entry of D must occur strictly below these blocks. For a column e_i the only such entry is the coefficient 1 of f_{i+1}. For columns in the middle block, self-adjointness gives a nonzero dual coefficient only for the column z, with coefficient -1 at e_{h-1}^*, since $H(De_{h-1}, z) = H(z, z) = -1$. (For a definite column its pairing with every De_i is zero.) For the dual columns, the remaining lower entries are prescribed by self-adjointness of the isotropic block: they are the coefficients 1 at (e_i^*, e_{i+1}^*). The dual blocks of a unitriangular unitary matrix are inverse adjoints. Thus the superdiagonal dual entries of g_s are minus the conjugates of the superdiagonal isotropic entries. For the last root the sign of the negative line makes the z, e_{h-1}^* entry the conjugate of e_{h-1}, z, as is seen by imposing

$H(g_s z, g_s e^*_{h-1}) = 0$. The middle rotation h_s acts trivially on the indicated frame and dual entries. Summing all contributions proves

$$\operatorname{tr} D(g_s - h_s) = \sigma(g_s) - \overline{\sigma(g_s)}.$$

On the middle definite block D is exactly the compact weight $\operatorname{diag}(0_X, a)$; the other diagonal entries of $h_s - 1$ vanish. Therefore $\operatorname{tr} D(h_s - 1)$ is the second-factor weight trace. This gives the required phase. Its real part is zero, its additivity follows from upper unitriangularity of the frame matrices, and the trace identity follows from the diagonal blocks of g_s. Finally a compact transformation of the common X is identified by xp with the standard unitary subgroup acting on those definite coordinates. Thus it is present in the two factors simultaneously. □

The source statement is `IsSteep.exists_isTwoFactor_besselPairHom` in `New/DiscreteSeries/BesselExistence.lean`. Its hypotheses on the actual parameter pattern are supplied by `Data.exists_isSteep` in `New/Exchange/SteepParameters.lean`; establishing the inequalities between the actual weight parameters is separate from this interpolation calculation. The first projection of the Bessel graph after conjugation is proper because the subgroup UH is closed and the compact factor of the graph is a continuous function of the middle block. The second component has operator norm one. Accordingly the quantitative Laplace argument above applies once its gap and size hypotheses are verified.

Combined with the quotient integral construction (15.15)–(15.16) and the peak-to-cut argument of Section 15.6, this supplies the principal analytic mechanism of `Data.exists_besselSystem` (`New/Filling/FillingData.lean`). The representation-theoretic input of a unit highest vector of the compact Schur type, identification of the lowest-type coefficient and its formal degree, and the differentiated convolution bound are still used as explicitly indicated. For the remaining Fourier-floor portion of this declaration, the arguments just proved do not replace the local gauge, bandlimit and open-cell estimates in `New/RealPlace/`. We return to that distinction at the end of the chapter.

16.7 Marked integral lattice and the exact finite level indices

Let E be the CM field, E^+ its totally real subfield, $\mathcal{O} = \mathcal{O}_E$, $[E : \mathbb{Q}] = 2f$. The Hermitian space is

$$V = L \oplus M \oplus L^\vee, \qquad L = \langle e_0, \dots, e_{h-1} \rangle_E, \qquad M = X \oplus Ez, \quad h|_X = 1_k, \quad h(z,z) = -N,$$

where $L^\vee$ is isotropic dual to L and X is positive definite at every infinite place. Write $\mathcal{O}_f = \widehat{\mathcal{O}}$, and normalize the global additive measure using ordinary Lebesgue measure on $\mathbb{C}^f$; by the covolume-one convention on the adeles, $\operatorname{vol}(\mathcal{O}_f) = 2^f |D_E|^{-1/2}$. The constructions use the generator ϖ of the different, with $\bar{\varpi} = -\varpi$, and the primitive trace-one element of $\mathcal{O}$ available because E/E^+ is everywhere unramified at its finite places in the stipulated `StablyUnram` setting. The same setting supplies a unit $i_0 \in \mathcal{O}$ with $i_0^2 = -1$.

We distinguish the two halves of the split level prime N by writing $\eta\bar{\eta} = N$ for finite ideles, $\eta, \bar{\eta} \in \mathcal{O}_f$; at each of the pairs over N precisely one of them is a uniformizer and the other a unit. At primes away from N both are units. The marked lattice $\mathcal{L}$ in $V(\mathbb{A}_f)$ has the standard hyperbolic basis and, in the middle block, the basis

$$x_i \ (i \neq i_0'), \qquad v = x_{i_0'}/\eta + z/N, \qquad z.$$

Here i_0' is merely the chosen index in the positive definite basis, unrelated to the imaginary unit i_0. The Gram matrix of (v, z) is

$$\begin{pmatrix} 0 & -1 \\ -1 & -N \end{pmatrix};$$

indeed $h(v,v) = 1/N - N/N^2 = 0$ and $h(v,z) = -1$. Hence $\mathcal{L}$ is unimodular. Let K be the integral stabilizer of this lattice and $H = U(X)$ the subgroup of $U(V)$ fixing the hyperbolic pairs and z. Write K_X for the stabilizer of the standard finite-adelic lattice in X.

Lemma 16.11 (The marked-lattice index)**.** *The subgroup $K \cap H(\mathbb{A}_f)$ is contained in K_X, and its index there is $(N^k - 1)^f$.*

Proof. At a finite place write $\mathcal{L}_X = \mathcal{L} \cap X$ and $\mathcal{L}^X$ for the projection of $\mathcal{L}$ onto X. At the marked split places

$$\mathcal{L}_X = \sum_{i \neq i'_0} \mathcal{O}_w x_i + \bar{\eta}\mathcal{O}_w x_{i'_0}, \qquad \mathcal{L}^X = \sum_{i \neq i'_0} \mathcal{O}_w x_i + \eta^{-1}\mathcal{O}_w x_{i'_0}.$$

Indeed a combination $tv + sz$ lies in X if and only if $s = -t/N$, i.e. $t \in N\mathcal{O}_w$. Away from N both lattices are standard, while at each split place over N one or the other is standard. Any transformation of X preserving $\mathcal{L}$ preserves both lattices; it therefore lies in K_X.

For $x \in K_X$ the condition that its action preserve $\mathcal{L}$ is exactly the finite marked-vector condition. To check it apply the action to the unmarked vectors, v and z. At a place where η is the uniformizer, $\mathcal{L}_X$ is standard and the only new condition is

$$x(v) - v = (xx_{i'_0} - x_{i'_0})/\eta \in \mathcal{L}_X,$$

so the distinguished vector is fixed modulo η. At its conjugate, η is a unit and the other factor $\bar{\eta}$ is a uniformizer; preservation means that the distinguished coordinate row of x agrees with the identity row modulo $\bar{\eta}$. These are equivalent under the unitary inverse-adjoint relation between the two split local factors. The same test on the inverse gives equality, not just containment; the vector condition is preserved by inverse.

The split local unitary stabilizer is identified with $\mathrm{GL}_k(\mathcal{O}_w)$ at the selected place of each pair, and reduction onto $\mathrm{GL}_k(\mathbb{F}_N)$ is surjective: lift the entries of an invertible matrix, whose determinant remains a unit, and fill its conjugate local factor by inverse adjoint. The index of the stabilizer of a nonzero vector in $\mathrm{GL}_k(\mathbb{F}_N)$ is the cardinality $N^k - 1$ of the orbit of that vector. Multiply over the f pairs of split places over N. □

This establishes `Data.relIndex_latStabilizer` of `New/Filling/FillingAdmissible.lean`. It is important that the condition fixes a *vector* rather than merely a line; the latter gives a different index by the factor $N - 1$.

The two auxiliary rational primes q_0, q_1 split in E and are different from N; each comes with a choice of f places, one from every conjugate pair. Reduction of K at these places gives

$$\mathfrak{G} = \prod_{j=0}^{1} \prod_{w \mid q_j \text{ chosen}} \mathrm{GL}_r(\mathbb{F}_{q_j}).$$

This map is surjective by the same split lifting as in Lemma 16.11; its kernel is the level $K(g)$, where $g = q_0 q_1$. The marking at N does not interact with this reduction, since local factors can be set to the identity at the other finite places. The column level is $K_c = K(g) \cap H(\mathbb{A}_f)$, the preimage of the congruence level under the inclusion of H in the marked lattice basis.

Let $\bar{U}, \bar{H}, \bar{S}_B$ denote the images under reduction of U, H, and $S_B = UH$, intersected with K. At the split level primes, the image of U is the block unitriangular group of block sizes $1, \dots, 1, k+1, 1, \dots, 1$ (h blocks of size one at each end). The conjugate factor is filled by inverse adjoint, so there are no extra equations on the selected factor. Its number of free residue-field entries is

$$\frac{(k+2h+1)^2 - (k+1)^2 - 2h}{2} = h(2k+2h+1).$$

The image of H is the ordinary k-dimensional linear group inside the middle block. Thus $\bar{S}_B = \bar{U} \rtimes \bar{H}$, and

$$|\bar{U}| = g^{fh(2k+2h+1)}, \qquad |\bar{S}_B| = |\bar{H}|\, g^{fh(2k+2h+1)}. \tag{16.6}$$

For the count, lifts of unitriangular elements are obtained by lifting their free entries and the inverse-adjoint factor; outside the marked split places take the identity. Conversely a product uh is integral only if both factors are integral: its middle diagonal block determines h, and multiplying by h^{-1} recovers u. The same argument applies to congruence levels. This is the content of `card_reducedImage_besselS` in `New/Filling/GlobalToLocal.lean`; it also shows directly why the exponent of g depends on the full hyperbolic structure and not just the dimension of the upper triangular frame group.

In the positive definite group write

$$\mathcal{C} = H(E^+)\backslash H(\mathbb{A}_f)/K_c, \qquad \mathfrak{m}_k = \sum_{d \in H(E^+)\backslash H(\mathbb{A}_f)/K_X} |\operatorname{Aut}(d)|^{-1}.$$

The usual class mass calculation is recalled in Chapter 5. Since the deeper level K_c has trivial rational stabilizers, the number of classes over d is $[K_X : K_c]/|\operatorname{Aut}(d)|$. The stabilizer acts freely on the finite quotient: a fixed coset would give a conjugate stabilizer at level K_c. By Lemma 16.11 and surjectivity of reduction on the middle group,

$$|\mathcal{C}| = |\bar{H}|(N^k - 1)^f \mathfrak{m}_k. \tag{16.7}$$

This gives `Data.card_classes_filling`. Together with the cardinality $D_I = \prod_v |I_v|$ of the local Bessel index tuple, the column dimension is $D_c = |\mathcal{C}| D_I$. Notice the cancellation between the $|\bar{H}|$ factors in the class count and the order of $\bar{S}_B$ in the energy bounds below.

16.8 Haar coordinates and the arithmetic unipotent covolume

We give the exact unipotent covolume calculation, including its dependence on the marking. Write $T = \operatorname{diag}(1_k, -N)$ for the middle form before archimedean rescaling. Relative to the ordered blocks $(L, M, L^\vee)$, an element of U is determined by three upper blocks as follows:

$$u = \begin{pmatrix} A & B & C \\ 0 & 1 & D \\ 0 & 0 & (A^*)^{-1} \end{pmatrix}, \qquad D = -T^{-1}B^*(A^*)^{-1},$$

$$M_0 + M_0^* = -A^{-1}BT^{-1}B^*(A^*)^{-1}, \qquad M_0 = A^{-1}C,$$

with A upper unitriangular. These formulas follow by multiplying the Gram matrix against the three block columns. At an infinite place we will use the normalized negative entry -1. The free complex entries may be taken to be, in row e_a:

$$e_b,\ e_b^*\ (b > a), \quad \text{all } k+1 \text{ middle columns}, \qquad 0 \le a < h.$$

In addition take the real number $\Im C_{aa}$ for every a. To see that these are actual coordinates, fix A, B and the upper triangular entries C_{ab} for $a < b$. Solve $M_0 = A^{-1}C$ column by column, starting with $b = h-1$ and descending: the lower entries $M_{0,jb}$, $j > b$, are determined by the already computed columns and the displayed Hermitian relation. The real part of the diagonal of M_0 is fixed by this relation, its imaginary part is determined by $\Im C_{bb}$ after subtracting known later entries, and then the entries $M_{0,ab}$ for $a < b$ are solved in descending order a. All these operations are affine triangular with diagonal coefficient one in the free coordinates.

We normalize Haar measure on U_∞ to be Lebesgue in these free entries. This agrees with the source's `Shape.besselHaar`, which uses Euclidean measure on the unitriangular frame and on the skew-matrix and definite variables (`Lib/DiscreteSeries/Bessel/UnipotentCoords.lean`, `RealMeasures.lean`). To verify the agreement, parametrize M_0 by its skew-Hermitian part plus the known quadratic term in A, B. The change of its skew coordinates into the free entries is triangular, of determinant of modulus one, by the preceding calculation. Left multiplication by a fixed group element has triangular coordinate Jacobian one: it multiplies A by an upper unitriangular matrix, replaces B by that matrix times B plus a fixed vector, and replaces M_0 by itself plus a function of the other coordinates, as follows from the block multiplication law. Hence the measure is left Haar and also right Haar, since this group is unipotent. Using the compact subgroup H's Haar probability gives the Bessel graph measure, as in Section 16.6.

For a class c choose a representative $h_c \in H(\mathbb{A}_f)$ and use the lattice $\mathcal{L}_c = h_c\mathcal{L}$, with the basis matrix $B_c = h_c B_{\mathcal{L}}$ in rational coordinates. The matrix B_c is the identity on the two isotropic blocks and invertible on the middle block. Its generalized index is

$$[\det B_c] = N^{-f}.$$

Here $[a] = \prod_{w \text{ finite}} |a_w|_w^{-1}$, and $|\cdot|_w$ is the normalized local module. Indeed $[\eta] = N^f$, the determinant of the middle basis change is η^{-1}, and the determinant of h_c has norm one. The Gram matrix of $\mathcal{L}_c$ is the same unimodular Gram matrix as that of $\mathcal{L}$. Let

$$\Gamma_c = U(E^+) \cap h_c K(g) h_c^{-1}.$$

Lemma 16.12 (Integral row translations)**.** *Let l be a rational linear row form supported on the free complex entries of row a. Suppose lB_c is integral at all finite places. Given $s_0 \in \mathcal{O}$ with $\bar{s}_0 = -s_0$, there is $\gamma \in \Gamma_c$ whose free off-diagonal coordinates are zero except in row a, where they equal gl, and when $l = 0$ its only free coordinate is the skew diagonal entry gs_0 in that row.*

Proof. The Hermitian form identifies a row form with a vector: choose v with $h(v, \cdot) = l$. The condition on l says exactly that the finite image of v lies in the integral dual of $\mathcal{L}_c$; since $\mathcal{L}_c$ is unimodular, that dual is $\mathcal{L}_c$. Put $x = e_a$; it also lies in $\mathcal{L}_c$, is isotropic, and $h(v, x) = 0$. For an integral trace-one element τ ($\tau + \bar{\tau} = 1$) use the Eichler transformation

$$u = 1 + g(xh(v, \cdot) - vh(x, \cdot)) + (gs_0 - g^2\tau h(v, v))xh(x, \cdot).$$

It is unitary: the square of the first-order skew part is a multiple of $xh(x, \cdot)$; expanding u^*Ju reduces the quadratic identity to $\tau + \bar{\tau} = 1$ and $s_0 + \bar{s}_0 = 0$. It is upper unipotent relative to the flag because l vanishes on each e_b, $b \le a$, and the first-order non-row term has only the column e_a^*. Equivalently the matrix and its inverse are block upper triangular with diagonal blocks the identity; the possible nonfree entries in row $b > a$ are confined to the column e_a^*. Its difference from identity sends the class lattice into $g\mathcal{L}_c$, by integrality of x, v, their pairings, and τ, s_0. Thus it lies in Γ_c. Inspecting the entries of the displayed formula gives the assertion. □

Here trace one cannot always be replaced by $\tau = 1/2$: the latter need not preserve an integral lattice at residue characteristic two. Lemma 16.12 proves the specialized lifting statement `exists_eichler_classUnip` in `New/Filling/BoxPartition.lean`.

For each row set $d_a = k + 1 + 2(h - a - 1)$, the number of its free complex entries. The rational row lattice of those entries is

$$\mathcal{R}_a = \{l : \operatorname{supp} l \subseteq \text{free positions of row } a, lB_c \in \mathcal{O}_f^r\}.$$

The matrix on these positions is the restriction of B_c to them. It contains the entire middle block and is otherwise identity. Therefore its determinant has generalized index N^{-f}. In the ordinary mixed embedding $(\Phi_v)_{v|\infty}$ its covolume is

$$\frac{N^f}{\mathrm{vol}(\mathcal{O}_f)^{d_a}}.$$

For this formula, if D is any finite adelic invertible matrix, the lattice of rational row vectors x for which xD is integral is commensurable with $\mathcal{O}^d$. The ratio of covolumes is the inverse of the generalized index of $\det D$, by taking Smith forms locally (or multiplying by an integer to compare finite-index sublattices). This argument does not choose a principal generator for the resulting global lattice; it also applies when h_c has nontrivial class.

The rational skew lattice of a diagonal entry is $\{s : \bar{s} = -s,\ s \in \mathcal{O}\} = i_0\mathcal{O}_{E^+}$. The equality follows by dividing by the imaginary unit i_0, which is integral and a unit. The covolume of its imaginary parts at the f complex places is the covolume of the totally real integer lattice, namely $|D_{E^+}|^{1/2} = |D_E|^{1/4}$, since the stipulated CM field is unramified over E^+ at its finite places. This verifies the skew-lattice formula of `Filling.exists_fundamental_domain_im_skew`, `Lib/General/NumberTheory/FundamentalDomains.lean`.

We next justify multiplying these covolumes. Grade the free-coordinate groups in descending row order, using first the complex coordinates in a row, then its skew diagonal coordinate. If all earlier coordinates of $\gamma \in \Gamma_c$ vanish, right multiplication $u \mapsto u\gamma$ adds the free coordinates of γ at the current grade, without changing earlier ones. Here is the block verification: for rows $b > a$ already zero in their free coordinates, the corresponding rows of A and B in γ are identity and zero, respectively. The submatrix of C on two indices $b, a' > a$ is zero by the relation on $M_0 + M_0^*$, since the upper triangular and imaginary diagonal entries have already been set to zero. The D entries in the corresponding dual columns also vanish, and the corresponding columns of $(A^*)^{-1}$ are identity columns. Thus the matrix product cross term $(u-1)(\gamma-1)$ vanishes at all the free entries of row a. At the next grade, when row a's complex coordinates also vanish in γ, the same argument makes the cross term vanish at its diagonal skew coordinate. The arithmetic integrality condition for that row is precisely $lB_c \in \mathcal{O}_f^r$ after division by g; restriction to the free coordinates preserves integrality because the full middle block is included and the other blocks are identity. For the skew grade, when the complex entries are zero the diagonal entry is actually skew, by the same Hermitian relation. Lemma 16.12 realizes every permitted shift with all earlier coordinates zero.

Choose a half-open fundamental parallelepiped for each row lattice and for each real skew lattice. Successively in the above order translate to those parallelepipeds by an element of Γ_c. Each step exists by Lemma 16.12, and does not alter previous steps. For uniqueness, if two resulting elements are related by a group element of Γ_c, consider its earliest nonzero coordinate: the two coordinates in that parallelepiped would differ by a nonzero lattice shift, impossible. Injectivity of the free entries gives uniqueness. Thus the product of the parallelepipeds, in these graded coordinates, is a measurable strict fundamental domain. Inversion preserves Haar, so its volume is also the left covolume. This is the graded-box argument of `Filling.existsUnique_mul_mem_box` in `Lib/General/NumberTheory/FundamentalDomains.lean`.

Finally set

$$\kappa_v = \frac{g\,\Phi_v(\varpi)}{2\pi i} \in \mathbb{R}^\times, \qquad l_{va} = \frac{\kappa_v^{h-a}}{\sqrt{N}}, \quad s_{va} = l_{va}^2 \quad (0 \le a \le h).$$

The infinite-place change of basis scales e_a by l_{va}, e_a^* by l_{va}^{-1}, definite coordinates by one, and z by $l_{vh} = 1/\sqrt{N}$. It carries the middle negative norm to -1. In a free complex entry at row a, the rescaling multiplies by the column scale divided by l_{va}. The pairs e_b, e_b^* for $b > a$

cancel in the product of all numerator scales; the only remaining numerator is l_{vh} in column z. A complex scaling changes Lebesgue measure by its squared absolute value. The skew diagonal changes Lebesgue measure by l_{va}^{-2}. Including the level factor g and the preceding lattice covolumes, the volume contribution of row a is

$$\Big(\prod_v l_{vh}^2\, l_{va}^{-2(d_a+1)}\Big)\frac{g^{2fd_a}N^f}{\mathrm{vol}(\mathcal{O}_f)^{d_a}}g^f|D_E|^{1/4}.$$

As $\prod_v l_{vh}^2 = N^{-f}$, the row factors N^f cancel. Put

$$d_0=\sum_{a=0}^{h-1} d_a = h(k+h), \qquad \mathcal{S}=\prod_v\prod_{a=0}^{h} s_{va}^{k+2(h-a)}.$$

Since $d_a+1 = k+2(h-a)$, $\prod_{v,a<h} s_{va}^{-(d_a+1)} = N^{-kf}/\mathcal{S}$. The complete covolume is therefore

$$\boxed{\nu=\frac{(g^{\,2d_0+h}/N^k)^f\,|D_E|^{h/4}}{\mathrm{vol}(\mathcal{O}_f)^{d_0}\mathcal{S}}.} \tag{16.8}$$

This is exactly `Filling.unipCovolume` in `Setup.lean`. Its independence of c follows from the determinant index of B_c, not from assuming the class lattices are globally isometric.

Every measurable nonnegative cutoff β_c satisfying $\sum_{\gamma\in\Gamma_c}\beta_c(\gamma y)=1$ then has integral ν over each left U_∞-orbit: integrate the identity over the fundamental domain just constructed and use left Haar invariance and Tonelli. That is the precise assertion `lintegral_classCutoff_eq_unipCovolume` in `New/Filling/UnipotentCovolume.lean`. It supplies the constant used in the unfolding and energy normalization of Chapter 14.

16.9 Cutoffs, approximation and the quotient character

The measure calculation did not need a particular cutoff. The counting argument does need one whose support is boxed in the frame variables when the Bessel-size cutoff is active. We explain the construction and distinguish its arithmetic size estimate.

First the rational image of every Γ_c under the archimedean embedding is discrete. For a bounded set of entries at all archimedean places, the condition of lying in $h_cK(g)h_c^{-1}$ confines the finite entries to a compact set. An entry belongs to a denominator lattice of E with bounded archimedean values; there are finitely many such entries by discreteness of this lattice. Thus the set of matrices under these bounds is finite. The same argument applies to any prescribed submatrix. In particular if Γ_c^0 is the kernel of the map Γ_c to the frame unitriangular group, it is a closed discrete subgroup of U_∞^0. It has a measurable nonnegative cutoff of sum one on its left orbits (use a countable local fundamental covering of the discrete proper action and a partition of unity, or a Borel fundamental domain). Write β^0 for such a cutoff.

There is a useful structural fact about bounded frame boxes. Fix $y\in G_\infty$. For a frame unitriangular matrix A let u be any lift in Γ_c, so the vectors $(u^{-1}y)^{-1}f_a = y^{-1}uf_a$ depend only on A. There are finitely many possible A for which all these vectors have bounded norms at all infinite places: multiplying by the fixed y_v, the entries of A are bounded and the finite positions are also bounded in the class lattice. Let $\mathcal{Q}_c=\Gamma_c/\Gamma_c^0$ and let $I_y(A)$ be the indicator that these frame vectors have norm at most

$$H(y,v)=b^{60h+92}\mathrm{bs}(y_v) \qquad \text{at each } v.$$

The source proves that at least one such A exists for every y. Assuming this numerical reduction assertion for a moment, a cutoff on the quotient is given by

$$\omega(y) = \frac{I_y(1)}{\sum_{A \in \mathcal{Q}_c} I_y(A)}. \tag{16.9}$$

Each summand is measurable, the denominator is finite and positive, and $\omega \le 1$. The identity $I_{uy}(A) = I_y(u^{-1}A)$ gives the required weighted partition by reindexing the quotient. Also I_y and hence ω are invariant by the left action of $U_\infty^0 H_\infty$, which fixes the frame and leaves Bessel size unchanged. If $\omega(y) \neq 0$ and $\mathrm{bs}(y_v) \le b^{6n+2}$, the vector bound is at most $b^{60h+92+6n+2} \le b^{2n^2}$ because $h \ge 10^5$. Multiplying by β^0 gives the cutoff β_c of sum one used in `Filling.classCutoff`. The box and the discrete counting are therefore compatible with the covolume formula (16.8).

We now prove the required uniformly bounded representative. The following elementary weak transference estimate is convenient.

Lemma 16.13 (Covering-radius estimate)**.** *For a full lattice Λ in a Euclidean space of real dimension m, if its Euclidean dual has shortest nonzero vector of length at least $\rho > 0$, then every vector of that space is within distance m^2/ρ of Λ.*

Proof. Minkowski's first theorem in Λ and its dual gives

$$\lambda_1(\Lambda)\lambda_1(\Lambda^*) \le 4/\operatorname{vol}(B_1)^{2/m} \le m,$$

since the Euclidean unit ball contains the cube of side $2/\sqrt{m}$. Choose a shortest vector v of Λ, which is primitive, and project orthogonally along its span. The projected lattice has rank $m-1$, since a primitive vector extends to a lattice basis; its dual is $\Lambda^* \cap v^\perp$ and has shortest vector of length at least ρ. By induction on dimension, approximate the projection of the given vector by a projected lattice element to within $(m-1)^2/\rho$. Choose a lift and round along v to make the other component at most $\|v\|/2 \le m/(2\rho)$. The sum of the two errors is at most m^2/ρ. This proves the assertion, including the one-dimensional base case. □

In the mixed embedding of the CM field, the Euclidean dual of the lattice $2\varpi^{-1}\mathcal{O}$ is $\mathcal{O}$ itself: the Euclidean pairing with $2\varpi^{-1}a$ is the trace pairing with $\varpi^{-1}a$, and the trace dual of $\mathcal{O}$ is the inverse different by definition. The sign of the generator under conjugation does not change the integral dual. The shortest nonzero vector of $\mathcal{O}$ has length at least one, as the product of its squared complex absolute values is the nonzero integer norm. Lemma 16.13, with $m = 2f$, gives for every tuple (z_v) an element $x \in \mathcal{O}$ such that

$$|z_v - \Phi_v(x)| \le (2f)^2 |\Phi_v(\varpi)| \quad \text{for all } v.$$

Indeed apply the lemma to the tuple $z_v/\Phi_v(\varpi)$ and the lattice $2\varpi^{-1}\mathcal{O}$, then multiply each coordinate by $\Phi_v(\varpi)$. More generally a finite adele can first be approximated modulo $\mathcal{O}_f$ by a rational element, then this argument is applied to the remaining infinite coordinates. This derives `exists_rat_near_of_mem_differentIdeal` in `Lib/Adelic/Transference/MassBasic.lean`.

For any $t = (t_v) \in U_\infty$ apply the following row algorithm to its frame matrices. To reduce the entry (a, b) of the unitriangular frame matrix, $a < b$, choose an integral scalar x by the preceding approximation and multiply on the right by the Eichler row generator of Lemma 16.12 with free row supported in the column f_b and scalar $-Nx$. Its row-integrality condition holds: for an isotropic column of the frame the lattice basis is unchanged; for $f_b = z$ the coefficient Nx clears the $1/N$ of the marked basis vector. Hence the generator belongs to

Γ_c. Its normalized frame matrix is the elementary column operation subtracting from column b the normalized scalar

$$R_v\Phi_v(x), \qquad R_v = gN\, l_{vb}/l_{va} = gN\, \kappa_v^{a-b},$$

times column a. Apply the approximation to the tuple of entries divided by R_v. This makes the new entry (a,b) at most

$$|R_v|(2f)^2|\Phi_v(\varpi)| \le b^{23+60n+3} < b^{60h+91},$$

using $g \le b^{23}$, $N \le b^{60n}$, $2f \le b$, $|\kappa_v| \ge 1$ and $|\Phi_v(\varpi)| \le b$. Treat rows in descending order: this operation cannot change any row of larger index, because column a has zero entries below a. Treat all columns of the current row, then move to the next row. This constructs one rational $\gamma \in \Gamma_c$ for which all entries of $\mathrm{frameMatrix}(t_v\gamma_v)$ satisfy the stated bound.

For arbitrary y choose $t_v \in U_\infty$ with $\|t_v y_v\|_{\mathrm{op}} \le \mathrm{bs}(y_v)$: by definition some Bessel translate has norm plus one less than $\mathrm{bs}(y_v)+1$, and remove its compact factor without changing the operator norm. For the rational γ just found, the normalized frame of $\gamma_v^{-1}y_v$ is $y_v^{-1}\gamma_v f_b = (t_v y_v)^{-1}(t_v\gamma_v f_b)$. Its length is at most $\|t_v y_v\|_{\mathrm{op}}\, nb^{60h+91} \le H(y,v)$. Thus the denominator of (16.9) is positive everywhere, as claimed. This completes the construction of the box partition `isBoxPartition_epsilon` in `New/Filling/BoxPartition.lean` with its explicit bound. Mere density of E in its finite adeles would not have sufficed to obtain this archimedean reduction.

In contrast, for strong approximation of the unipotent group with no archimedean bound, the elementary coordinates above suffice. The group U is generated by Heisenberg-type Eichler transformations along the successive isotropic rows. Over a field this follows by cancelling its free entries in the same triangular order as in the proof of (16.8); the skew part is arbitrary, since one can write any skew element as $t - \bar{t}$ (use $t = s/2$ over a field of characteristic zero). The same polynomial parametrizations over the finite adeles are continuous, even at the dyadic components. The diagonal copy of E is dense in its finite adeles by the ordinary weak approximation and Chinese remainder theorem, allowing denominators at finitely many places. Approximating the entries of the finitely many generating parameters therefore shows $U(E^+)$ is dense in $U(\mathbb{A}_f)$. Consequently for any adelic $x \in U(\mathbb{A}_f)$ and finite-adelic $y \in G(\mathbb{A}_f)$ there is a rational $u \in U(E^+)$ with $y^{-1}x^{-1}uy \in K(g)$, because this is an open neighborhood condition. This is `Filling.exists_besselU_approx` in `New/Filling/NearTerms.lean`, in the appropriate marked basis.

The normalized additive character used in that statement is important. Let ψ_f be the usual finite trace character, taken so that its product with the standard infinite trace character is trivial on E. Let

$$\chi(u) = \psi_f(-\sigma(u)/(g\varpi)),$$

where σ is the sum of the simple frame-root entries. The character $x \mapsto \psi_f(-x/\varpi)$ has conductor $\mathcal{O}_f$, by the definition of the different (the annihilator of the integral lattice for the unscaled trace pairing is the inverse different). On integral U its values modulo $K(g)$ are therefore well defined. The same is true in the marked lattice coordinates: the change of basis fixes the frame, so the simple root sum is unchanged. Conjugating by H also leaves this sum unchanged since H fixes the frame. For rational u, after the infinite rescaling the simple-root sum at place v becomes $\Phi_v(\sigma(u))/\kappa_v$. Consequently

$$\prod_v \mathrm{wave}(\Phi_v(\sigma(u))/\kappa_v) = \exp(2\pi i\ \mathrm{Tr}_{E/\mathbb{Q}}(\sigma(u)/(g\varpi))) = \chi(u).$$

We have used $\bar{\varpi} = -\varpi$ and $\kappa_v = g\Phi_v(\varpi)/(2\pi i)$. In particular the finite and infinite characters cancel under the adelic product with the required opposite signs. If x and a rational

approximation u differ in the above congruence neighborhood for a class representative, then the values of χ agree: the quotient of the values is the scaled character of an integral simple-root sum in $g\mathcal{O}_f$. This gives `unipotentChar_eq_of_mem_principalLevel` and `exists_rat_approx_unipotentChar_eq` of `NearTerms.lean`. It is exactly the approximation needed for the orbital vanishing and for moving a Bessel element between class representatives.

16.10 Orbital cutoffs and Bessel transport between classes

We first give the precise quotient integration needed by the moment argument. For classes c, c' and $s \in \mathfrak{G}$, the set of terms is

$$\mathcal{T}(c, s, c') = \{\gamma \in G(E^+) : h_c^{-1}\gamma_f h_{c'} \in K,\ \mathrm{red}(h_c^{-1}\gamma_f h_{c'}) = s\}.$$

If T_0 is a subset of $G(E^+)$ stable under left and right $U(E^+)$ multiplication, set

$$A_{T_0}(c, s, c') = \sum_{\gamma \in T_0 \cap \mathcal{T}(c,s,c')} \int_{G_\infty} \beta_c(y)\beta_{c'}(\gamma_\infty^{-1}y)W(y) \otimes \overline{W(\gamma_\infty^{-1}y)}\, dy$$

where W is the product of the full or the cut local functions. We use the same absolute convergence of these orbital sums as in Section 14.9. The cutoff manipulations below are made first for nonnegative integrands of absolute values using Tonelli, and then for the vector integrals by absolute integrability. In particular, they do not constitute an independent proof of the uniform weighted summability used for the full moments.

A weight on a proper action of a countable discrete group can be changed without changing the integral of an invariant function, when its orbit sum has the same value. If F is invariant and both weights b_1, b_2 have sum one, insert $\sum_\delta b_2(\delta y)$ into the integral with weight b_1, interchange sum and integral, change variable by δ, and recover the integral with weight b_2. The same calculation holds for weights whose total sums on orbits agree. Apply this on each rational double coset for the commuting left and right actions of $\Gamma_c, \Gamma_{c'}$: if γ is its representative, its stabilizer is

$$I_\gamma = \Gamma_c \cap \gamma\Gamma_{c'}\gamma^{-1}.$$

Accounting for all the representatives of a double coset is equivalent to replacing the two cutoffs by one cutoff for I_γ. To check this, first sum over right translates of γ and use the cutoff identity for $\beta_{c'}$; then sum over the left translates modulo I_γ and use the corresponding cutoff identity for β_c. The covariance of the Bessel kernel (both characters are trivial on the congruence subgroups) allows the changes of variable in the integrals. Thus that double coset contributes the integral of

$$y \longmapsto W(y) \otimes \overline{W(\gamma^{-1}y)}$$

on $I_\gamma\backslash G_\infty$ with its quotient measure. The equality is initially one for finite unions of cosets with absolutely integrable summands; the absolute-value calculation extends it to the summable series. This is the argument of `hasSum_amplitudeSum_change_of_weight` in `New/Filling/AmplitudeSums.lean`.

Suppose $u \in U(E^+)$ and $u' = \gamma^{-1}u\gamma \in U(E^+)$. On the kernel in the last display, left multiplication by u_∞ contributes the character ratio $\chi(u)/\chi(u')$ (or its conjugate according to the Bessel sign convention). For the vanishing of the double coset integral it is necessary to use a map preserving the quotient measure on $I_\gamma\backslash G_\infty$: a rational element of the intersection need not normalize I_γ.

Here is a quotient argument avoiding such an assumption. Let $U_\gamma = U_\infty \cap \gamma_\infty U_\infty \gamma_\infty^{-1}$. This is the infinite group of the rational unipotent intersection, and I_γ is an arithmetic lattice

in it. Its finite covolume follows from the same triangular-coordinate argument as for a unipotent group: a unipotent group over a characteristic-zero field has a central filtration with vector-group quotients, and intersecting any compact open subgroup at the finite places gives full rank lattices in the rational quotient coordinates. One way to obtain this filtration here is to take logarithms of the upper unitriangular matrices (a finite polynomial), intersect their nilpotent Lie algebras, and use the central series and the exponential polynomial; a lattice in each quotient admits bounded fundamental sets, and lifting bounded sets successively gives finite covolume. Rational approximation in these polynomial coordinates also holds. The local congruence condition defining I_γ is compact open in the finite intersection.

The function on G_∞ is invariant under I_γ. Upon disintegrating its quotient integral along $I_\gamma\backslash U_\gamma$, the inner integral on each orbit is a constant vector times the integral of the character ratio on that compact quotient. Since the ratio is trivial on I_γ, it descends to the quotient; if it is nontrivial on U_γ, its integral is zero by right Haar invariance of the quotient (which does not require I_γ normal). Rational approximation of finite characters and the equality of infinite and finite characters above show that nonagreement on rational points gives such nontriviality. Hence stable sets of terms for which the two characters never agree have zero total amplitude. In the notation of the source the agreement condition `Filling.CharsAgree` is exactly equality of $\chi(u)$ and $\chi(u')$ for every rational pair $u\gamma = \gamma u'$ in the unipotent groups.

The next transport identity concerns multiplying a term by a rational Bessel element. For $s \in \bar{S}_B$, fix c and lift s to $k \in K \cap S_B(\mathbb{A}_f)$. Decompose $h_c k = u_f h_f$ under UH. By the class decomposition of h_f, there are a rational $h \in H(E^+)$, a class d, and a $k_c \in K_c$ with $h_f = h h_d k_c$. Approximate u_f rationally by an element of U modulo the required neighborhood conjugated by h_f; this is the strong approximation proved in Section 16.9. Then the product $\delta = uh$ is a rational Bessel element lying in $\mathcal{T}(c, s, d)$. Thus for every c a Bessel lift over s exists.

For $s = 1$, a rational term in S_B forces $c = d$ and lies in Γ_c. Indeed its factor in H has position in K_c (integrality and congruence of the H factor follow by the block-diagonal argument of (16.6)), hence the two classes agree and the rational H element fixes a deeper-level class. Its stabilizer is trivial. This proves `eq_and_mem_classUnip_of_besselS` and, with the preceding approximation, `exists_besselS_mem_termsOver` in `New/Filling/NearTerms.lean`. In general the map $c \mapsto d$ produced by a lift of s is a permutation: if two classes have the same d, multiplying one lift by the inverse of the other gives a Bessel term over the identity, forcing equality of the classes, and a finite injective self-map is bijective.

As a consequence of the quotient weight-change argument, transporting terms over s by such Bessel lifts permutes the class blocks and multiplies the corresponding coordinate factors by unitary compact type matrices and complex phases of absolute value one. More explicitly, if $\delta \in \mathcal{T}(c, s, d) \cap S_B$, right multiplication by δ^{-1} maps the appropriate sets of terms over $(c_0, s_0 s, d)$ bijectively to those over (c_0, s_0, c); it conjugates the right unipotent congruence subgroup at d to the one at c because the finite position is in K and $K(g)$ is normal in K. The replacement of the cutoff for d by the transported cutoff for c is valid by orbit sums. The tensor factor in the integral changes by the representation of the compact factor of δ and its character. Both have norm one. This yields precisely the covariance needed for the equality of Hilbert–Schmidt block sums below. It is the mathematical content of `amplitudeSum_move` in `New/Filling/AmplitudeSums.lean`.

There is no Bessel-subgroup contribution over 1 except the class-diagonal integral. Indeed for $c = d$ the surviving terms are exactly Γ_c, and summing the two cutoff factors reduces the integral to ν times the normalized Bessel inner products on $U_\infty\backslash G_\infty$. Schur orthogonality with the chosen normalization of the Bessel function gives the identity on the coordinate

space, so

$$\nu^{-1}A_{S_B}(c,1,d)_{ij} = \delta_{c,d}\delta_{i,j} \quad \text{for the full functions.}$$

This is `amplitudeSum_besselS_one`, and explains the identity term in the first moment. The same covariance and weight-change argument works for the cut functions, although the coordinate inner product need then no longer be exactly the identity.

16.11 Agreeing characters and propagation of a shared flag vector

We now prove the algebraic elimination of degenerate terms, which was an input in Section 14.6. The notation here is rational: $F = (f_0, \dots, f_h)$ is the standard frame, $f_h = z$, and Fx denotes the corresponding linear combination of its columns. For a rational γ put

$$Z = F^* J_N \gamma F, \qquad D_j = \det Z_{\{0,\dots,j-1\}}, \qquad \mathcal{D}(\gamma) = D_{h+1}\bar{D}_{h+1} - N^2 D_h \bar{D}_h,$$

where J_N is the standard rational Gram matrix with middle entry $-N$ on z (and $D_0 = 1$).

Lemma 16.14 (Shared flag-vector propagation)**.** *Suppose the unipotent characters agree for every rational pair $u\gamma = \gamma u'$. If $Fx = \gamma Fy$ and both coefficient vectors vanish at indices $> a$, with $x_a = c$ and $y_a = c'$, then for every $b \geq a$ such a common vector exists with leading coefficients still c, c'. In particular $c\bar{c} = c'\bar{c}'$.*

Proof. It suffices to propagate from $a < h$ to $a + 1$. Let ℓ_0 be any rational linear form on V vanishing on the subspace spanned by $f_i, \gamma f_i$ for $i \leq a$. Use nonsingularity of J_N to write $\ell_0(v) = h(w, v)$. With $v_0 = Fx = \gamma Fy$, the Eichler transformations

$$E(r) = 1 + r v_0 h(w, \cdot) - \bar{r}\, w h(v_0, \cdot) - r\bar{r}\,\tau h(w, w) v_0 h(v_0, \cdot), \qquad r \in E,$$

lie in U, and their conjugates by γ^{-1} also lie in U. Here $\tau + \bar{\tau} = 1$, and no integrality is required. To verify the assertion, the support of v_0 in the isotropic frame and the vanishing of ℓ_0 on the initial flag imply upper unipotence in the block order. Expanding the Gram form gives unitarity as in Lemma 16.12, with v_0 isotropic and $h(v_0, w) = 0$. On the simple superdiagonal entries of the frame only the term $r v_0 h(w, \cdot)$ contributes, since $h(v_0, f_j) = 0$ for every frame vector. Thus the difference of the two simple-root sums is

$$r\,\ell_0\left(\sum_{i \leq a} x_i f_{i+1} - \sum_{i \leq a} y_i \gamma f_{i+1}\right).$$

Agreement of the rational characters for all r forces the parenthesized vector to lie in the annihilator of every such ℓ_0: the finite character on $\mathbb{A}_f$ is nondegenerate, and E is dense in $\mathbb{A}_f$, so if $\psi_f(rt/(g\varpi)) = 1$ for all $r \in E$ and $t \in E$, then $t = 0$. The annihilator of the forms vanishing on a subspace is that subspace itself. Consequently the displayed vector belongs to $\langle f_i, \gamma f_i : i \leq a \rangle$. Moving these lower terms to the two sides gives a common vector supported at indices $\leq a + 1$ with leading coefficients c, c' at $a + 1$, as required.

Iterate up to $b = h$. The norm of a vector in the frame span is $-N$ times the norm of its last coefficient, because all other frame entries are isotropic and orthogonal to the entire frame. The two equal vectors at that last stage therefore have $-Nc\bar{c} = -Nc'\bar{c}'$, proving the last assertion. □

This is `Filling.hasCommonFlagVector_of_le` of `New/Filling/CrudeEnergy.lean`, using `exists_mem_besselU_eq_eichlerMap` of `New/Filling/EichlerFlag.lean`. In particular no assumptions about the norm of the original rational matrix are needed in this propagation lemma.

There is always a simple nonnegativity relation for $\mathcal{D}(\gamma)$. Suppose $Zy = Df_{\mathrm{last}}$ in coefficient notation (that is, all entries except the last of Zy vanish). Put $t = \gamma Fy$. Its e_a^* coordinates vanish, and its z-coordinate is $-D/N$. Write t_0 for its component on the positive definite X. Unitarity gives, at every infinite embedding,

$$N\|t_0\|_v^2 = |\Phi_v(D)|^2 - N^2|\Phi_v(y_h)|^2. \tag{16.10}$$

In particular equality of the right-hand terms at any place implies $t_0 = 0$ (all X-coordinates are rational and the ordinary positive norm is definite). Taking $y = \mathrm{adj}(Z)(0,\dots,0,1)^t$ shows $D = D_{h+1}$ and $y_h = D_h$. Therefore $\mathcal{D}$ is nonnegative at every infinite place. If $\mathcal{D} \neq 0$, then $D_{h+1} \neq 0$. This proves the discriminant and normalization content of `norm_sq_middle_part_eq`, `map_termDisc_eq_norm` and `hasLevel_last` in `NearTerms.lean`.

Proposition 16.15 (Levels of all relevant degenerate terms)**.** *Under rational character agreement, if $\mathcal{D}(\gamma) = 0$, there is an integer $j \in \{0,\dots,h\}$ for which $D_j \neq 0$ and a scalar $\alpha\bar{\alpha} = 1$ such that for every $a \geq j$ the two frames share a vector with leading coefficients $\alpha, 1$ at index a. Together with $\mathcal{D}(\gamma) = 0$ these are precisely the conditions of `Filling.HasLevel` for level j in `CrudeBoundDefs.lean`. If $\mathcal{D} \neq 0$, the term has level $h+1$.*

Proof. If $\mathcal{D} = 0$, first find a nonzero vector shared by the frame spans. When $D_{h+1} \neq 0$, use the adjugate vector just used in (16.10). Its image γFy lies in the span of F since its positive definite component vanishes. If $D_{h+1} = 0$, choose any nonzero y in the kernel of Z instead. Formula (16.10) now implies $y_h = 0$ and $t_0 = 0$; again the image is a nonzero vector in the span of F.

Write such a relation as $Fx = \gamma Fy$ and let a be the largest index with either coefficient nonzero. By Lemma 16.14, the leading coefficients have equal norms, so both are nonzero. Their ratio $\alpha = x_a/y_a$ has norm one. After normalizing by y_a, the propagated relation shows that the same ratio occurs at every index $\geq a$. In fact the ratio of leading coefficients of any nonzero shared relation is unique: propagate two such ratios to index h, normalize their last y-coefficients to one, and subtract. If the ratios were unequal the resulting common vector would have one last coefficient zero and the other nonzero, contradicting equality of norms. Thus let j be the least index for which a shared relation occurs; it has the stated common ratio at every index $\geq j$.

We show $D_j \neq 0$. If $j = 0$ this is the convention $D_0 = 1$. Otherwise suppose the $j \times j$ leading block is singular and let $y \neq 0$ be a vector supported on indices $< j$ in its kernel. Put $t = \gamma Fy$. Its pairings with $f_0,\dots,f_{j-1}$ vanish. For every $a \geq j$ there is a common vector whose coefficient at f_a on the standard side is $\alpha \neq 0$, with other terms only at lower indices. The pairing of that vector with t is zero on the γ-frame side, since $y_h = 0$. By induction up the standard frame this implies $h(f_a, t) = 0$ for all a, including the last vector. Hence $Zy = 0$. Formula (16.10) shows $t \in \operatorname{span} F$; write $t = Fx$. Taking the largest index of a nonzero coefficient in x or y, the propagation lemma says neither leading coefficient can vanish. It follows that this largest index is $< j$, contradicting minimality. Thus $D_j \neq 0$.

If $\mathcal{D} \neq 0$ the leading full minor is nonzero as noted above, the normal-form condition at level $h+1$ is empty, and the condition on $\mathcal{D}$ is exactly the top condition. □

The preceding proposition is `Filling.CharsAgree.exists_hasLevel` in `New/Filling/CrudeEnergy.lean`. The set of terms of each level is stable under the two unipotent actions: these

actions change the frame columns by unitriangular matrices, preserving leading-coefficient relations up to lower terms and preserving every leading principal determinant of Z by multiplying on the corresponding sides by unitriangular matrices. In particular the discriminant $\mathcal{D}$ is preserved. This also gives the stability statements in `New/Filling/StableCells.lean` required to split the energy according to levels.

16.12 The link as a Hermitian Lagrangian and its congruence defect

The link is the image of $L^\perp \cap (\gamma L)^\perp$ in

$$(L^\perp/L) \oplus ((\gamma L)^\perp/\gamma L) \simeq M \oplus M$$

under $t \mapsto ([t], [\gamma^{-1}t])$, using γ^{-1} to identify the second middle quotient. Denote it by $\mathcal{L}_\gamma$. Give $M \oplus M$ the Hermitian form $T \oplus (-T)$, $T = \mathrm{diag}(1_k, -N)$. For two vectors in the defining intersection, their first Hermitian pairing agrees with their second by unitarity of γ. Thus this image is isotropic. If $a = \dim(L \cap \gamma L)$, the dimension of the defining intersection is

$$r - \dim(L + \gamma L) = r - (2h - a) = k + 1 + a.$$

The kernel of the map is precisely $L \cap \gamma L$. Consequently the image has dimension $k + 1$, half the ambient dimension. It is therefore equal to its orthogonal complement. This proves `me m_link_iff_orthogonal` in `StableCells.lean` and `finrank_link` in `BoxPartition.lean`. Invariance under U on both sides follows because U preserves L, acts trivially on $L^\perp/L$, and applying a left translate of a defining vector gives a corresponding translate for the new link. On a big-cell representative $w_p\sigma$, with $\sigma \in U(M)$ and the Weyl element w_p exchanging L and $L^\vee$, the common orthogonal intersection is M; the link is the graph of σ^{-1} and equality of such links is equivalent to equality of the middle transformation.

Let $\ell_z : M \to E$ be the form $h(z, \cdot)$. The source puts the positive definite metric on $M \oplus M$

$$Q_{\mathrm{link}} = \tfrac{1}{2}(T \oplus T) + N^{-1}(\ell_z \oplus \ell_z)^*(\ell_z \oplus \ell_z).$$

Here the second summand is the sum of the squares of the two separate forms, not the square of their sum. In the standard coordinates this matrix is diagonal: the positive definite coordinates have weight $1/2$ and each z-coordinate has weight $-N/2 + N = N/2$. Hence it is positive at every infinite embedding. In particular its Gram determinants on any nonzero subspace are nonzero; it is the positive height used in the pair-counting estimates, not the isotropic link pairing itself.

We prove the particular residue-field defect bound that controls the congruence gain in those estimates. Fix a split level place w (or its conjugate) with residue field $\mathbb{F}_q$, and let $\bar{\gamma}$ denote the reduction of the integral finite position of γ in the marked basis. Write F also for the reduced frame. The relevant form matrices at the two factors are related by the adjoint inverse; one may choose whichever factor is specified by the vector-fixer condition.

Lemma 16.16 (Residue-field propagation into the two last columns)**.** *Suppose the reduced position has agreeing characters in the sense that any two integral unipotent elements with reductions conjugate by it give the same unipotent character. Suppose a nonzero frame vector is fixed by $\bar{\gamma}$ at one specified split factor. Let P be any matrix of two row forms on the reduced space each vanishing on the standard and translated isotropic frame vectors. Then the 2×2 matrix of their values on $z, \bar{\gamma}z$ is singular.*

Proof. Work at the specified split factor; a vector fixer written using inverse adjoint simply means work at its conjugate factor. If the 2×2 determinant were nonzero, any shared relation $Fx = \bar{\gamma} Fy$ would have $x_h = y_h = 0$, by applying the two forms and subtracting. The fixed nonzero frame vector supplies such a shared nonzero relation with equal coefficient vectors.

Suppose one has a shared relation with both last coefficients zero and largest nonzero standard coefficient x_a at $a < h$. Let m be any residue-field linear form annihilating f_i for $i \leq a$ and $\bar{\gamma} f_i$ for $i < h$. The transvections

$$1 + u(Fx)m, \qquad 1 + u(Fy)m\bar{\gamma} \quad (u \in \mathbb{F}_q)$$

are conjugate by the reduced position and strictly upper block unitriangular off the identity. Indeed the left vector is supported in earlier isotropic rows and the right form vanishes on the initial flag through those rows. To lift them to integral U, lift these vectors and forms separately at the specified split factor, keeping their indicated zero entries zero. Their rank-one cross terms square to zero. At the opposite split factor set the unitary inverse adjoint, and at every other finite place use the identity. This produces integral unipotents. The opposite factor of a transformation $1 + u(Fx)m$ has its nonidentity entries in columns dual to the support of Fx; these are isotropic dual columns, not frame columns. Indeed the inverse is $1 - u(Fx)m$, and the adjoint inverse has its changed columns given by $(Fx)^* J$, which is supported on the dual columns. Thus at that conjugate factor the simple-root sum of the frame entries vanishes, and similarly for the second transvection.

At the selected factor the additive character is primitive: the different-normalized finite character has conductor $\mathcal{O}_f$, and multiplication by $1/g$ tests the full residue field at that place. Agreement on varying u consequently implies

$$m \left(\sum_{i \leq a} x_i f_{i+1} - \sum_{i<h} y_i \bar{\gamma} f_{i+1} \right) = 0.$$

Here the sums are the simple superdiagonal entries of the two rank-one operators. As m ranges over the indicated annihilator, the parenthesized vector is in the span of the standard frame through a and the translated isotropic frame. Moving these terms to the two sides gives a new shared relation with standard coefficient at $a + 1$ equal to $x_a \neq 0$, other standard coefficients of lower index, and translated coefficients allowed through the last vector. The alleged invertibility of the two-form determinant forces both last coefficients zero in every shared relation. Iterating the propagation gives a relation with nonzero standard last coefficient, a contradiction. Hence the determinant is singular. □

This is the residue-field argument in `Filling.le_genIndex_det_of_fixesNormAtLeast`, `New/Filling/BoxPartition.lean`. Notice its distinction from Lemma 16.14: over one split factor it is not necessary to recover a norm relation or to propagate with exactly matching upper supports on the two sides. The invertible two-form determinant was used instead to exclude the last coordinates of every shared relation.

If the two forms in Lemma 16.16 are finite adelic integral forms annihilating the two isotropic frames, put δ equal to the determinant of the matrix of their values on the two last frame vectors. For every place at which the reduced position satisfies the corresponding vector-fixer condition and has agreeing characters, δ is divisible by that prime. Therefore

$$[\delta] \geq \prod_{w \text{ fixed}} q_w \tag{16.11}$$

if δ is an idele; here the generalized index on the left is the absolute ideal index. This follows either by taking the local valuations or by mapping the quotient of $\mathcal{O}_f$ by $\delta \mathcal{O}_f$ onto the

product of the corresponding residue fields. Requiring agreement in the finite reductions, rather than only agreement on rational intersections, is what permits the use of the separate residue-field tests.

Finally we record a direct comparison of this type of ideal index with the positive link height. For two rational rectangular matrices A, B of the same full column rank in a positive definite Hermitian space with form Q, at every infinite embedding the exterior-power Cauchy–Schwarz inequality gives

$$|\det(A^*QB)|^2 \le \det(A^*QA)\det(B^*QB). \tag{16.12}$$

If their adelic changes of basis are $a, b \in \mathrm{GL}_s(\mathbb{A}_f)$, multiplicativity of the determinant and the product formula give

$$[\det(a^*A^*QBb)]^2 \le [\det a]^2[\det b]^2|\operatorname{Norm}_{E/\mathbb{Q}}\det(A^*QA)||\operatorname{Norm}_{E/\mathbb{Q}}\det(B^*QB)|.$$

Indeed the finite generalized index of a nonzero rational number is the absolute field norm, and taking the product of (16.12) over all complex embeddings, together with their conjugates, gives the required norm inequality. This is `liouville_gram` in `BoxPartition.lean`. It explains the appearance of powers of positive sublattice discriminants when a finite index such as (16.11) is large. The remaining precise comparison of the pair invariants and the class-size estimate in `LinkPairing.lean` and `BlockBound.lean` are not automatic from this inequality alone, since those bounds require controlling the ranks and saturation indices of several simultaneous lifts.

16.13 Saturated lattice heights and their zeta sum

We prove the subspace counting estimate needed both for crude link weights and for the sharp global sieve. Let $V_0 = E^m$ with its standard positive definite Hermitian form, and let $\mathcal{M}$ be a finite adelic unimodular Hermitian lattice in V_0. Write $L_{\mathcal{M}} = V_0 \cap \mathcal{M}$. If $W_0 \subset V_0$ is a rational subspace of dimension s, choose an E-basis matrix Y for W_0, and a finite adelic matrix $a \in \mathrm{GL}_s(\mathbb{A}_f)$ such that $Ya\mathcal{O}_f^s = W_0(\mathbb{A}_f) \cap \mathcal{M}$. Such a basis exists locally by the elementary divisor theorem over each discrete valuation ring; outside finitely many places the rational basis is already integral and primitive. Define the height

$$d_{\mathcal{M}}(W_0) = [\det a] \prod_{v|\infty} \det\left(\Phi_v(Y)^*\Phi_v(Y)\right). \tag{16.13}$$

The expression does not depend on the choices: changing the rational basis by $P \in \mathrm{GL}_s(E)$ multiplies the product of Hermitian determinants by $|\operatorname{Norm}\det P|$ and the finite index by its reciprocal, while changing an integral local basis has finite index one. In the source this is the discriminant $\sqrt{[\det(Ya)^*(Ya)]}$, `sublatticeDisc` of `New/Filling/GramLiouville.lean`. The equalities follow by taking determinants and the product formula at all complex places.

The Euclidean covolume in $W_0(\mathbb{R})$ of its rational lattice $W_0 \cap L_{\mathcal{M}}$ equals

$$d_{\mathcal{M}}(W_0)\operatorname{vol}(\mathcal{O}_f)^{-s}. \tag{16.14}$$

Indeed the preimage of its lattice in E^s has finite coordinates $a\mathcal{O}_f^s$, hence ordinary covolume $[\det a]\operatorname{vol}(\mathcal{O}_f)^{-s}$. At each complex place the real Jacobian of the injective map Y onto its image is $\det(Y^*Y)$ (singular values occur twice in real coordinates). Multiply these Jacobians. This proof handles arbitrary nonprincipal lattice modules by working locally.

Two elementary observations are useful.

- If Ya' is another integral full rank array of the same subspace then $Ya' = YaP$ with P integral at the finite places, since the columns Ya form a saturated integral basis. The corresponding Gram determinant is multiplied by $[\det P]^2$. If in addition a linear image of the columns of Ya' lies in some finite-index submodule of the image lattice of Ya, the index i of that submodule divides the index of the column sublattice; in particular $i \le [\det P]$. Therefore the squared Gram index of Ya' is at least $(d_{\mathcal{M}}(W_0)i)^2$. This proves the primitive-basis and index calculation in `sublatticeDisc_mul_index_sq_le`.

- Heights of complementary orthogonal subspaces agree:

$$d_{\mathcal{M}}(W_0) = d_{\mathcal{M}}(W_0^{\perp}). \tag{16.15}$$

  To check it locally, extend the saturated integral basis of W_0 to a basis matrix M of $\mathcal{M}$. This extension is possible because a saturated submodule over a discrete valuation ring has torsion-free quotient and is a direct summand. The last $m - s$ columns of the Hermitian dual basis give a basis of the orthogonal-complement lattice: they pair to zero with the first s columns, and unimodularity gives their integrality. Their Gram determinant is the complementary block minor of the inverse full Gram matrix. By Jacobi's identity it equals the Gram determinant of the first s columns divided by the full Gram determinant. The latter is a local unit at every finite place by unimodularity. Take the generalized indices and square roots, which is legitimate because these determinants are nonzero (the rational spaces are positive definite). This gives (16.15).

Lemma 16.17 (A quantitative Siegel lemma for the orthogonal complement)**.** *Suppose* $\dim_E W_0 = m - s$, $1 \le s \le m$. *Then its orthogonal complement has an E-basis $x_1, \dots, x_s$ in $L_{\mathcal{M}}$, with Euclidean mixed-embedding norms satisfying*

$$\left(\prod_{i=1}^{s} \|x_i\|\right)^{2f} \le (8fs\ \mathrm{rd}_E)^{fs} d_{\mathcal{M}}(W_0), \qquad \mathrm{rd}_E = |D_E|^{1/(2f)}.$$

Proof. Apply Minkowski's second theorem to the Euclidean lattice $W_0^{\perp} \cap L_{\mathcal{M}}$ in real dimension $2fs$. If its successive minima are $\lambda_1, \dots, \lambda_{2fs}$ for the ordinary Euclidean ball, choose successively E-linearly independent vectors of norm at most $\lambda_{2f(i-1)+1}$ for $1 \le i \le s$. This is possible because the previously chosen E-span has real dimension $2f(i-1)$, whereas the next successive minimum gives more independent real vectors. Since the minima are nondecreasing, the product of their chosen powers $2f$ is at most the product of all the real successive minima. By Minkowski's second theorem and (16.14)–(16.15),

$$\prod_{j=1}^{2fs} \lambda_j \le \frac{2^{2fs}(fs)!}{\pi^{fs}} \left(\frac{|D_E|^{1/2}}{2^f}\right)^s d_{\mathcal{M}}(W_0) \le (8fs\ \mathrm{rd}_E)^{fs} d_{\mathcal{M}}(W_0).$$

In the last step use $(fs)! \le (fs)^{fs}$, $|D_E|^{s/2} = \mathrm{rd}_E^{fs}$, and $\pi \ge 1$. This proves the result. □

This recovers the statement `Filling.siegel_lemma` in `New/Filling/NearTerms.lean`; the field-degree powers are important, but no uniform dependence on the ideal class is required from a separate choice of generators.

Theorem 16.18 (Sublattice height zeta bound)**.** *With the same notation, over all E-linear subspaces of codimension s,*

$$\sum_{W_0} d_{\mathcal{M}}(W_0)^{-(m+1)} \le (200fm\ \mathrm{rd}_E)^{f(m+1)s}.$$

For $s = 0$ the sum consists of the full space with height one.

Proof. Every nonzero vector x of $L_{\mathcal{M}}$ has mixed-embedding norm at least one. Indeed its Hermitian norm $h(x,x) \in E^+$ is totally positive and integral, since its pairing with the unimodular integral lattice is integral. Hence the product of its positive real conjugates is an integer at least one, and their sum is at least one by the arithmetic-geometric mean inequality. Distinct lattice vectors are also separated by distance one.

In a Euclidean space of real dimension d, a set of unit-separated points in a ball of radius R has at most $(1+2R)^d$ elements, by packing disjoint balls of radius $1/2$ and comparing volumes. Nonzero lattice points with norms in $[2^j, 2^{j+1})$ consequently number at most $5^d 2^{jd}$. For $p \geq d+2$ this gives

$$\sum_{0 \neq x \in L_{\mathcal{M}}} \|x\|^{-p} \leq 5^d \sum_{j \geq 0} 2^{-j(p-d)} \leq 5^p.$$

Use $d = 2fm$, $p = 2f(m+1)$.

For each subspace W_0 select one ordered basis $\boldsymbol{x}$ of its orthogonal complement as in Lemma 16.17. The selection is injective in W_0, because the span of the tuple determines $W_0^{\perp}$. Rewrite that lemma as

$$d_{\mathcal{M}}(W_0)^{-(m+1)} \leq (8fs \ \mathrm{rd}_E)^{fs(m+1)} \prod_{i=1}^{s} \|x_i\|^{-2f(m+1)}.$$

Sum over the selected tuples, and enlarge to all tuples of s nonzero lattice vectors. The separated-point sum bounds this by $(8fs \ \mathrm{rd}_E)^{fs(m+1)} 5^{2fs(m+1)} \leq (200fm \ \mathrm{rd}_E)^{fs(m+1)}$. All summands are nonnegative and the bound holds on every finite subset, which also proves convergence of the infinite sum. □

This proves `tsum_sublatticeDisc_inv_le`, `New/Filling/NearTerms.lean`. It accounts for the denominator lattices in the discriminant counts, not just for points with integral coordinates. In particular rational lines in E^2 can have varying nonprincipal denominator ideals without invalidating this summation.

16.14 The global frequency estimate and explicit sieve budget

We apply Theorem 16.18 to the nonzero rational T-integers y used as frequencies, where T is a conjugation-stable finite set of finite places. Write

$$[y]_T = \prod_{w \in T} |y|_w^{-1}, \quad d_y = [\mathcal{O} : \mathcal{O} \cap y^{-1}\mathcal{O}], \quad L_K(t) = t^K \min(1,t)^6.$$

The denominator index d_y only involves places in T, since y is integral elsewhere. The height of the line through $(1,y)$ in E^2 with its standard Hermitian lattice is

$$d(E(1,y)) = d_y \prod_v (1 + |\Phi_v(y)|^2). \tag{16.16}$$

Indeed its saturated column lattice consists exactly of the vectors $x(1,y)$ with $x \in \mathcal{O} \cap y^{-1}\mathcal{O}$. This proves the formula also for nonprincipal denominator ideals, by taking the local basis of that fractional ideal and then the generalized index as in (16.13).

Suppose $|c_v|^{-1} \leq |x_v| \leq B$ with $B \geq 1$. The product formula and integrality off T give $[y]_T^{1/2} \leq \prod_v |\Phi_v(y)|$. Also

$$\min(1, B/t)^6 \leq ((1+B^2)/(1+t^2))^3 \quad (t > 0).$$

Hence for every integer $K \geq 0$,

$$\frac{[y]_T^{K/2}}{d_y^6} \prod_v L_K(|c_v \Phi_v(y)|^{-1}) \leq \left(\prod_v |x_v|^K (1+B^2)^3 \right) d(E(1,y))^{-3}.$$

Here we used $d_y^{-6} \leq d_y^{-3}$. The line map $y \mapsto E(1,y)$ is injective. Apply Theorem 16.18 with $m=2$, $s=1$. Since $400f \leq b$ and $1+B^2 \leq bB^2$ under the filling hypotheses, every finite sum of these frequency terms is bounded by

$$\left(\prod_v |x_v|^K (bB^2)^3 \right) (b \operatorname{rd}_E)^{3f}. \tag{16.17}$$

This derives `sum_frequencies_le` in `New/Filling/GlobalSieve.lean`, with its exact damping exponent and the correct denominator index.

We now verify the numerical input into the one-sided sieve induction of Section 15.10; in particular the two broad bounds in that section are no longer left as independent estimates. We use $n = h+1$ as above, $K_1 = k - n^7$, $K_0 = n^6$. The standing hypotheses include

$$b \geq 1000, \quad n \geq 10^5 + 1, \quad k \geq 1000(2n)^9, \quad k \leq f \leq b/1000, \quad \operatorname{rd}_E \leq b, \quad b^{19} \leq g \leq b^{23},$$

as well as $N \leq b^{60n}$, $1 \leq |\kappa_v| \leq b^{23}$. Put $s_{va} = \kappa_v^{2(h-a)}/N > 0$ as in (16.8); the source calls these `frameScaleSq`. The per-entry Jacobian in the Poisson argument is

$$C_a = \operatorname{vol}(\mathcal{O}_f) / \prod_v s_{va}^2.$$

The nonzero frequency argument at step a is $(c_v \Phi_v(y)/s_{va})_v$ at infinity and y at the finite places. Indeed the character on one normalized infinite variable is $\operatorname{wave}(z) = \exp(z - \bar{z})$, and the rational global character with frequency y is obtained at infinity by $2\pi i\, \Phi_v(y)/\Phi_v(\varpi)$. Thus the Poisson choice satisfies $|c_v|^{-1} = |\kappa_v|/g$. For unequal entries the difference is $y = g^{-1}(t^{-1} - u^{-1})$, and the infinite factor has $c_v = g$, so again $|c_v|^{-1} \leq |\kappa_v|/g$. These are the two cases of the frequency test (15.20). In either case the inverse magnitude of the infinite scale, after division by s_{va}, is bounded by

$$|x_v| = |\kappa_v|^{2h+1}/(gN) \leq b^{46n}.$$

This is the quantity used in the numerical hypothesis `prod_band_le_inv_deg`.

The archimedean frequency majorant given by the local bandlimit input has one factor

$$\Theta\, b^{2000n^9} L_{K_1}(|c_v \Phi_v(y)/s_{va}|^{-1}), \qquad \Theta = \operatorname{vol} U(k)\, \pi^{(2h+1)k} / \operatorname{vol} U(r).$$

Its finite factor is $[y]_T^{K_1/2} d_y^{-6} b^{n^9 f}$. Applying (16.17) with $B = b^{46n}$, the entire frequency sum is at most

$$\varepsilon = \left(\prod_v \Theta |x_v|^{K_1} \right) b^{(2001n^9 + 276n + 9)f}. \tag{16.18}$$

The exponent $276n+9$ contains the product of the $B^6 b^3$ term and $(b \operatorname{rd}_E)^3$ at each place.

For the head-weight estimate (15.21), let A be the finite set of rational units used in the induction, whose elements and inverses are T-integral. For a head $(p_0, \dots, p_{i-1}) \in A^i$, put $D_j = p_0 \cdots p_{j-1}$ for $1 \leq j \leq i$. These partial products uniquely determine the head. The local weight is the product over j of damped powers with exponent K_0 at inverse absolute value of $D_j \prod_{a<j} s_{va}$, and the finite weight is the product of $[D_j]_T^{K_0/2}$ if each D_j is integral at T, zero otherwise. In the nonzero case D_j is globally integral, so its denominator index is one. Thus

the sum over heads is bounded by the product of the sums over the allowed individual D_j, by injectivity of the partial-product map.

For each such sum use (16.17) with inverse scale bounded by N^i, because each $s_{va} \geq N^{-1}$. Set $B = N^i \leq b^{60ni}$. With $K = K_0$ each sum is at most $b^{f(9+60ni(K_0+6))}$. For the square of the head sum the exponent is therefore at most $f(18i + 120ni^2(K_0 + 6))$. On the other hand

$$C_a \geq b^{-(92n+1)f}$$

because $s_{va} \leq b^{46n}$ and $\mathrm{vol}(\mathcal{O}_f) = (2/\operatorname{rd}_E)^f \geq b^{-f}$. Dividing the squared head bound by $\prod_{a<i} C_a$, and using $i \leq h < n$, $K_0 = n^6$, gives an exponent at most

$$f(120n^9 + 720n^3 + 92n^2 + 19n) \leq 121fn^9.$$

Thus one may take $\Gamma = b^{121fn^9}$, including the empty head at $i = 0$.

By the admissibility condition for the product of band constants,

$$\left(\prod_v \Theta|x_v|^{K_1}\right) b^{7kf} \leq d_\infty^{-1}, \qquad d_\infty = \prod_v d_{\lambda_v}.$$

Combining this with (16.18), the bound for Γ, and $k \geq 1000(2n)^9$, we obtain

$$\Gamma\varepsilon \leq d_\infty^{-1} b^{-7kf+(2122n^9+276n+9)f} \leq d_\infty^{-1} b^{-7kf+2123n^9 f}.$$

Since $2h + 1 \leq b$ and $(k - 2123n^9)f \geq 11$,

$$(2h+1)\Gamma\varepsilon \leq d_\infty^{-1} b^{-6kf-10}. \tag{16.19}$$

The bottom constant of the product sieve system is the product of $(1 + b^{-10})/d_{\lambda_v}$ at infinity and the specified finite bottom constant e. Applying the one-sided Poisson induction of Section 15.10, with the parameters just proved, yields

$$\Re \sum_{t\in U,\ p,q\in A^h} J'_{h+1}(p,t;q,t) \prod_v J_{v,h+1}(p,t;q,t) \leq \left(\prod_{a=0}^{h} C_a\right)\left(e\prod_v \frac{1+b^{-10}}{d_{\lambda_v}} + \frac{b^{-6kf-10}}{d_\infty}\right),$$

where the arguments have the rational and scaled infinite interpretations of the source and $U \subseteq A$ is arbitrary. This is `Filling.sieve` in `New/Filling/GlobalSieve.lean`. In particular the global frequency estimate, partial-product bound and resulting power-saving error are derived here. The existence of the actual local positive kernel systems with the stipulated bottom and oscillatory bounds is a separate input, not a consequence of this global estimate.

16.15 The finite zero form and central projections

We discuss the zero-frequency finite group estimate separately because it is also needed for the projection onto good finite types in the first moment. Consider a conjugation-stable finite product of local fields containing the selected pairs over g. In it denote the integral compact subgroup by K_T, the level group by $K_T(g)$, and choose the local group Haar measure M. The function called `IntegralPlace.phase` in `New/Filling/TopFloor.lean` is

$$\phi(x) = \begin{cases} \psi(\sigma(u)/g), & x \in uK_T(g),\ u \in U, \\ 0, & \text{otherwise.} \end{cases}$$

It is well defined by the conductor of the self-dual additive character and the additivity of the simple-root sum. Restricted to K_T, it is the finite function on $\mathfrak{G}$ equal to $\chi(s)$ on $\bar{U}$ and zero outside $\bar{U}$. The zero form for an intermediate compact subgroup $K_T(g) \subset Z \subset K_T$ is

$$\mathcal{K}(Z) = \int_{G_T} \left| \int_Z \phi(yxy^{-1})\, dM(y) \right|^2 dM(x). \tag{16.20}$$

The following bottom estimate is still a quantitative finite-place Fourier-floor input, `floorZeroForm_le_main_term` in `TopFloor.lean`. For $t + 10^{10}n^4 \le k$ and $t \ge 0$, it says

$$\mathcal{K}(Z) \le M(Z)M(K_T(g))(M(Z) + e_T^{t+1}2^{-n}), \qquad e_T = \mathrm{vol}(R_T \setminus R_T^{\times}). \tag{16.21}$$

Here $R_T = \prod_{w \in T} \mathcal{O}_w$ with additive volume one in the local coordinate normalization. Also $M(K_T) \ge 2^{-n}$ when T consists of chosen level pairs; the source obtains this from the normalized box measure comparison. We retain that comparison as an input in applying (16.21). On the other hand the reverse bound

$$\mathcal{K}(Z) \ge M(Z)^2 M(K_T(g)) \tag{16.22}$$

is immediate since for $x \in K_T(g)$, all its conjugates by Z belong to $K_T(g)$ and have phase one.

We first show how the zero form yields the saving for fixed vectors on finite factors. At one selected pair above q_0, let Z be the preimage of the fixer of a nonzero frame vector (standard or inverse adjoint). Then $M(K_T)/M(Z) = q_0^r - 1$ by the transitive linear action on nonzero vectors. The measure of the nonunits at any product of the chosen split pairs is at most $4fb^{-10} \le b^{-9}$: the set where some coordinate is a nonunit is the union of the corresponding maximal ideals, each of measure $1/q_j$. For the single pair, use (16.21) with $t = k - 10^{10}n^4$. In the standing size range $10^{10}n^4 \le n^6$, $k \ge 1000(2n)^9$, $b^{10} \le q_0 \le b^{11}$. Thus

$$\frac{M(Z)}{M(K_T)} = \frac{1}{q_0^r - 1} \le \frac{2}{b^{10r}}, \qquad e_T^{t+1}\frac{2^{-n}}{M(K_T)} \le b^{-9(k-10^{10}n^4+1)} \le b^{-8k}.$$

Both bounds are at most $\frac{1}{2}b^{-6k}$ under the given scales. Combining the upper estimate for $\mathcal{K}(Z)$ with (16.22) for K_T proves

$$\frac{\mathcal{K}(Z)}{M(Z)} \le b^{-6k}\frac{\mathcal{K}(K_T)}{M(K_T)}.$$

The phase factors over separate pairs, since the finite conductor character and unipotent membership factor, as does the integral (16.20). Haar normalization constants cancel in this ratio inequality (the zero form is cubic in the Haar normalization, while the denominator is linear). Choosing this intermediate subgroup at each of a set Σ of fixed-vector places and the full integral subgroup elsewhere gives the factor $b^{-6k|\Sigma|}$. Thus the numerical comparison in `IsFillingLevel.exists_intermediateGroup_lines` of `New/Filling/FixedLines.lean` follows from the bottom floor estimate, with all finite vector orbit counts explicit.

Lemma 16.19 (Central average controlled by the zero form)**.** *Let T be the union of the chosen split pairs and let ϕ denote the finite phase function there. Define*

$$\overline{\phi}(s) = |\mathfrak{G}|^{-1} \sum_{t \in \mathfrak{G}} \phi(tst^{-1}).$$

Under (16.21) for $Z = K_T$, $t = 0$, and $M(K_T) \ge 2^{-n}$, one has $\sum_s |\overline{\phi}(s)|^2 \le 2$.

Proof. Restrict the nonnegative integral (16.20) to $x \in K_T$. Since the phase depends only on the residue class there, its inner integral is $M(K_T)\overline{\phi}(\operatorname{red} x)$. Each residue class has measure $M(K_T(g))$. Thus

$$M(K_T)^2 M(K_T(g)) \sum_s |\overline{\phi}(s)|^2 \le \mathcal{K}(K_T).$$

Divide the upper bound (16.21) by the factor on the left. As $e_T \le 1$ and $M(K_T) \ge 2^{-n}$ the resulting bound is at most two. □

This also gives a direct derivation of the trace estimate in `New/Filling/Moments.lean`. Let $\mathfrak{G}_i$ be the group at the i-th rational prime and $\jmath_i$ its insertion into the product group. We use the following exact residue support of the Bessel-subgroup amplitudes:

$$\nu^{-1} \sum_c \sum_j A_{S_B}(c, \jmath_i(s)^{-1}, c)_{jj} = D_c \phi(\jmath_i(s)). \tag{16.23}$$

Indeed a rational Bessel term $\gamma = uh$ with same initial and terminal class and reduction identity at all places of the *other* factor has h integral in the class lattice and congruent to identity at those places. Such a rational definite unitary h must be one. Here is an elementary verification of the freeness: if it is not one, the rational algebraic integer $a = \operatorname{tr}(h) - k$ is nonzero, because at any complex embedding a unitary matrix has trace of real part at most k, with equality only at the identity. Integrality follows by taking the trace of its integral conjugate. If it reduces to the identity at f places of norm q_{1-i}, its nonzero absolute norm is at least q_{1-i}^f; but its absolute value at every complex place is at most $2k$, so the absolute norm is at most $(2k)^{2f} < q_{1-i}^f$. The contradiction proves freeness. Hence the Bessel term is unipotent, and none exists if its reduction is not in $\bar{U}$. For a reduction in $\bar{U}$, strong approximation and the quotient covariance of Section 16.10 identify its amplitude with the unipotent translate of the identity amplitude. With the source convention $W(uy) = \chi(u)^{-1} W(y)$ for rational u, the term over the inverse residue gives phase $\phi(\jmath_i(s))$. Then the identity amplitude contributes ν on each diagonal coordinate. This proves (16.23); the sign of the phase would be inverted under the conjugate Bessel convention, with identical norm estimates.

If $z = \sum_s z_s[s]$ is a central self-adjoint idempotent of $\mathbb{C}[\mathfrak{G}_i]$, then the identity coefficient satisfies

$$z_1 = (z^* z)_1 = \sum_s |z_s|^2 \ge 0.$$

Its centrality implies that the sum $\sum_s z_s \phi(\jmath_i(s))$ does not change on replacing ϕ by its full conjugation average: conjugation by the other factor acts trivially and conjugation by its own factor permutes the coefficients of z. Thus Lemma 16.19, the identity (16.23), and Cauchy–Schwarz give

$$\left| \nu^{-1} \sum_s z_s \sum_{c,j} A_{S_B}(c, \jmath_i(s)^{-1}, c)_{jj} \right| \le D_c \sqrt{2 z_1}. \tag{16.24}$$

When $z_1 \le b^{-8}/2$ this is at most $b^{-4} D_c$, precisely `re_trace_besselS_le` in `Moments.lean`. Thus the required trace estimate is a consequence of the bottom finite local kernel estimate and the proved Bessel transport identities, not a separate use of a global trace formula.

16.16 The projection covariance and the complement-coset error

Let T_0 be the complement of the rational Bessel subgroup and let $A(s)$ denote the normalized block matrix with entries $\nu^{-1} A_{T_0}(c, s, c')$. Bessel transport under right multiplication by

$\sigma \in \bar{S}_B$ permutes its column classes and replaces each column-coordinate block by a unitary factor, as proved in Section 16.10. In other words for a unitary block permutation V_σ one has $A(s\sigma) = A(s)V_\sigma$ (or the transpose convention, if the indices are reversed). It is essential here that V_σ depends on the lift of σ and the column class, but not on s; this follows from the term bijection by the same right multiplication for every s. The complement of the rational Bessel subgroup is stable under this multiplication. It follows for every central self-adjoint projection $z \in \mathbb{C}[\mathfrak{G}_i]$, with

$$M_z(s) = \sum_t z_t A(\jmath_i(t)^{-1}s),$$

that $M_z(s\sigma) = M_z(s)V_\sigma$ and hence $\|M_z(\sigma)\|_{\mathrm{HS}} = \|M_z(1)\|_{\mathrm{HS}}$.

For clarity the set on which to use the complement energy bound must be invariant under multiplication by $\jmath_i(\mathfrak{G}_i)$. The correct set is

$$\Omega_{1-i} = \{s : \ s_{1-i} \text{ is the } (1-i)\text{-component of some } \sigma \in \bar{S}_B\}.$$

It is the set involving the *other* factor, not the i-th factor. It contains $\bar{S}_B$ and is invariant under left $\jmath_i(\mathfrak{G}_i)$ multiplication. Convolution by z is an orthogonal projection in the regular Hilbert space, also with Hilbert–Schmidt matrix-valued coefficients, so restrict the function A by the indicator of Ω_{1-i} and apply its norm contraction. This proves

$$|\bar{S}_B|\|M_z(1)\|_{\mathrm{HS}}^2 \le \sum_{s\in\Omega_{1-i}} \|M_z(s)\|_{\mathrm{HS}}^2 \le \sum_{s\in\Omega_{1-i}} \|A(s)\|_{\mathrm{HS}}^2.$$

The complement energy estimate of `energy_compl_besselS_le` bounds the last sum by $b^{-12}|\bar{S}_B|D_c$. Thus the sum of the Euclidean row norms of $M_z(1)$ is at most

$$\sqrt{D_c}\|M_z(1)\|_{\mathrm{HS}} \le b^{-6}D_c \le b^{-4}D_c.$$

This supplies the remaining covariance-convolution step of `sum_norm_amplitudeSum_le` in `Moments.lean`. The use of the opposite factor is needed for the contraction argument and is independent of the fact that the corresponding energy estimates hold symmetrically for each factor.

We can also verify the reduction of the complement energy estimate to the crude stratified comparison. In the latter write

$$D = \frac{D_c^2}{d_\infty \mathfrak{m}_r}, \qquad D_1 = |\bar{S}_B|D_c, \qquad \tau = b^{kf/4}, \qquad \alpha = (b/g)^{kf}(2\tau^2).$$

The standing size hypotheses are $\tau^{44}D_1 \le D$ and $\tau^{18}\alpha D \le D_1$. The crude energy inequality `energy_le_of_linkEnergy_le` of `CrudeEnergy.lean`, recalled in Section 15.13, bounds the energy by $(1+b^{-4})^2 w + b^{-12}C$ when the retained links have energy at most w, and the parameters satisfy the stated thresholds: $C \ge D_1$; if a level-zero normal form occurs, $C \ge \tau^{20}D_1$; and if links are omitted, $C \ge \tau^{12}D/X^k$, where X is the product of residue-field sizes at which a nonzero frame vector is fixed.

On Ω_i, every place in both halves of every split pair above q_i fixes the first frame vector, since this is true for the reduction of the Bessel group. Thus we may take $X = q_i^{2f}$. The frame of a level-zero normal form is carried onto itself by an upper triangular matrix of constant diagonal ζ with $\zeta\bar{\zeta} = 1$: this follows by assembling the shared relations for each index into two upper triangular matrices, the first of diagonal ζ and the second of diagonal one. In the marked integral position the frame is unchanged by the class basis, so ζ is integral and invertible at every finite place. Its reduction at the specified factor agrees with the identity diagonal of $\bar{S}_B$; if $\zeta \ne 1$, its difference from one has norm at least q_i^f but absolute value

at every complex place at most 2, a contradiction. Thus $\zeta = 1$; the term is in the Bessel subgroup. It follows that no level-zero normal form occurs for the complement on Ω_i.

Take all links omitted and $w = 0$, $C = D_1$. The missing-link size requirement holds because

$$\tau^{12}/X^k \le b^{3kf-20kf} = b^{-17kf}, \qquad \tau^{18}\alpha = 2b^{5kf}(b/g)^{kf} \ge 2b^{-17kf}.$$

The other requirements are immediate. Applying the crude comparison proves $\mathcal{E}(A_{T_0}; \Omega_i) \le b^{-12}D_1$, as claimed in `energy_compl_besselS_le`. This derivation makes explicit its dependence on the inverse index savings and the level-zero freeness. No coherent main-term estimate is used for this error term.

16.17 Assembly of the sharp energy bound

We conclude by tracing the quantitative thresholds and the remaining inputs in the full second moment. Let Σ be the specified fixed-vector places above q_0 and set $y_\Sigma = b^{-6k|\Sigma|}$. The set `Filling.fixingSet` consists of reductions with the specified fixers and, unless all selected places above q_0 belong to Σ, with at least one selected place at which *all* leading frame minors are units. For Σ equal to the whole set, the latter condition is vacuous. The theorem of Section 14.8 then bounds the full positive-type energy using this intermediate set; that induction does not itself require unit minors at every place simultaneously.

The normalized coherent cut energy, with links retained when Σ is proper and no links retained when it is the whole set, must satisfy

$$w \le (1 + b^{-4})y_\Sigma D. \tag{16.25}$$

The proof in the source uses the top packet comparison (`sum_blocks_le_topKernel` in `New/Filling/FiniteSystem.lean`), the global-to-local limiting identity (`exists_hasSum_coSetAmplitude` in `GlobalToLocal.lean`), the archimedean and finite local band systems, the fixed-vector bottom factor, the global sieve just proved, and the packet constant inequality of the filling datum. We have supplied the complete global frequency and Poisson budget in Section 16.14, and the elementary norm comparison and multiplicative saving for fixed-vector groups from the finite bottom estimate. The top packet comparison must still be supplied as a specialized input, including the integral-point count on the definite quadric and its normalization. In particular one cannot use the crude bound on the number of rational elements in each link as a replacement for that comparison.

Given (16.25) and the crude comparison, the energy estimate is elementary. Take $C = y_\Sigma D$. Its level-zero requirement holds since $y_\Sigma \ge b^{-6kf} = \tau^{-24}$, so $y_\Sigma D \ge \tau^{20}D_1$ by the standing size hypothesis. If links are omitted, Σ is the whole set and $X \ge q_0^f$, so

$$\tau^{12}D/X^k \le b^{-7kf}D \le y_\Sigma D.$$

For proper Σ that requirement is vacuous. Thus the crude comparison with w as in (16.25) yields

$$\mathcal{E} \le ((1 + b^{-4})^3 + b^{-12})y_\Sigma D \le (1 - b^{-4})(1 + 9b^{-4})y_\Sigma D.$$

For the last inequality set $u = b^{-4}$: subtracting the left coefficient from the right gives $5u - 12u^2 - 2u^3 \ge 0$ for the stated b. This is the assembly of `energy_fixingSet_le` in `New/Filling/Moments.lean`. Its main analytical content remains precisely the local estimates and the top packet comparison giving (16.25); its numerical conclusion requires no further input.

The central projection calculation then gives the first moment of Section 14.9: use $z = 1$ in the covariance estimate and the identity Bessel amplitude to bound the rows of $P^*P - 1$. For

each bad central-type projection use (16.24) and the covariance estimate; since its identity coefficient is at most $b^{-8}/2$, the total discarded trace is at most $2b^{-4}D_c$ for that factor. For the two commuting projections the total discarded trace is at most $4b^{-4}D_c$. The positive-type induction of Section 14.8 gives the second moment from the energy estimate just assembled. Thus the Hilbert–Schmidt variance and spread argument of Section 14.10 can be applied without additional finite-type covariance or central-trace inputs.

16.18 What remains of the filling argument

The new proofs have separated the following mathematical obligations from the global filling theorem.

- The normalized Flensted–Jensen identity, its change of splitting and triangular shifts, the actual large-gap determinant decay, the commutator-sensitive minor expansion, the quantitative Gaussian estimate, the common compact correction and the interpolation realization of the Bessel phase are now related by proofs with explicit constants. The classical representation-theoretic identification, formal degree and differentiated convolution theory are still invoked in their stated forms. The cut construction uses the peak and its quotient norm bounds in Section 15.6.
- The marked lattice and the exact class and group-order factors, unipotent coordinate Haar measure and covolume, arithmetic box construction, finite and infinite character normalization, transport between Bessel classes, and flag propagation eliminating the vanishing-discriminant terms are reconstructed here from integral block matrices, density, and elementary lattice geometry. The positive link metric and its congruence defect have also been checked.
- The lattice-height zeta sum and global frequency and head sums are reconstructed in their numerical strength, not only as convergence statements. Given the local systems, the one-sided sieve now follows with its error at most $b^{-6kf-10}/d_\infty$. The projection and complement-energy deductions from the crude comparison and bottom form have also been supplied.

Several genuinely specialized local estimates still remain. They include the construction of the floor kernels and the uniform Fourier bounds for the actual real Bessel system (`New/RealPlace/LocalTheorem.lean`, `BandLimits.lean`), the finite gauge/quotient measure comparisons and bottom form (`New/Filling/StandardLevel.lean`, `TopFloor.lean`, `FiniteSystem.lean`), the quantitative joint lattice-saturation and weighted pair bounds (`LinkPairing.lean`, `BlockBound.lean`, `CrudeEnergy.lean`), and the normalized coherent top packet comparison (`GlobalToLocal.lean`, `FiniteSystem.lean`). The class and degree size inequalities of `Filling.IsAdmissible` must be checked simultaneously for the actual weight and number field parameters using the mass formula and degree estimates. Some of these numerical identities follow directly from the class count and covolume derived here, but the full simultaneous verification is not asserted by these two identities alone. In particular (16.25) is still conditional on the named local and packet results. This distinction is necessary because a fixed-rank matrix coefficient estimate or a qualitative orbital integral convergence statement would not supply the constant-one second moment or its uniform error.

Part IV

Lower bounds: periods, congruences and crystalline defects

Chapter 17

The determinant lower bound

Chapters 18 and 19 develop the complementary frames, finite square formulas, integral gluing, Rallis unfolding, comparison operators and the bounded Hecke annihilator arguments needed below. The genuine global period identities and their specialized nonvanishing remain explicit inputs when not proved.

We turn to the lower bound which is used in conjunction with the shortness and spanning theorems. In this part there is a second pair of mechanisms, both involving a comparison with theta lifts for an auxiliary finite-order character. The first comparison converts a nearly integral family of rows into a height estimate for the ratio of two full-dimensional determinants. The second compares the unknown determinant of a chain of lifts with determinants of single lifts, and then replaces each of these unknown determinants by a product of a special-value ratio and an integral comparison determinant. It is the last step which repays the cost of the ordinary (or idle) differentiation directions in the upper bound. The existence of adequate twists, a separate arithmetic theorem, will be considered subsequently.

We write $n = \dim_{\mathbb{C}} W > 0$, and retain $B = fa(\log b)^3$, $L = B \log b$, $H = mk\mathcal{U}$. Here the cost of one direction is the particular real number

$$\mathcal{U} = \sum_{\tau \in \Phi} (J(w_0(\tau)) - A(w_{\mathrm{nom}}(\tau), \tau)). \tag{17.1}$$

In particular it is not a positive-definite quadratic form or a norm. None of the inequalities below assumes its nonnegativity. Put $B_0 = B/(mk)$. The object of this section is the following precise estimate.

Theorem 17.1 (C)**.** *Let $y_1, \dots, y_q \in W$ be independent, where $q \geq (1 - 4b^{-2})n$. Suppose that* $\mathrm{IndexLe}(y, d_{\mathrm{ind}})$ *holds and that, in the coefficient-conjugate metrics of the setup,*

$$\mathrm{covol}^2_\sigma(\Theta_1 y_1, \dots, \Theta_1 y_q) \leq e^V \quad (\sigma \in \mathrm{Aut}\,\mathbb{C}).$$

Then

$$n(H - 3L) \leq V + 2d_{\mathrm{ind}}. \tag{17.2}$$

The statement to be reconstructed is `Data.integral_forms_long` in `New/Endgame/LongForms.lean`. The precise estimates below designated *Formal-source input* are taken from named declarations of the compact development. We give ordinary mathematical proofs of the intervening algebraic and determinant implications, and isolate the statements of the more demanding automorphic and local computations. In particular the occurrence of a theorem from the formal source as an input is not an assertion that its independent verification or its translation from a published theorem has been completed.

17.1 Conventions for norms of conjugate coefficient vectors

Fix a level on a definite unitary group. In the coefficient model of algebraic modular forms, a form is a family of vectors in weight representations indexed by a finite adelic class set, with the appropriate covariance. For the Weyl coordinate representation used in the compact code, the pairing has a rational positive-definite matrix G (the matrix denoted `weylGramOver`). Write the class-sum pairing in schematic form as

$$\langle u, v\rangle = \sum_{h\in X} a_h\, \overline{u(h)}^{\,T} G v(h), \qquad a_h = |\operatorname{Aut}(h)|^{-1}.$$

Here X is a choice of class representatives at the specified level; different levels give different class sets and pairings. At the levels where the class stabilizers are trivial, $a_h = 1$. At other levels one clears the stabilizer orders by a uniform integer. The normalization concerning class sums is essential: it is not the usual probability measure on the finite class set.

For $\sigma \in \operatorname{Aut}(\mathbb{C})$ the norm used in the construction is

$$\|u\|_\sigma^2 = \sum_h a_h\, \overline{\sigma(u(h))}^{\,T} G \sigma(u(h)),$$

with σ applied entrywise. If u is a tuple of forms, the associated positive Gram matrix is formed from these conjugated vectors. It is in general incorrect to identify it with σ applied to the original Gram matrix. With κ denoting complex conjugation and $\sigma^* = \kappa\sigma\kappa$ as field automorphisms, one has instead

$$\sigma\langle u, v\rangle = \sum_h a_h\, \overline{\sigma^*(u(h))}^{\,T} G \sigma(v(h)). \tag{17.3}$$

This follows from the rationality of the coefficient pairing and from $\sigma(\bar z) = \overline{\sigma^*(z)}$. Thus

$$|\sigma\langle u, v\rangle| \le \|u\|_{\sigma^*}\|v\|_\sigma.$$

Much of the emphasis below on simultaneous estimates for *all* field conjugates is due to this distinction, not to the number of complex embeddings of a fixed coefficient field. These conventions match `LevelForm.normSqAt`, `LevelForm.pairingMatrix` and `LevelForm.covolumeSq` in `Lib/Adelic/CuspForms/LevelForms.lean`.

Lemma 17.2 (Mixed Gram determinant and completion)**.** *Let $v = (v_i)_{i=1}^q$ be a tuple in a finite coefficient representation as above. If its squared covolume is bounded by e^V at every conjugate, then*

$$|\sigma \det(\langle v_i, v_j\rangle)| \le e^V \quad (\sigma \in \operatorname{Aut}\mathbb{C}).$$

If $z_1, \dots, z_h$ satisfy $\|z_j\|_\sigma^2 \le e^A$ for every σ, the corresponding bound for the tuple (v, z) is e^{V+hA}.

Proof. Regard the scaled coefficient vectors as vectors in one ordinary finite-dimensional Hermitian space, so that the form displayed above is its Gram form. Equation (17.3) identifies the conjugate of the determinant with the scalar product of the exterior products of the σ^*-tuple and the σ-tuple. Cauchy–Schwarz in the exterior power bounds its modulus by the square root of the product of their squared covolumes, at most e^V. Equivalently, one uses

$$|\det(X^*Y)|^2 \le \det(X^*X)\det(Y^*Y)$$

for two rectangular matrices of equal column size. For the completion bound, if the conjugated v_i are dependent, the extended Gram determinant vanishes. Otherwise perform Gram–Schmidt

beginning with these q vectors. Each subsequent squared length is bounded by the original squared length of the added vector. Hence

$$\det \mathrm{Gram}_\sigma(v, z) \le \det \mathrm{Gram}_\sigma(v) \prod_{j=1}^{h} \|z_j\|_\sigma^2.$$

The source proves these assertions as `norm_map_det_gram_le` and `covolumeSq_elim_le` in the same file. □

We will also use the lattice property of the coefficient pairing. Let $\phi : \mathbb{C} \to \mathbb{C}_\ell$ be a field embedding and Λ_ϕ the lattice of forms of Weyl gauge at most one at ϕ. The coordinate lattices at the individual points can be calculated in integral adelic frames, rather than in an arbitrary fixed complex coordinate basis. Changing a rational frame is compatible with the Gram pairing. Consequently, for $u \in \Lambda_{\phi\circ\kappa}$, $v \in \Lambda_\phi$ the local pairing of the two coordinate vectors is integral. At the levels under consideration the class stabilizers are trivial and this implies

$$|\phi\langle u, v\rangle|_\ell \le 1. \tag{17.4}$$

Indeed after passing to integral frames the pairing becomes a sum of terms $\phi(\bar{u}_i)G_{ij}\phi(v_j)$ with integral entries; the integral class-sum property then follows from the ultrametric inequality. The source identifies these frames by `exists_mul_mem_integralAbove` and its Weyl gauge change law. We use their existence and compatibility as *Formal-source inputs*: `gramForm_mem_embValRing`, `norm_inner_le_one` (`Lib/Adelic/CuspForms/LevelForms.lean`), and `Data.classAut_eq_bot_of_le_chosenLevel` (`New/Data/ChosenPlaceLevels.lean`, used in `New/Endgame/LongForms.lean`). Notice that integrality at just one embedding for both arguments would not justify (17.4).

17.2 Heights with varying finite-place entries

In the direct application, different primes can require different twists. The height device allows the local entries of a family to vary, while preventing the arbitrary insertion of prime-dependent normalizing factors in the archimedean term. Let

$$I = \{\iota = (\ell, \iota_\ell) : \ell \text{ a rational prime, } \iota_\ell : \mathbb{C} \simeq \overline{\mathbb{Q}}_\ell \text{ a field isomorphism}\}.$$

The ℓ-adic absolute value is normalized by $|\ell|_\ell = \ell^{-1}$. These isomorphisms are not topological. Their use as field isomorphisms follows from the axiom of choice and cardinalities of transcendence bases: both algebraically closed fields have cardinality and transcendence degree $2^{\aleph_0}$ over a countable algebraic closure of $\mathbb{Q}$. In particular any embedding of a fixed number field in either algebraically closed field extends to the requisite field isomorphism.

Definition 17.3. For a family $u = (u_\iota)_{\iota\in I}$ of complex numbers and $c \in \mathbb{R}$ write $\mathrm{Ht}_{\le c}(u)$ if there exist $Q \ne 0$ in $\mathbb{C}$, a real number α, and a real family h_ℓ indexed by primes such that

$$\begin{aligned} |\sigma Q| &\le e^{\alpha} && (\sigma \in \mathrm{Aut}\,\mathbb{C}),\\ |\iota_\ell(Qu_\iota)|_\ell &\le e^{h_\ell} && (\iota \in I),\\ \alpha + \sum_{\ell<M} h_\ell &\le c && (M \text{ sufficiently large}). \end{aligned}$$

The last sum runs over primes. We allow the h_ℓ to be signed; this is an eventual bound for prime cutoffs, not an assertion of absolute summability. This is `HeightLe` in `New/Heights/Heights.lean`.

Thus Q is a single nonzero number, uniform in ι. Even if a height bound is negative, it need not contradict a product formula when the entries u_ι differ. The construction requires precisely such a negative bound for the inverse volume ratios, to which the common determinant calculation will attach their reciprocal.

Lemma 17.4 (Height calculus)**.** *The following assertions hold.*

1. *If* $|\iota(u'_\iota)|_\ell \le |\iota(u_\iota)|_\ell$ *for all* ι*, then* $\mathrm{Ht}_{\le c}(u)$ *implies* $\mathrm{Ht}_{\le c}(u')$*; the constant* c *can be increased.*
2. *From* $\mathrm{Ht}_{\le c}(u)$ *and* $\mathrm{Ht}_{\le c'}(u')$ *follows* $\mathrm{Ht}_{\le c+c'}(uu')$*. In particular estimates add under finite products, and* $\mathrm{Ht}_{\le jc}(u^j)$ *for* $j \in \mathbb{N}$*.*
3. *If* $|\iota(u_\iota)|_\ell \le e^{h_\ell}$ *for every* ι *over* ℓ*, and* $\sum_{\ell\in F} h_\ell \le c$ *for all finite sets of primes* F*, then* $\mathrm{Ht}_{\le c}(u)$*.*
4. *If* $\mathrm{Ht}_{\le c}(1)$ *holds for the constant family, then* $c \ge 0$*.*

Proof. The first assertion uses the same witness. The second uses the product witness QQ' with archimedean bound $\alpha + \alpha'$ and finite-place bounds $h_\ell + h'_\ell$. The eventual prime cutoff works by taking the maximum of the two cutoffs. The third uses witness 1. For the last, the boundedness of all field conjugates of Q already implies that Q is algebraic: if it were transcendental, isomorphisms defined by extension of transcendence bases could send Q to a transcendental element of arbitrarily large absolute value. Put $F_0 = \mathbb{Q}(Q)$, $s = [F_0 : \mathbb{Q}]$, and $z = N_{F_0/\mathbb{Q}} Q \in \mathbb{Q}^\times$. Taking products over all embeddings of F_0 into $\mathbb{C}$ and into $\overline{\mathbb{Q}}_\ell$ respectively gives

$$\frac{1}{s}\log|z| \le \alpha, \qquad \frac{1}{s}\log|z|_\ell \le h_\ell.$$

Choose the eventual cutoff M large enough that $|z|_\ell = 1$ for every $\ell \ge M$. The rational product formula gives

$$0 = \frac{1}{s}\left(\log|z| + \sum_{\ell<M} \log|z|_\ell\right) \le \alpha + \sum_{\ell<M} h_\ell \le c.$$

The named source lemmas are `HeightLe.mono`, `HeightLe.mul`, `HeightLe.prod`, `HeightLe.pow`, `HeightLe.of_le` and `HeightLe.nonneg`. □

A useful consequence, repeatedly used below, is that if Q_ι is a positive integer prime to ℓ then inserting an arbitrary integral power of Q_ι into u_ι does not change its height bound: its image has local absolute value one. No new archimedean witness is required. In the application Q_ι is the norm of the twist conductor, not a fixed normalization scalar.

17.3 The determinant lemma with index

Let $T_\theta : W \to V_\theta$ be injective linear maps of finite-dimensional complex spaces with the prescribed metrics, for example Θ_θ. Define the squared volume

$$Z_\theta = \det(T_\theta^* T_\theta) > 0.$$

For a basis $v = (v_i)_{i=1}^n$ of W write $G_\theta(v) = \mathrm{Gram}(T_\theta v)$. A change of basis from an orthonormal basis to v shows

$$\det G_\theta(v) = Z_\theta \det \mathrm{Gram}(v), \qquad \frac{Z_\theta}{Z_1} = \frac{\det G_\theta(v)}{\det G_1(v)}. \tag{17.5}$$

Here it is important to use the same basis of the row space for every twist. The proof of Theorem C uses the following form of the determinant argument; its applicability does not require an integral basis of all row forms at a fixed finite place independent of the twist.

Lemma 17.5 (One determinant)**.** *Suppose v is a full basis and $|\sigma \det G_1(v)| \le e^{\alpha}$ for every σ. Fix a family θ_ι of twists and a family of nonzero scalars γ_ι such that $\mathrm{Ht}_{\le g_0}((\gamma_\iota\bar\gamma_\iota)^{-1})$. Let h be a nonnegative integer. Suppose the following holds for every ι over ℓ and every embedding $\phi : \mathbb{C} \to \mathbb{C}_\ell$: a matrix M_ϕ combines $T_{\theta_\iota} v$ into coefficient-integral rows at ϕ, with*

$$|\phi(\det M_\phi)|_\ell \ge e^{-e_\ell} |\phi(\gamma_\iota)|_\ell^h, \qquad \sum_{\ell \in F} e_\ell \le d_{\mathrm{ind}}$$

for every finite prime set F. Assume that the integral lattices pair as in (17.4). Then

$$\mathrm{Ht}_{\le \alpha + 2d_{\mathrm{ind}} + hg_0}\left(\frac{Z_{\theta_\iota}}{Z_1}\right). \tag{17.6}$$

Proof. Take $\phi = \iota_\ell$ and its composition with κ; the two matrices are not required to be the same. After applying ι_ℓ, the matrix of mixed pairings of the new rows is integral by hypothesis, and in the complex model it is

$$\overline{M_{\phi\circ\kappa}} G_{\theta_\iota}(v) M_\phi^T.$$

All its entries have image of absolute value at most one, hence so does its determinant by the Leibniz formula. Evaluating the determinant product and using the two determinant lower bounds gives

$$|\iota_\ell(\det G_{\theta_\iota}(v)(\gamma_\iota\bar\gamma_\iota)^h)|_\ell \le e^{2e_\ell}.$$

In view of (17.5), using $\det G_1(v) \ne 0$ as witness shows that $(Z_{\theta_\iota}/Z_1)(\gamma_\iota\bar\gamma_\iota)^h$ has height at most $\alpha + 2d_{\mathrm{ind}}$. Multiply by the hth power of the assumed inverse-square family, and apply Lemma 17.4. This is `one_determinant` in `New/Heights/OneDeterminant.lean`. □

The coefficient 2 in front of the index is therefore forced by the Hermitian determinant and is not a generous error allowance. The source treats the general situation in which the chosen lattices vary with ι, θ_ι, and ϕ. The quantifier over *all* embeddings ϕ here is intentional. In particular, an index bound at ι_ℓ only, in the absence of one at $\iota_\ell \circ \kappa$, is insufficient to reach (17.6).

17.4 Special-value data and fit twists

It will be helpful to give the exact analytic quantities to which the determinants are compared. Let M denote the set of chosen split places above $g_0 g_1$ (not the analytic cutoff on primes), and for a twist ϑ let $Q(\vartheta)$ denote its support, disjoint from M and its conjugate in the sense of the tame level datum; write $N(Q) = \prod_{w \in Q} Nw$. A *pair* $p = (\nu, z)$ of the proof consists of a tame unitary character ν of the norm-one torus of weight $-c_\tau$, unramified away from M, and an integer z with $1 \le z \le t$ and $z \le c_\tau \le 4b$. Tame part is denoted ψ. Define

$$\rho_p(\vartheta, s) = \frac{|L^{M\cup Q}(z, (\nu\vartheta^s)^{-1})|}{|L^M(z, \nu^{-1})|}, \qquad s \in \{-1, 1\}. \tag{17.7}$$

The removed local factors are to be understood in the source convention, including the conjugate places for the Hecke L-functions on the norm-one torus. The lower point p_v is the same pair with z replaced by v, for $1 \le v \le z$. For an integer $e \ge 0$ put

$$\Pi_S(\eta, e) = \left(\prod_{w \in S} Nw\right)^{e^2} \prod_{v=1}^{e} |L^S(v, \eta)|^2.$$

This is precisely `lValueProd` of `Lib/SiegelWeil/Tower/LargeSideConstantTerm.lean`; the ratio of two products at e satisfies

$$\frac{\Pi_{M\cup Q}((\nu\vartheta^s)^{-1},z)}{\Pi_M(\nu^{-1},z)} = N(Q)^{z^2}\prod_{v=1}^{z}\rho_{p_v}(\vartheta,s)^2. \tag{17.8}$$

Indeed each factor agrees and the first prefactor is a ratio of products of residue cardinalities over disjoint sets. The definitions are `Data.Pair.ratio`, `Pair.lowerPoint` and the identity `Pair.lValueProd_eq` in `New/Exchange/LValueRatio.lean` and `New/Exchange/TriangleLevelExponent.lean`.

The needed nonvanishing of these special values uses the weights of the pairs and their particular Hecke characters. We do not use GRH for them. As a *Formal-source input*, the declaration `ratio_ne_zero_and_real` of `New/Exchange/TriangleLevelExponent.lean` gives that every ratio here is nonzero and real; this follows in the code from `AutChar.HasWeight.L_ne_zero`. This property allows reciprocal families of ratios to enter the height calculus without exceptional zero entries.

The local projection condition. For an auxiliary pair $O=(W_j,V_j)$ with $\dim_E W_j = r$, let V_O denote its algebraic form space on the definite group. At each split place w outside $M\cup Q$ and its conjugate, there is a normalized Hecke polynomial operator $H_w(X)$ on V_O. The test point is

$$X_w = (a_O(w)\,Nw^{\mathrm{rankDiff}(O)-1})^{-1},$$

where $a_O(w)$ denotes the Frobenius character value in the normalization of `AuxPair.aFrobValue` for rank r. We use the precise operator in `AuxPair.heckePolyOnForms`. The definition of X_w is `AuxPair.testValue` in `New/Rows/StackedDeterminants.lean`. If e_0 denotes the chosen central projector and $Z\subset V_O$ the space of lifts, set

$$F_w = \operatorname{im} e_0 \cap \ker H_w(X_w)\cap Z^{\perp}.$$

The points X_w are designed so that genuine lifts belong to the first two factors. A vector is *tested* if it is orthogonal to F_w at every such place; the false-lift subspaces depend on w. This is `AuxPair.IsTested` in `New/Selmer/FitTwists.lean`. The orthogonal complement $Z^{\perp}$ is computed with the ordinary complex pairing.

For a fixed isomorphism ι_ℓ, Z has denominator bound β when the following holds for $d_Z=\dim Z$: for every $u_1,\dots,u_{d_Z}$ integral at every ℓ-adic embedding, and every $v_1,\dots,v_{d_Z}$ integral at ι_ℓ and tested,

$$\left|\iota_\ell \det(\langle u_i, p_Z v_j\rangle)\right|_\ell \le e^{d_Z r\beta}. \tag{17.9}$$

Here p_Z is the complex orthogonal projection. This is the actual declaration `DenominatorLe` in the same file. The projection is only applied to the second family; equality with the pairing of both projections may be used at the ordinary complex metric. In this definition it would be substantially stronger to ask for an entrywise denominator bound for every integral vector, or to drop the word tested. Neither strengthening is assumed.

An m-fold twist θ has first m powers still nontrivial at every prime of its support, each power with the same level support. The family θ_ι is fit if:

1. $\ell\nmid N(Q(\theta_\iota))$ for ι over ℓ;
2. for every $j\in[1,m]$ there exist numbers β_ℓ (allowed to depend on j) such that $\sum_{\ell\in F}\beta_\ell\le B_0$ for all finite prime sets F, and the good single-lift subspace at twist θ_ι^j has denominator bound (17.9); the same β_ℓ bounds the rank-one lift lines of all pairs of rank-one type at this index. For the latter the rank in (17.9) is one;

3. for all j, both signs s, and all pairs p ramified exactly at M,

$$\mathrm{Ht}_{\le x^{999}}\left(\rho_p(\theta_\iota^j, s)^{-1}\right).$$

Here rank-one type means $m < z \le r$, tame part ω^j, and $w_j + r - t \le c \le w_j$ at all infinite places (with the rank index condition in `Data.Pair.IsRankOne`); note that the term rank-one refers to a lift *from* $U(1)$ and not the dimension of its algebraic coefficient representation. At a fixed j the same finite-place costs must handle all these auxiliary pairs. These conditions are precisely `Data.IsFitFamily` in `New/Selmer/FitTwists.lean`. The existence result `exists_isFitFamily` in `New/Selmer/FitTwistsExist.lean` is a *Formal-source input* here; its substantial arguments are developed separately later in the paper.

Definition 17.6 (Cost of an exchange factor)**.** For a complex scalar $c(\theta)$ depending on the m-fold twist, say it has cost s squares if for every fit family one has

$$\mathrm{Ht}_{\le 2stx^{999}}\left(|c(\theta_\iota)/c(1)|^{-2}\right).$$

This description is used below only for nonzero factors. It is `Data.CostsSquares` in `New/Exchange/CostOfSquares.lean`.

Lemma 17.7. *Costs add under multiplication (and may be increased). A factor whose squared modulus equals a fixed positive number times an integral power of $|\mathcal{O}_E/\mathfrak{m}(\theta)|$, for $\mathfrak{m} = M^a Q(\theta)$ with fixed a, costs zero squares. If a factor c satisfies*

$$|c(\theta)|^2 = C\Pi_{M\cup Q(\theta)}((\nu\theta^j)^{-1}, z),$$

with $C > 0$ independent of θ, and $p = (\nu, z)$ ramified exactly at M, then c costs one square.

Proof. The first assertion is Lemma 17.4. The cardinality of the residue ring is a fixed integer times $N(Q(\theta))$, so in the second assertion the squared modulus ratio is a local unit at the relevant ℓ; the constant height witness 1 suffices. For the last assertion apply (17.8) and the fit-family condition to each of the z inverse ratios, twice. The factor $N(Q)^{-z^2}$ is a local unit and does not alter the bound. Since $z \le t$, the resulting constant is at most $2tx^{999}$. These are the rules expressed in `CostsSquares.mul`, `CostsSquares.of_norm_sq_eq_mul_card_zpow` and `exists_isTameFor_costsSquares_one`. □

17.5 The correspondence of complementary rank and the finite square

There are two uses of an exchange identity: finding a spanning family of crude rows, and factoring the volume on the good row space. We give their common input in some detail. For Hermitian spaces V, V', A let $I_A^V(\xi, \Phi)$ be the theta integral over $[U(V)]$, with real datum ξ and finite datum given by Φ on the matrix residue ring. The source makes this precise as a mixed real–finite datum `modDatum`, with splitting characters on the two members having product the prescribed rank twist. The input character ν has weights μ_τ and tame finite character χ modulo a level ideal $\mathfrak{m}$. The normalized character on the large space of rank N is denoted $\mathrm{rankTwist}(\nu^{-1}, N)$; it includes a power $\lfloor N/2 \rfloor$ of the fixed splitting character. We do not suppress this normalization in comparing the special values.

Proposition 17.8 (Complementary-rank exchange; formal-source input)**.** *Let $N = \dim_E A = \dim_E V + \dim_E V'$, $\dim_E V = \dim_E V' + 2e$, $e \ge 0$. At the infinite places the negative signs of*

V, V' add up to those of A; write this number as q_A. Assume $2q_A + 2 \le N$ and $N \le 2\mu_\tau + 2q_A$ at every infinite place. Fix the index identification of the complementary blocks and a rapid real datum v. Then there are a rapid real datum x and a constant C independent of the varying modulus and character such that, for every admissible tame character $(\mathfrak{m}, \nu, \chi)$, there is $c \neq 0$ with

$$|c|^2 = C\Pi_{\mathfrak{m}}(\nu, e), \qquad I_A^V(x, \Phi) = cI_A^{V'}(v, \Psi)$$

whenever the finite data Φ, Ψ are partners. Here $\Pi_{\mathfrak{m}}$ means the product defined above with the places of $\mathfrak{m}$.

This is `exchange_theorem` in `New/Exchange/ExchangeTheorem.lean`. Its proof in the formal development reduces the nonzero constant at equal rank to the calculation `exists_equalRankThetaRatio_value` in `New/Exchange/EqualRankValue.lean`, and then applies `exchange_of_partners` in the same file. Along a hyperbolic plane the surviving constant term of the lifted character has a scalar given by local Gauss sums and a Tate integral; the size of the total scalar under repeated reduction is precisely the product Π. Its factors have the conductor exponent e^2 since adjoining successive isotropic lines adds exponents $1, 3, \dots, 2e-1$. The actual matching of these integrals and nonconstant terms is not established by this counting identity. The tower and constant-term results in `Lib/SiegelWeil/Tower/LargeSideConstantTerm.lean`, `New/Exchange/ExchangeDifference.lean` and `New/Exchange/EqualRankValue.lean` are essential background to the displayed theorem. We take the complementary-rank theorem with its stated hypotheses as an input, rather than simply appeal to a general Siegel–Weil identity without verifying normalization or ramified data.

The passage from such an equality of integrals to an exchange of operators is important. With two exterior Hermitian spaces U_1, U_2, combine them to play the role of A. The two intermediate spaces have ranks differing by $2e$. Suppose the characters on all four edges are matched so that their products are the rank twist on $U_1 \oplus U_2$. Opposite choices of signs are used in the descent edges. Pairing the composition of a lift and a descent with a functional on the exterior target gives an integral on the product of the exterior groups, with a theta integral over the intermediate group. This is the situation of Proposition 17.8. Inserting the complementary-rank equality changes the order to descent then lift through the other intermediate group.

The source takes special care about integrability in this argument, because some of the involved groups are indefinite. Its lemma `exists_integrable_kernel_bound` bounds, for rapid vectors on the exterior groups, the integral of their products against a theta kernel by a constant times a weighted norm of the real datum. The permitted real datum on the other side is obtained as a limit in that norm and approximated by finite sums of elementary tensors. The closedness of a one-dimensional span in the resulting space of periods then gives a common scalar. The approximating datum is not necessarily a single product of two edge vectors. This is why the resulting scalar independent of the twist is not asserted nonzero in the general square law. The exact square assertion is `one_square` in `New/Exchange/OneSquare.lean`; we use it, including this potentially vanishing scalar, as a *Formal-source input.*

Here is the finite-ring part of this assertion. Let $R = \mathcal{O}_E/\mathfrak{m}$. For a prescribed pair of rectangular matrix index sets and an identification σ_0 of complementary dimensions, partners are the finite data in `IsPartner`; they are defined with the character twist and complementary full-rank blocks. The function of the edge datum on one side of a square is $F(X)$, and on the opposite side $F'(Y)$ (these are complex numbers after applying a specified functional). In the

source a square law of scalar c_0 consists of the following three conditions:

$$\begin{aligned} F(kX) &= \chi(\det k)F(X), \\ F'(Yl) &= \chi(\det l)F'(Y), \\ \sum_Y \Psi(Y)F'(Y) &= c_0 \sum_X \Phi(X)F(X) \quad ((\Phi,\Psi) \text{ partners}). \end{aligned}$$

Here k, l are invertible matrices of the appropriate sizes. In the application $c_0 = \kappa_0/c(\theta)$ with $c(\theta) \neq 0$ and κ_0 independent of θ. This definition is `SquareLaw` in `New/Exchange/SquareLaw Fourier.lean`.

The actual periods to be moved use point data and inclusion matrices. The square law is applied not to a single arbitrary point on one side and a single arbitrary point on the other, but to appropriate averages of points. For example the finite transformation called `nextDatum` sends a matrix function p to the function which, on block matrices

$$Y = \begin{pmatrix} y & * \\ 0 & I \end{pmatrix},$$

takes value $p(-y)/|\operatorname{Mat}_{n'\times b'}(R)|$, and is zero on other matrices; here n', b' are just the relevant block sizes. For this particular change the identity `SquareLaw.sum_eq_nextDatum` reads

$$|\operatorname{GL}_{a'}(R)| \sum_y p(y) F'\begin{pmatrix} y \\ I_0 \end{pmatrix} = c_0 |\operatorname{GL}_{n'+b'}(R)| \sum_x p^+(x) F\begin{pmatrix} I_0 & x \end{pmatrix}.$$

The inclusion matrices I_0 have the sizes prescribed by the block decomposition. This is a *Formal-source input* from `New/Exchange/SquareLawFourier.lean`. Its shape reflects invariance under elementary block changes and averaging over the unconstrained $*$ block. Crucially, a determinant character is evaluated only on invertible blocks; their values are roots of unity.

To make the transport of a finite family of data precise, consider a grid of spaces and put one residue matrix on each oriented edge $e = (P, P')$. If X_e is the matrix on that edge, a family of vertex matrices k_P acts by

$$X_e \longmapsto k_{P'}^{-1} X_e k_P.$$

A functional Λ on tuples of residue matrices is called tidy of order v on a set of edges $\mathcal{E}$ if it satisfies character covariance for admissible families k of vertex matrices,

$$\Lambda(f \circ k) = \left(\prod_{e\in\mathcal{E}} \chi_e\big(\det k_{P'} / \det k_P\big) \right) \Lambda(f),$$

and if it is a sum of at most Q_*^v point evaluations with nonzero coefficients roots of unity, supported on tuples whose matrices on the distinguished ramified edges have left inverses. Here

$$Q_* = |R|^{(2n_m)^2}, \qquad n_m = k + m(t+1)$$

is the bound for the number of matrices on any one edge. The precise admissible vertex families are those satisfying the endpoint level conditions and the prescribed diagonal block relation between the two grids. This is `IsTidy`, whose short-sum component is `IsShortSum`, in `New/Exchange/GridMove.lean`. The distinction between the full tidy datum and a merely level-invariant datum matters at the next square.

One elementary observation in the formal source is that orders add: if, for each point S, its evaluation is replaced by at most Q_*^w evaluations with root-of-unity coefficients supported

on valid new tuples, then a tidy short-sum of order v becomes a short-sum of order $v+w$, provided covariance is checked for the new tuple action. Indeed expand the original sum and then the replacement sum. The number of terms is at most $Q_*^v Q_*^w$, and all nonzero new coefficients are products of the permitted roots of unity. The covariance must be verified separately; it does not follow from this count. The counting statement is `IsShortSum.bind` in the same file.

The actual replacement at a square is `IsTidy.exists_square_move` (`New/Exchange/GridMove.lean`), a *Formal-source input*. Its hypotheses are the square law, the matching identifications on the edge blocks, and the compatibility of an admissible vertex family with those block embeddings. Its proof uses the finite kernel δ of `graphKernel` in `New/Exchange/SquareLawFourier.lean`. For a matrix $k \in \mathrm{GL}_N(R)$ and a decomposition $N = a' + b'$, write x_k for the last a' rows of k^{-1} and y_k for the first b' columns of k. The kernel supported on such pairs is

$$\delta(x_k, y_k) = \frac{\chi(\det k)}{|\,\mathrm{GL}_{b'}(R)|}.$$

It has the following useful property: if $\Psi(yl) = \chi(\det l)^{-1}\Psi(y)$ for $l \in \mathrm{GL}_{b'}(R)$, then

$$\sum_y \delta(x_k, y)\Psi(y) = \chi(\det k)\Psi(y_k).$$

In fact for fixed x_k its complementary columns span $\ker x_k$, a free direct summand, and its full-rank column frames are $y_k l$ as l runs through $\mathrm{GL}_{b'}(R)$. Replacing k by $k\,\mathrm{diag}(l, I)$ keeps x_k fixed, transforms y_k by l, and multiplies the determinant character by $\chi(\det l)$. The two character factors cancel term by term, proving the formula. Changes in the other complementary block that preserve both data change k by a block unipotent, so the definition is unambiguous. This is `sum_graphKernel_mul` in the source. It explains the cancellation of the averaging denominator: without the covariance, individual coefficients $1/|\,\mathrm{GL}_{b'}(R)|$ would eventually destroy the lattice estimate. The square move uses further block support and equivariance identities to normalize the graph kernel so that nonzero coefficients are roots of unity. This last assertion is included in the formal-source square-move input.

17.6 The two-grid exchange and the crude rows

We now describe the endpoint transport used for the first of the two height estimates. The grid has two families of vertices; in the notation of the source a vertex is (z, i, j) with $z = 0, 1$ and $0 \le i, j \le m$. Their ranks are

$$N(z, i, j) = r + z + i(t + z) + j(t - z). \tag{17.10}$$

The grid with $z = 0$ has equal rank increments t on both axes, whereas in the second grid the increments are $t + 1$ and $t - 1$. Thus the two diagonal vertices $(0, j, j)$ and $(1, j, j)$ have ranks differing by one. The sign count at offset zero is $\lfloor (m - j + z)/2 \rfloor$; in particular the endpoint $j = m$ is definite in both grids. The initial and final diagonal pairs differ by a line of sign $(-1)^j$. At the real places, the Harish–Chandra parameters are the initial r-block together with blocks built from the weights w_l of the construction and their negatives. The spacing $h = 2r^8$ ensures the required gap conditions and the terminal parameters interlace. These rank, sign, gap and interlacing assertions are the *Formal-source inputs* `card_gridIndex`, `card_neg_gridSigns`, `gridDiagEquiv_spec`, `hasGap_gridParam` and `gridParam_interlace` in `New/Exchange/GridParameters.lean`; (17.10) itself follows by counting the disjoint blocks.

In grid z, the path of stage j starts at $(z, 0, m)$, descends through the vertical edges to $(z, 0, j)$, then lifts along horizontal edges to (z, j, j). The union of these two paths is the path of stage j. A pair of forms on the two starting definite groups and a tidy datum on the path gives a period at the two diagonal endpoints: one of the forms is restricted through the codimension-one Hermitian embedding, with the real translation and line multiplier specified by the grid, and then multiplied by the other and integrated.

The reason for using the second grid is the following separation result.

Proposition 17.9 (Starting periods; formal-source input)**.** *There exist tidy data Λ_θ of order $4m$ on the stage-zero path, a fixed real translation, a scalar α, and a linear surjection R_0 from the specified endpoint form space on the second grid onto the full initial principal-level row form space, such that, writing ϕ for a form in the domain and Ξ for a form used to test the first endpoint,*

$$[K_f : J_m(\theta)] \int_{[U(Y_m)]} \Xi(u)\ \mathrm{lift}_\theta(R_0\phi)(u)\, du = \alpha\ \mathrm{gridPeriod}_0(\Lambda_\theta; \Xi, \phi).$$

Here $J_m(\theta)$ in the index is the level `Data.levelY` *at m as in the period statement; the integral uses exactly the source's bilinear, not Hermitian, period normalization. The map R_0 and real translation are independent of θ.*

The declaration is `Data.exists_isTidy_integral_eq_gridPeriod` (`New/Exchange/CrudeRows.lean`). The surjectivity input in its proof is stronger than existence of nonzero periods for some forms: by `Data.eq_zero_of_orthogonal_gridLevelForms` in `New/Exchange/RowsReached.lean`, there is one real translation for which a form orthogonal to all the restricted forms on the second grid is zero. It employs the real gap construction `Data.exists_isSteep` and the restriction theorem `eq_zero_of_petersson_restrict_eq_zero`. The descent from the definite endpoint onto that grid's level forms is onto by `Data.gridLevelForms_eq_map_iotaPlus`; one uses injectivity of the adjoint weighted lift chain to prove surjectivity. Thus a fixed linear surjection, rather than a different map for each row, is available.

For $j < m$, transporting the two paths from stage j to stage $j + 1$ requires commuting a descent past j lifts in each grid and moving across the remaining diagonal configurations. The source packages this precisely in `Data.hasRoute_succ` in `New/Exchange/CrudeRows.lean`, a *Formal-source input*: the cost and increase in the tidy order are at most $4j + 2$, and the initial period equals a scalar independent of the twist times the inverse of a nonzero factor of that cost times the new period (in the direction fixed by the source's `HasRoute`). Its proof uses the row ladder `Data.descent_past_gridPathMap` in `New/Exchange/RowLadder.lean` and the elementary square move. Iterating the statement is elementary: the sums of costs and of orders are bounded by

$$\sum_{j=0}^{m-1} (4j + 2) = 2m^2.$$

Products of scalar factors respect the bound by Lemma 17.7. Therefore the start period can be transferred to the terminal stage with one factor $c(\theta)$ costing $2m^2$ squares, and with tidy endpoint data of order $4m + 2m^2$. This is `Data.gridPeriod_climb`; in particular the switch to the terminal endpoint is not achieved with exponentially many independent normalization factors.

At the terminal diagonal pair (z, m, m) both unitary groups are definite. The interlacing noted after (17.10) gives a bilinear branching pairing of their Weyl coefficient spaces, represented by an integer matrix. The terminal composite consists of two definite chains, each beginning at $(z, 0, m)$ and lifting to (z, m, m). Let $\mathrm{endPeriod}(\Lambda, y)$ be the resulting

class-integral pairing, linear in the tidy datum and bilinear in the endpoint forms $y = (y_0, y_1)$, as defined in `Data.endPeriod` (`New/Exchange/GridMove.lean`). The terminal levels are chosen so that the coefficient integration can be computed as a finite class sum with the branching matrix. For later use the following are the precise uniform estimates.

Proposition 17.10 (Terminal periods; formal-source input)**.** *There is a positive integer D_*, independent of the twist, with*

$$D_* \le \exp(x^{99}b^2)$$

such that $D_*\,\mathrm{endPeriod}(\Lambda, y)$ is integral at every $\ell \nmid N(Q)$ whenever Λ is tidy on the terminal path and both endpoint forms have gauge at most one and are fixed by the terminal endpoint levels. For the trivial twist and tidy data of order $4m+2m^2$, if every coefficient-conjugate value of the two input forms has squared norm at most S_0, S_1, respectively, then at any $\sigma \in \mathrm{Aut}\,\mathbb{C}$,

$$|\sigma(\mathrm{endPeriod}(\Lambda, y))| \le \exp(x^{200}b^2)\sqrt{S_0 S_1}.$$

These are `Data.exists_mul_endPeriod_mem_embValRing` and `Data.endPeriod_le` in `New/Exchange/CrudeRows.lean`. The first clears the denominators of the successive theta kernels with the standard vacuum integrality statements, and the final branching factor with an integer involving powers of the level and N_1. The bound `boxes_endWeightData_le` bounds the algebraic weights and their box counts. For the second estimate the same finite branch weights bound the size of all derivatives and the branching pairing; the class numbers at each endpoint are bounded using `card_stdClasses_le_exp_mul_log`, and the tidy-data bound counts the residue terms. Although the trivial twist is used for this archimedean estimate, every nontrivial twist is covered by the finite-place estimate away from its own conductor. Thus one does not need an archimedean norm bound growing uniformly in all the permissible twist conductors.

We record the result of the terminal-period calculation with its precise quantifiers.

Proposition 17.11 (Crude rows; formal-source input)**.** *There is a nonzero function $c(\theta)$ of the twists costing $2m^2$ squares, and a spanning set $\mathcal{C}_0$ in the full principal-level row space (before good projection), such that for $y \in \mathcal{C}_0$:*

1. *$(c(\theta)/c(1))\,\mathrm{lift}_\theta(y)$ is coefficient-integral at every embedding over every $\ell \nmid N(Q(\theta))$;*
2. *every coefficient-conjugate value of the untwisted lift has squared Weyl norm $\le \exp(x_1^2 b^2)$.*

The declaration is `Data.crude_rows` (`New/Exchange/CrudeRows.lean`). Its construction combines Proposition 17.9 and the route calculation with Proposition 17.10, choosing spanning short integral forms at the second definite endpoint. For the first definite endpoint use delta forms as coefficient tests: if the pairing with every integral coefficient test of a vector lies in the valuation ring, the dual Weyl lattice criterion gives its integrality. The branching matrix and denominators in the dual pairing are accounted for in the common rescaling of the row forms. The source obtains a common normalization independent of θ besides the ratio $c(\theta)/c(1)$; this uniformity is part of the proposition and indispensable for the one-determinant argument.

To apply the proposition, project the spanning set to good rows with the integral normalized projector. Its integrality assertion is `goodProjTranslate_mem_integralForms` and its metric bound is `normSqAt_goodProjTranslate_le` in `New/Rows/MonomialCoefficientVector.lean` (both are used in `New/Endgame/LongForms.lean`). These are *Formal-source inputs* concerning the given finite levels. Select $n - q$ projected members extending $y_1, \dots, y_q$ to a basis; this is possible by choosing a basis of the quotient of W by their span from the projected spanning set. Their integral normalization is $\gamma_\iota = c(\theta_\iota)/c(1)$. Its inverse squared norm admits height bound

$$g_0 = 4m^2 t x^{999}.$$

Complete the local recombining matrix for the original q rows by a diagonal block $\gamma_\iota I_{n-q}$ to make the new basis integral for the same twist. The determinant condition of Lemma 17.5 holds, with $h = n - q$.

The initial archimedean upper bound for the extended full determinant is of the form $V + (n-q)A_*$, obtained from Lemma 17.2 after conversion of the crude value bound to the class-sum norm. The exact size estimate, which is a calculation with the class-count bound, the normalized projector and the admissible magnitudes, gives

$$A_* + 4m^2 t x^{999} \le \tfrac{1}{4} b^3.$$

This is the final per-vector numerical input in `Data.heightLe_volume_rows` (`New/Endgame/LongForms.lean`); it is a *Formal-source input* as regards the bounds for the coefficient models. Its scale is natural: the leading crude exponent $x_1^2 b^2$ is $o(b^3)$ since $x_1^{10^6} \le b$, and the remaining polynomial losses and class logarithms are much smaller than b^3. The final step does not require these terms separately to be smaller than the sharp error L.

Write from now on $Z(\theta) = \det(\Theta_\theta^* \Theta_\theta)$; injectivity, used for the determinant denominator, is proved below. Applying Lemma 17.5 and using $n - q \le 4n/b^2$ gives

Proposition 17.12 (Forward ratio)**.** *For the indicated family of short rows and fit family θ_ι, one has*

$$\mathrm{Ht}_{\le V + 2d_{\mathrm{ind}} + nb}\left(\frac{Z(\theta_\iota)}{Z(1)}\right). \tag{17.11}$$

Proof. The preceding determinant calculation gives a bound $V + 2d_{\mathrm{ind}} + (n-q)(A_* + g_0)$. Bound its last term by $(4n/b^2)(b^3/4) = nb$. This is the implication from the precise crude-row theorem to the forward-ratio estimate. □

17.7 The triangle of weighted lifts

The chain determinant is not a multiplicative determinant of ordinary square matrices: the Hermitian spaces increase in dimension and the lift maps take their values in proper subspaces. We need an exchange computation to make the determinant on the initial row space factor. The formal source uses an abstract triangle which we now describe. Fix a twist and let $F_{i,v}$ be spaces on the triangular array $i + v \le m$. Use arrows

$$a_{i,v} : F_{i,v} \longrightarrow F_{i,v+1}, \qquad d_{i,v} : F_{i,v+1} \longrightarrow F_{i+1,v}.$$

The $v = 0$ spaces contain isometric copies $j_i(U)$ of a common Hilbert space U of forms. Let $w_v > 0$ be the weights (integral finite group orders in the application). Define

$$A_{i,v} = (w_{v-1} a_{i,v-1}) \cdots (w_0 a_{i,0}), \qquad A_{i,0} = 1.$$

The first-arrow adjoint law, on $j_i(U)$ and all target vectors z, is

$$\langle z, a_{i,0} j_i(u) \rangle = \langle d_{i,0} z, j_{i+1}(u) \rangle.$$

Moreover $d_{i,0} a_{i,0} j_i(u) = j_{i+1}(B_i u)$ for an injective $B_i : U \to U$. For the other squares the hypotheses contain a scalar $\beta_{i,v}$ with

$$w_v a_{i+1,v} d_{i,v} = \overline{\beta_{i,v}}^{-1} w_{v+2} d_{i,v+1} a_{i,v+1}. \tag{17.12}$$

The square scalar is potentially zero in the initial exchange statement. For norms, the source has an additional adjoint version of the square, including an adjoint vector for a horizontal

arrow, and a factor of cost zero. All these conditions together are precisely `IsTriangle` in `New/Exchange/TriangleNorms.lean`; (17.12) is interpreted with the source's zero-inverse convention when necessary. The induction below on injectivity does not divide an unknown zero scalar.

The weighted chain satisfies the commuted-descent identity

$$A_{i+1,v}d_{i,0}a_{i,0} = \frac{w_{v+1}}{w_1 w_0 \overline{\beta}}\, d_{i,v}A_{i,v+1} \tag{17.13}$$

for a product of square factors (with the normalizations of the source). Here at $v = 0$ the factor is $1/w_0$. This follows by applying (17.12) as one moves the descent past each arrow in order; the finite-rank weight ratios telescope. Successive adjoint laws in the norm case give the compatible identity

$$\langle z, a_{i,v}A_{i,v}j_i u\rangle = \kappa\langle d_{i,v}z, A_{i+1,v}j_{i+1}u\rangle \tag{17.14}$$

with the product $\kappa\beta$ costing at most v squares. In the general injectivity case only (17.13) is used. These are the two assertions of `IsTriangle.weightedChain_descent`. The source proves their simultaneous induction beginning with the first-arrow adjoint law; each adjoint square supplies one factor, and the extra zero-cost factor leaves the prescribed budget unchanged. We take the full normalized adjoint identities as *Formal-source inputs* from that declaration.

Lemma 17.13 (Injectivity of the weighted chain)**.** *For $i + v \le m$, the restriction of $A_{i,v}$ to $j_i(U)$ is injective.*

Proof. Use induction on v. This is clear for $v = 0$. If $A_{i,v+1}j_i u = 0$, apply (17.13): $A_{i+1,v}d_{i,0}a_{i,0}j_i u = 0$. Since its argument is $j_{i+1}(B_i u)$, the induction hypothesis and injectivity of B_i give $u = 0$. This proves `IsTriangle.eq_zero_of_weightedChain`. □

Lemma 17.14 (Gram identity of the triangular array)**.** *Suppose a subspace $W \subset U$ is preserved by the first-arrow Gram operators S_i characterized on W by*

$$w_0\langle a_{i,0}j_i u, a_{i,0}j_i v\rangle = \langle u, S_i v\rangle.$$

In the adjoint case of `IsTriangle`, for $i + v = m$ there is a scalar η of cost at most v^2 such that the Gram form of $A_{i,v}j_i$ on W equals

$$\eta\frac{w_0 w_1^v}{w_v}\langle u, T_i v'\rangle, \qquad T_i = S_i S_{i+1}\cdots S_{m-1},$$

where v' is the second vector, and the product is the identity for $v = 0$. In particular $\det T_i = \prod_{h=i}^{m-1} \det S_h$.

Proof. The first-arrow adjoint law implies $w_0 B_i = S_i$ on W when paired against vectors of W (equivalently one uses its preserving operator on the full common space). For the inductive step take the Hermitian pairing of the last arrows, use (17.14) on one side and (17.13) on the other. The factors of w_v, w_0, w_1 combine to give $w_0 w_1^{v+1}/w_{v+1}$, with scalar multiplier contributed by $\kappa\beta$ at cost at most v. Apply the induction hypothesis in row $i + 1$ to S_i and the remaining chain, using self-adjointness of its Gram form to put S_i in the first argument of the endomorphism product. The determinant is independent of the resulting product order, so even without commutativity the conclusion about determinants holds. Starting with cost zero at length zero and increasing the bound by no more than $(v + 1)^2 - v^2$ proves the stated cost. This is the induction in `IsTriangle.inner_weightedChain_eq`; the adjoint normalizations and the exact scalar in this calculation are the ones in (17.13)–(17.14). □

For the application one uses the triangle on the left vertical strip of the preceding grids, with a varying offset i. At offset i and stage v the space of forms is that of the real parameter `gridParam` and signature `gridSigns`; the initial forms are identified by negating signs and using the isometry `formNegEquiv`. The weights are $w_v(\theta) = \mathrm{triangleOrder}(\theta, v)$, computed by the order of the finite group of the relevant modulus. The prescribed real vectors harmonics and local inclusion data define the arrows `triLift`. The first-arrow adjoint law follows from the usual pairing of lift and descent; the finite normalizing order is important for the square identity. The first-arrow injectivity uses the Rallis-type norm calculation of `hasSum_inner_triLift`, with the chosen large rank. The theorem `exists_isTriangle` in `New/Exchange/TriangleArrows.lean` verifying these requirements is a *Formal-source input*. The additional adjoint condition for the Gram identity holds in the $z = 0$ grid.

To relate these operators to the single lifts we need the Hecke transfer between groups with different signatures. For $1 \le j \le m$, let U_j denote the good form space on the auxiliary group of rank r with $m - j$ negative signs, and let $V_j(\vartheta)$ denote the definite algebraic form space of the auxiliary pair of ranks $(r, r + t)$. Write

$$u_{j,\vartheta} : U_j \to V_j(\vartheta)$$

for the good lift `Data.goodLift`, which is the adjoint of a normalized theta descent. We reserve the letter W for the good rows on Y_0, not for the Hermitian space of the same symbol in the code. The following precise transfer statements are used.

Proposition 17.15 (Rallis pairing and transfer; formal-source inputs)**.** *For every j, the map $u_{j,\vartheta}$ is injective for every allowed twist ϑ, and $\dim_{\mathbb{C}} U_j = n$. The pairings defining the squared norms of the first arrow in row $j - 1$ of the triangle and the auxiliary lift u_{j,θ^j} are, up to the explicit constants and finite group orders, sums over the same family of Hecke terms. In particular they are represented by one endomorphism in the common off-level Hecke algebra acting on the good form spaces. For every element P of that algebra, the characteristic polynomials of its actions on W and U_j agree.*

The first-arrow formula is `hasSum_inner_triLift` (`New/Exchange/TriangleArrows.lean`). The auxiliary-lift formula is `Data.auxLift_injective_and_hasSum_inner` (`New/Exchange/CharacterLiftNorm.lean`): there is $A_j \neq 0$ independent of ϑ with

$$\langle u_{j,\vartheta} x, u_{j,\vartheta} y \rangle = A_j \delta_{j,\vartheta} \sum_{\gamma} R_{j,\vartheta}(x, y; \gamma).$$

Here $\delta_{j,\vartheta}$ is `AuxPair.gramFactor`, and the terms `Data.rallisTerm` are explicitly the Hecke coset matrix coefficients, with their character factors and divided by the number of $(r + t) \times r$ matrices over the level residue ring. The Rallis identity has a stable-range hypothesis and uses a chosen sharp local datum; the usual abstract nonvanishing of a theta lift alone would give neither this sum nor its transformation under level. We rely on the source's precise identity.

The fact that the series of Hecke terms represents one element simultaneously on the two spaces is `exists_common_hecke_rallis`, in `New/Trace/GoodFormTransfer.lean`. This follows from an elementary finite-dimensional closure step once the local coset expression is known: the images of the common algebra in the product of the two endomorphism spaces, and their corresponding pairings, are finite dimensional and closed. Each common coset term is in this image after averaging over the principal-level subgroup. The convergent sum stays in the image. It need not define a convergent sum in an infinite-dimensional topology on the Hecke algebra itself.

For the characteristic polynomials the essential theorem is `trace_heckeOnGood_eq` in the same file. Its proof isolates each occurring eigensystem by a projector in the common

algebra, and uses primitive label rigidity (`New/Trace/LabelRigidity.lean`) to rewrite the good idempotent, on the eigensystem, in terms of functions supported at classes where the simple trace formula applies. The trace comparison is a transfer result on the specific inner forms, with the cuspidal type and gap assumptions; it is not implied by the same rank of two groups. Once the traces agree for all P, however, the rest of the conclusion is ordinary algebra. Taking $P = 1$ shows equality of dimensions. The traces of all powers P^h agree and determine its characteristic polynomial by Newton's identities. This proves the determinant comparison `det_heckeOnGood_eq` and the rank identity `finrank_rows_eq_finrank_goodForms`. The assertions of trace comparison and primitive label rigidity here are formal-source inputs; the elementary polynomial implication was just proved.

Let

$$\lambda_j(\theta) = \det(u^*_{j,\theta^j} u_{j,\theta^j}) > 0.$$

The positive quantities λ_j cannot be assumed algebraic merely because they are determinants of spaces of forms; the height inequalities are statements about chosen ratios with uniform normalization scalars. The inner product on the final form space of the main chain is a class-sum inner product. Relative to the Petersson metric of the triangular chain there is a conversion factor

$$J(\theta)\, D_m, quad J(\theta) = [K_0 : J^{\mathrm{res}}_m(\theta)], \quad D_m = \#\,\mathrm{WeylIndex}(Y_m, \kappa_m).$$

The identity `stdLevelIndex_lastResLevel_eq` in `New/Exchange/ChainNorms.lean` expresses $J(\theta)$ as

$$CN(Q)^{k_0} \delta^m_{1,\theta} w_m / (w_0 w_1^m).$$

Here k_0 denotes the (integral) exponent called k in that declaration; it is unrelated to the rank of columns. The constants C and k_0 do not vary with the twist. This count compares block parabolic stabilizers in the quotient and is a *Formal-source input*; it cancels the triangular Gram factor $w_0 w_1^m / w_m$ exactly. The formula `inner_iotaAtLevel` in the same file gives the metric conversion. Injectivity follows from Lemma 17.13. Taking determinants of the resulting Gram comparison and applying Proposition 17.15 stage by stage yields the following identity.

Proposition 17.16 (Volume factorization)**.** *For every allowed m-fold twist, Θ_θ is injective; there exists a function $c(\theta)$ of cost m^2 squares such that*

$$Z(\theta) = |c(\theta)|^n \prod_{j=1}^{m} \lambda_j(\theta) > 0. \tag{17.15}$$

This is `Data.volume_iota_eq_prod_liftVolume` (`New/Rows/RowVolume.lean`). To check the scalar exponent, multiplying a Hermitian Gram form by a scalar multiplies its determinant by the nth power of that scalar; the factor in (17.15) is $|c|^n$, not $|c|^{2n}$. The residual factors resulting from the metric conversion, finite group orders and *delta* factors have squared norm equal to a fixed constant times an integral power of $N(Q)$, so they cost zero by Lemma 17.7. All that remains is the triangular cost of m^2. This explains both the injectivity presupposed in the forward-ratio calculation and the factorization required for the reverse ratio. The normalized scalar identities in this calculation are precisely those of the cited formal-source theorem.

17.8 Comparison forms for one lift

We next establish the height estimate for $\lambda_j(1)/\lambda_j(\theta_\iota)$ for each j. Fix j and abbreviate $u_\vartheta = u_{j,\vartheta}$, $Z_\vartheta = \mathrm{im}\, u_\vartheta$. In the next few paragraphs Z_ϑ is a subspace, not the scalar $Z(\theta)$.

The strategy is to use prescribed algebraic forms in the definite target, not to estimate the unknown Petersson norms of every automorphic representation separately.

Three character pairs are required for this particular j. Let $e_+ = m$, $e_- = m/2$, and $e_0 = 7(m/2) + k/2$ (the admissible parities make these integers). Write $p_i = (\nu_i, e_i)$ for $i \in \{0, +, -\}$. The pairs $p_\pm$ have character $\upsilon^{-e_\pm}\chi^{\pm j}$, where υ is the fixed rank-normalizing character and χ is the chosen tame character. The character in p_0 is $\mu_j^0 \upsilon^{j+\lfloor (t-11m-1)/2 \rfloor}$; here μ_j^0 is the untwisted character determining the jth lift. Use the twist ϑ positively for p_0, p_+ and negatively for p_-. Put

$$\Gamma_j(\vartheta) = \prod_{i=0,+,-} \Pi_{M\cup Q}((\nu_i \vartheta^{s_i})^{-1}, e_i), \quad s_0 = s_+ = 1, \quad s_- = -1. \tag{17.16}$$

This is `Data.heckeValueProd`, from `New/Rows/Separation.lean`; that file gives the precise pair definitions as `Data.linePair`. The long string p_0 will yield the term $k\mathcal{U}$, whereas the two other strings contribute to the small error.

Proposition 17.17 (Comparison maps; formal-source input)**.** *There is a finite-dimensional complex table space T, independent of ϑ, a surjection $A : T \to U_j$, and for every twist a linear map $g_\vartheta : T \to V_j(\vartheta)$ and $\gamma_\vartheta \neq 0$ such that*

$$u_\vartheta^* g_\vartheta = \gamma_\vartheta A, \tag{17.17}$$

$$|\gamma_\vartheta/\gamma_1|^2 = \Gamma_j(\vartheta)/\Gamma_j(1). \tag{17.18}$$

The vectors $g_\vartheta A^* v$ are tested relative to $Z_\vartheta^\perp$ for all $v \in U_j$. At every $\ell \nmid N(Q)$ and embedding ϕ, $g_\vartheta \zeta$ is integral whenever the conjugate tables $\bar\zeta$ are integral at $\phi \circ \kappa$ with the prescribed Weyl gauges of their classes. Finally

$$h_j \|g_1 \zeta\|_\sigma^2 \le e^{2rB_0} \sum_\alpha \|\sigma \zeta_\alpha\|_{\mathrm{Weyl}}^2,$$

where $h_j > 0$ is the hook number of the small algebraic weight.

This is exactly `exists_comparison_forms` in `New/Rows/ComparisonForms.lean`. We describe the detailed construction to specify its content. Split the definite Hermitian space in the auxiliary pair into $V'' \oplus V^\circ$ (the `DividedSplitting` data), and take tables indexed by a finite member index and by representatives of the small principal level $\bar\pi_g^2$ on V°. Their entries lie in the small Weyl module of weight given by `colsMinus`; thus h_j is the hook number of this module. To a member index corresponds a split lattice obtained by gluing the two blocks in the split model along the residue matrix. The same glued lattice can be chosen for every twist. The group in the large block which fixes the specified rows modulo $\bar\pi_g^2 Q$ has free class action; this is `classAut_rowFixingLevel_eq_bot`. If $y_{\vartheta,i}(h)$ is the line multiplier for member i, the comparison map is the algebraic form

$$g_\vartheta \zeta = g^{\mathrm{loss}} \sum_{(i,c)} \sum_{h \in X_i} \mathrm{productForm}_{L_i}((h,c), \overline{\beta_1 y_{\vartheta,i}(h)}, \zeta_{i,c}).$$

Here L_i is the glued lattice, X_i is the class set for the row-fixing level on the large block, β_1 is one nonzero common line scale, and the recipe called `productForm` applies the Cartan product of the representations in the divided splitting. This is the actual `comparisonForm`. The exponent loss is the one in the divided splitting; it clears the gluing loss from off-diagonal blocks. The source proves integral and small bounds for the line multiplier after applying β_1 in `Data.lineMultiplier_integral_and_small`, and controls the Cartan product in `New/Rows/CartanIntegrality.lean`. Multiplying a delta form by this product and summing

over a free class set preserves finite-place integrality. The archimedean estimates use the number of tables, classes and weights, and multiply the line bound with the squared Weyl norm of the input table. The output estimate is `normSqAt_comparisonForm_le` and the input integrality is `comparisonForm_mem_integralForms`. These assertions concerning the glued finite and infinite kernels are formal-source inputs.

The resulting inner product with a lift of $F \in U_j$ has a common period expression independent of the twist up to scalar: it is a weighted sum over members i of

$$\int_{[U(W_j)]} \overline{F(u)F_i^\natural(u)}\,\theta^\circ(\zeta_{i,\bullet})(u)\,du.$$

Here $F_i^\natural$ is a fixed natural form, and θ° the theta kernel from the small block. The scalar is determined by `Data.cosetKernelSum_eq_fNatural`, which gives $F_i^\natural$ as the coset kernel sum involving the line multiplier at every twist times κ_ϑ satisfying

$$|\kappa_\vartheta|^2 = C/\Gamma_j(\vartheta).$$

This is initially an assertion potentially with $C = 0$. The equality of period expressions is `inner_adjoint_comparisonForm`; the nonvanishing and surjectivity use the separation theorem `Data.eq_zero_of_integral_fNatural_tableTheta`. This states that if a good F gives zero in every displayed period, then $F = 0$. In fact if $\kappa_\vartheta = 0$, all $F_i^\natural = 0$ and the separation theorem forces $U_j = 0$, contrary to $n > 0$. If the range of $u_1^* g_1$ were not all of U_j, a nonzero orthogonal vector would annihilate all the periods, again contradicting separation. Set $A = u_1^* g_1$ and $\gamma_\vartheta = \overline{\kappa_1/\kappa_\vartheta}$. The period identity gives the first equation of Proposition 17.17; the quotient of the two norm identities for κ gives (17.18). This explains the proof of nonzero scalars and the fact that a *single* common surjection works for every twist. The equality of coset kernel sums and the separation theorem are *Formal-source inputs* in `New/Rows/Separation.lean`.

The tested condition in the proposition is equally substantive. Off level, the Hecke polynomial of a genuine lift has a product of t known roots and an affine rescaling of the original Hecke polynomial. Explicitly, if the auxiliary ranks are $r, r+t$, the polynomial at a split place w is

$$\prod_{v=0}^{t-1}(1 - q_w^{v+r/2}a_O(w)X)P(q_w^{t/2}a_O(w)b_O(w)^{-1}X)$$

in the character convention of `Data.lifts_auxLift` (`New/Exchange/CharacterLiftNorm.lean`); $a_O(w), b_O(w)$ in that display indicate the source's rank-normalized Frobenius values, and P is the source Hecke polynomial. In particular the test operator vanishes on the genuine lift, as expressed by `heckePolyOnForms_testValue_eq_zero`; one must use the test point in its precise source normalization, not insert a point from a different convention. A false lift satisfies this Hecke test and lies in the good part but is orthogonal to every actual lift. Its theta descent to the smaller group has zero good part (by adjointness and equivariance).

The crucial vanishing of the comparison periods for a false lift comes from enlarging the descent group by hyperbolic planes. The good polynomial labels imply that the reflection sums associated with the chosen tame character act by zero. Together with the vanishing of the initial descent and nontriviality of the character at the twist places, this kills the relevant boundary terms in the constant-term expansion. At a split place off level, a bounded descended form killed by a Hecke polynomial at a sufficiently small test point must vanish: expand the Hecke polynomial, use the sum of norms of the nonconstant terms strictly below one, and compare the sup norm. The precise statements `AuxPair.descentsVanish_of_reflectionSum`, `eq_zero_of_heckePolyOp_eq_zero` and `descentsVanish_add_planes` in `New/SiegelWeil`

`/AuxPairVanishing.lean` and `New/SiegelWeil/DescentsVanishPlanes.lean` formalize these steps. The period which results from the glued product is zero by `DividedSplitting.stdClassSum_inner_productForm_eq_zero` in `New/Rows/MemberOrthogonality.lean`. Applying that statement to the comparison construction gives `tableTheta_adjoint_eq_zero`: for a false lift Ξ, each small form $\theta^{\circ}((g_{\vartheta}^{*}\Xi)_{i,\bullet})$ vanishes. Consequently the common period formula gives

$$\langle \Xi, g_{\vartheta} A^{*} v \rangle = \langle A g_{\vartheta}^{*} \Xi, v \rangle = 0$$

for each v and every false-lift space. These special vanishing assertions are *Formal-source inputs*; the last implication is merely the adjunction identity. The source does not prove that arbitrary integral vectors satisfy the tested condition.

17.9 Compressed Gram determinants

Let p_{ϑ} be projection onto Z_{ϑ}, with respect to the class-sum inner product at the ordinary embedding. Define on Z_{ϑ}

$$D_{\vartheta} = \det(h_j p_{\vartheta} g_{\vartheta} g_{\vartheta}^{*}|_{Z_{\vartheta}}).$$

The key determinant calculation is independent of automorphic theory. Regarding u_{ϑ} as an isomorphism onto Z_{ϑ}, adjunction and Proposition 17.17 give

$$h_j u_{\vartheta}^{*} g_{\vartheta} g_{\vartheta}^{*} u_{\vartheta} = h_j |\gamma_{\vartheta}|^2 A A^{*}.$$

Taking determinants and using $u_{\vartheta}^{*} p_{\vartheta} = u_{\vartheta}^{*}$ gives

$$\det(u_{\vartheta}^{*} u_{\vartheta}) D_{\vartheta} = h_j^{n} |\gamma_{\vartheta}|^{2n} \det(A A^{*}). \tag{17.19}$$

Since A is onto, both sides are nonzero and g_{ϑ}^{*} restricted to Z_{ϑ} is injective. Dividing the identities at the trivial and arbitrary twist yields

$$\frac{\lambda_j(1)}{\lambda_j(\theta_\iota)} = \frac{D_{\theta_\iota^j}}{D_1} \left(\frac{\Gamma_j(1)}{\Gamma_j(\theta_\iota^j)} \right)^{n}. \tag{17.20}$$

This elementary identity is `det_compressedGramOp`, `adjoint_comp_eq_conj_smul` and their use in `heightLe_liftVolume_ratio` (`New/Rows/LiftVolumeHeight.lean`).

The next two estimates for the compressed determinant must be proved without assuming that projection to a complex automorphic summand is integral or that it is an orthogonal projection at all coefficient conjugates. We give the determinant-level argument.

Lemma 17.18 (Finite norm-adapted basis)**.** *Let V be a complex vector space of dimension d, let $(\ell_a)_{a \in A}$ be a finite separating family of complex linear functionals on V, and let ϕ be a field embedding into a non-Archimedean valued field. Then some d of these functionals are the dual basis of vectors v_i with $|\phi \ell_a(v_i)| \le 1$ for all a, i.*

Proof. Take any basis of V. Among all d-tuples of the given functionals, choose one maximizing the absolute value under ϕ of their determinant in this basis. The maximum is positive by separation. Use the chosen tuple as a dual basis. Express any other functional in this dual basis by Cramer's rule. Each coordinate is a quotient of a determinant obtained by replacing one of the selected rows by the new functional, by the maximizing determinant. Its *phi*-absolute value is at most one. This proves the assertion. The finite case of `exists_auerbach` in `Lib/General/LinearAlgebra/CommonEigenvectors.lean` is all that is needed here. □

At a fixed place ι, choose rational integral adelic frames for the small-block classes. Denote by f_a special table vectors given by the coordinate duals in those frames, scaled by the hook-number convention. Thus the functional on Z_ϑ is

$$\ell_a(z) = \langle f_a, g_\vartheta^* z \rangle.$$

The duals are integral in the Weyl gauge convention because the Fock Gram matrix times the relevant Weyl Gram matrix is the hook number times the identity (the Fock pairing of standard tableaux supplies the integral dual). In these frames the values $\ell_a(z)$ are exactly h_j times the conjugate-coordinate test of the table $g_\vartheta^* z$. Consequently, selecting a dual basis as in Lemma 17.18, each $h_j g_\vartheta^* z_i$ satisfies the requisite integral table test at the chosen place. The forms $g_\vartheta f_{a_i}$ are integral at all the embeddings used in the determinant test, and

$$v_i = h_j g_\vartheta g_\vartheta^* z_i$$

are integral at ι and tested: adjunction gives $g_\vartheta^* u_\vartheta = \bar{\gamma}_\vartheta A^*$, so the tested condition follows from the comparison theorem for $z_i \in Z_\vartheta$. In this step the assertion on the coordinate duals (including their gauge after conjugation) is the concrete integral-linear-algebra input of `norm_det_compressedGramOp_padic_le`; it relies on the integral frame lemma and `weylGram_mul_fockPairing` rather than on orthonormal coordinates having integral entries.

Apply (17.9) to these two families. The matrix of their pairings

$$\left(\langle g_\vartheta f_{a_i}, p_\vartheta(h_j g_\vartheta g_\vartheta^* z_h)\rangle\right)_{i,h}$$

is the matrix of the compressed operator in the chosen basis, by duality of ℓ_{a_i}. Its determinant is D_ϑ. Hence the determinant bound at twist $\vartheta = \theta_\iota^j$ is

$$|\iota D_\vartheta|_\ell \le e^{nr\beta_\ell}. \tag{17.21}$$

This derives the implication of the denominator test in `norm_det_compressedGramOp_padic_le`; it is precisely why determinant-level denominators suffice.

For the archimedean estimate, use the following property for all conjugate coefficient metrics: if y is tested and Z_ϑ lies in the good kernel of the Hecke polynomial at the appropriately chosen off-level place, then

$$\|p_\vartheta y\|_\sigma \le \|y\|_\sigma. \tag{17.22}$$

This is `AuxPair.normSqAt_starProjection_le` in `New/Selmer/FitTwists.lean`. Here is the mechanism. Write $y = z + x$ with $z \in Z_\vartheta$, x orthogonal to Z_ϑ. For a tested vector, x lies in the sum of the kernel of the good projector and the range of the Hecke polynomial (taking the complementary conjugate-place normalized operator as appropriate). This follows by taking orthogonal complements of the intersection defining the false-lift test, using the normality of the Hecke operators. Under application of a field automorphism σ, the adjoint of $H_w(X)$ for the coefficient-conjugate pairing is $H_{\bar{w}}(\sigma^{-1}(\overline{\sigma X}))$. This is `LevelPlaces.LevelForm.innerAt_heckePolyOnForms`, and the commuting Hecke operators and definiteness of the twisted metric imply that their relevant kernels are stable under taking the adjoint. In addition the good projector remains self-adjoint for every conjugate pairing because its underlying finite character action is unitary (character scalars are roots of unity). Hence z, x remain orthogonal for the σ-pairing. Positivity gives (17.22). The adjoint and projector laws with their precise local hypotheses are *Formal-source inputs* in `New/Selmer/FitTwists.lean`.

Apply this contraction to the vectors $g_1 g_1^* z$ occurring in the compression. One cannot simply apply σ to the spectral bound at the ordinary embedding. Instead, take an eigenvector

z with eigenvalue η, set $\zeta = g_1^* z \neq 0$, and pair the eigenvalue equation with the conjugate-compatible table obtained from ζ by the map $\kappa \circ \sigma^{-1} \circ \kappa \circ \sigma$ on coefficients. The mixed Cauchy–Schwarz inequality of (17.3), contraction (17.22), and the two table norm estimates in Proposition 17.17 then give

$$|\sigma\eta| \sum_{\alpha} \|\sigma\zeta_\alpha\|_{\mathrm{Weyl}}^2 \leq e^{2rB_0} \sum_{\alpha} \|\sigma\zeta_\alpha\|_{\mathrm{Weyl}}^2.$$

The sum is positive by injectivity of $g_1^*|_{Z_1}$. This is the eigenvalue calculation in `norm_det_compressedGramOp_arch_le` (`New/Rows/LiftVolumeHeight.lean`). Thus every conjugate of each eigenvalue has modulus at most e^{2rB_0}, and the determinant, the product of the n eigenvalues with multiplicity, satisfies

$$|\sigma D_1| \leq e^{2nrB_0}. \tag{17.23}$$

Using witness D_1 in the height definition, the local estimate (17.21), the global sum of β_ℓ, and (17.23), we obtain

$$\mathrm{Ht}_{\leq 3nrB_0}\left(\frac{D_{\theta_\iota^j}}{D_1}\right). \tag{17.24}$$

17.10 The height saving in a rank-one special value

It remains to show how derivative divisibility enters the lower bound, in the estimate for $\Gamma_j(1)/\Gamma_j(\theta_\iota^j)$. Fix a pair $p = (\nu, z)$ of rank-one type at j. The definite target of the lift of the trivial character of $U(1)$ has even rank $2z$ and its one-row coefficient weight at τ has degree $d_\tau = c_\tau - z$. Let $\Xi_{p,\vartheta}$ denote this lift, with the CM vacuum and jet scaling, namely `Data.Pair.trivialLift` in `New/Exchange/AlgebraicPartDefs.lean`. It is a scalar multiple of a sum of kernel forms over the classes of the norm-one torus, at the prescribed level. Put $d_{\mathrm{tot}} = \sum_\tau d_\tau$. The scaling of the theta kernel in this sum is

$$\kappa_0^{2z} (\overline{\pi_g}\,\mathrm{per})^{d_{\mathrm{tot}}},$$

where κ_0 is the common vacuum scale, π_g is the chosen level generator and per the jet scale. The integer

$$\mathrm{den}_p = g^{4fz} (gd)^{d_{\mathrm{tot}}}$$

clears its denominators outside its conductor. Let $M_* = \mathrm{chosenGroupOrder}$ be the fixed finite group order bounding the relevant class stabilizers. The normalized integral value used in the source is

$$\mathcal{A}_p(\vartheta) = M_* \,\mathrm{den}_p^2 \,\frac{\|\Xi_{p,\vartheta}\|^2}{N(Q)^{2z-1}}.$$

Nonzero proportionality of this value to $L^{M\cup Q}(z, (\nu\vartheta)^{-1})$ with a common constant independent of ϑ is the rank-one Rallis norm theorem `LValueData.inner_trivialLift_self_eq_L` (`New/Exchange/CharacterLiftNorm.lean`) and `intAlgPart_eq_mul_L` of `New/Exchange/TriangleLevelExponent.lean`. The factorial in the norm is exactly $\prod_\tau d_\tau!$, because a one-row standard monomial has Fock norm its factorial. The assertion that the residual normalization is common, including the indicated conductor power, is a *Formal-source input*; the ensuing ratios satisfy

$$\rho_p(\vartheta, 1) = \frac{\mathcal{A}_p(\vartheta)}{\mathcal{A}_p(1)} = N(Q)^{-(2z-1)} \frac{\|\Xi_{p,\vartheta}\|^2}{\|\Xi_{p,1}\|^2}. \tag{17.25}$$

In particular the required lift is nonzero. That normalized values are values of one group-ring sum with coefficients integral away from Q, when the twist varies at fixed modulus, is

`isIntegralMeasure_intAlgPart` in the latter file. It follows by expanding the squared class sum of kernels into the two torus class indices and pulling the character factor out of their two translated kernels; the remaining coefficients are integral by the norm-one lift denominator estimate and clearing the stabilizers by M_*. These precise integral-measure and norm assertions are formal-source inputs, not appeals to a general interpolation theorem for L-functions.

To estimate the untwisted norm, compare with a flat pair p^0 of the same rank, i.e. of all degrees zero, and with a permissible weight-zero character after rank normalization. Define

$$Q_0 = \mathcal{A}_{p^0}(1)/\|\Xi_{p,1}\|^2.$$

The ordinary rank-one norm itself need not be integral, but $Q_0\|\Xi_{p,1}\|^2 = \mathcal{A}_{p^0}(1)$ is integral at all finite embeddings. The precise archimedean comparison is

Proposition 17.19 (Rank-one conjugate norms; formal-source input)**.** *For fixed rank $2z$, at every automorphism fixing K there is a common c_σ independent of the pair such that*

$$\log|\sigma(\|\Xi_{p,1}\|^2)| = c_\sigma + e_{p,\sigma} + \sum_\tau J(d_\tau), \qquad e_{p,\sigma} \geq -4f.$$

For a flat pair $e_{p,\sigma} \leq 4f$ as well.

This is `Data.Pair.log_norm_map_norm_sq_trivialLift` in `New/Exchange/AlgebraicPartConjugates.lean`, obtained from the permutation and translation of the CM theta coefficients under field conjugation (`LValueData.map_trivialLift`). For automorphisms not fixing K, compose with complex conjugation to reduce the required modulus estimate to this case. Compare the upper bound for the flat numerator with the lower bound for the nonflat denominator. Because the weight of a pair of rank-one type has $d_\tau \geq w_0(\tau)$, the jet-size estimates compare $J(d_\tau)$ with $J(w_0(\tau))$. The extra logarithm of the fixed integer $M_* \operatorname{den}_{p^0}^2$ and the two losses of $4f$ have combined size at most B_0 (the exact bound is `D.le_of_hasLogBound` in `heightLe_norm_trivialLift_sq`). We deduce

$$\mathrm{Ht}_{\leq B_0 - \sum_\tau J(w_0(\tau))}(\|\Xi_{p,1}\|^2). \tag{17.26}$$

This is the declaration `heightLe_norm_trivialLift_sq` (`New/Rows/Separation.lean`). It is an estimate for the constant family with entries the specified squared norm, not for the absolute size of the norm at every place. The common unknown c_σ disappears on taking the ratio.

We need the finite-place lower bound on the pairing with a test form that restores the forced divisibility term. Put $F_\ell = \sum_\tau A_\ell(w_{\mathrm{nom}}(\tau), \tau)$, the quantity `Data.forcedPart`, so that the elementary prime-power formula gives

$$2\sum_\ell F_\ell \log \ell = \sum_\tau A(w_{\mathrm{nom}}(\tau), \tau).$$

All terms have finite support. The relevant depth indicator is `Data.depth_indicator_jet` in `New/Rows/GluedFormTheta.lean`. It is a *Formal-source input*: for the degree d_τ in the allowed interval of that declaration, one common nonzero scale of archimedean conjugates bounded by its prescribed small budget serves all places; locally there is a jet of the prescribed total degrees (distributed among twelve coordinates) with valuation at most the forced depth loss, after multiplication by the scale. The proof in the source uses tight frames and the inverse-theta normal form to realize a large coefficient of the needed jet. An earlier mere upper bound for the norms of all jets cannot replace this attainability statement.

In obtaining an actual test form, the twelve coordinates sit next to flat blocks of ranks $2z-20, 4, 4$. They permit the coefficient of the jet to be detected by a genuine coefficient pairing and absorb the tame character choices through the flat forms. The test-form existence formula is `exists_test_form_inner_kernelForm` (`New/Rows/HeckeValueHeight.lean`). It constructs a form by three class sums weighted by the flat theta lifts in these three blocks. Its pairing with the kernel is an explicit product of their three normalized values, a multinomial coefficient

$$\prod_{\tau}(d_\tau!)/\prod_{\tau,i}\alpha_{\tau,i}!$$

and the remaining rank-twelve monomial jet kernel (with the matching character). The left side contains the class-number normalization and the vacuum-scale power. It asserts the resulting form is integral at every embedding over ℓ and is tested for the rank-one false lifts. The tested assertion relies on the same extra-plane descent vanishing as the comparison form theorem, with the rank-one label projector. This is a separate *Formal-source input*, since the existence of an integral kernel coefficient alone is not enough to produce a class-sum pairing against an integral tested vector.

For clarity there are two elementary finite Fourier steps in passing from these pointwise test-form formulas to the uniform height witness.

Lemma 17.20 (Fourier selection and multiplicative selection)**.** *Let G be a finite abelian group and $s : G \to \mathbb{C}$ any function. For each embedding ϕ in a non-Archimedean field there is a character χ such that*

$$\left|\phi\left(\frac{1}{|G|}\sum_{g\in G}\chi(g)s(g)\right)\right| \geq |\phi(s(1))|.$$

Furthermore, let $a(\chi), b(\chi)$ be two families indexed by a finite character group $G^\vee$, and let Q satisfy $Q^{|G^\vee|} = \prod_\chi a(\chi)b(\chi)$. For every prescribed $\psi \in G^\vee$ there are χ_1, χ_2 with $\chi_1\chi_2 = \psi$ such that

$$|\phi(a(\chi_1)b(\chi_2))| \geq |\phi Q|.$$

The second assertion also holds for Archimedean embeddings.

Proof. Fourier inversion expresses $s(1)$ as the sum of the normalized Fourier coefficients. The ultrametric inequality bounds its absolute value by their maximum. For the second assertion, take the product of the numbers $a(\psi\chi^{-1})b(\chi)$ as χ ranges over all characters. It is $Q^{|G^\vee|}$ by the permutation of the character group, so one absolute value is at least $|\phi Q|$. □

In `exists_twisted_classSum_ge`, the finite group is the group of torus classes after translating the chosen base class to the identity; its stabilizers are trivial and its characters are realized by the permitted tame weight-zero characters. The second step is `exists_factorisation_le` in the same file. These two selections allow the flat class sums to be paired with the depth indicator at every local embedding while a single product of flat values is used globally as the height witness.

Here is the precise quantitative interface obtained by applying these constructions. Fix the rank-one pair p, the index j and the small flat pair p_2 (rank four, tame part the chosen primitive character ω). Let N_c be the torus class number at the principal chosen level. The combination of `depth_indicator_jet`, `exists_test_form_inner_kernelForm`, `exists_twisted_classSum_ge` and `exists_factorisation_le` in the proof of `heightLe_inv_ratio` (`New/Rows/HeckeValueHeight.lean`) yields a common $Q \neq 0$ such that $|\sigma Q| \leq e^{B_0}$ for every

σ, and for each ℓ and each embedding ϕ there exists a tested, everywhere ℓ-integral vector u_ϕ in the rank-one target satisfying, at $\vartheta = \theta_\iota^j$,

$$|\phi Q|_\ell \, \ell^{-F_\ell} \le \left| \phi \left(\frac{N_c^2 \langle u_\phi, \Xi_{p,\vartheta} \rangle}{\rho_{p_2}(\vartheta, 1)} \right) \right|_\ell . \tag{17.27}$$

The witness Q is an hth root of the product of the flat normalized values times the depth scale over the h tame characters, which gives the archimedean bound using the flat-conjugate estimates. For the fixed twist and a different finite embedding the test vector may change, but Q does not. Equation (17.27), including its conjugate-compatible version when one uses the field embedding composed with complex conjugation, is the local assertion called `K` in the proof of that declaration. We take this simultaneous assertion, which uses the attainable jet and all the class and gluing integrality facts just discussed, as a formal-source input. The following calculation reconstructs its numerical consequence.

For an arbitrary nonzero vector v in a complex Hilbert space, projection onto its line is

$$p_{\mathbb{C}v} w = v \langle v, w \rangle / \|v\|^2 .$$

The fit denominator condition for the rank-one lift line therefore bounds the pairings of two integral tested vectors with this projection by e^{β_ℓ}. In the source the strengthened form `DenominatorLe.rank_one_bound` in `New/Selmer/RankOneDenominator.lean` gives

$$|\iota(N_c \langle u, p_{\mathbb{C}\Xi} v \rangle)|_\ell \le e^{\beta_\ell}$$

for such vectors. The extra class number N_c is retained by replacing the second vector by a sum of its twisted central translates, which are still integral and tested, whereas the line of Ξ is fixed under this central action; this is again a *Formal-source input*. Apply (17.27) at the two embeddings giving the two factors of the rank-one projection pairing. The lower bound for their product and the projection upper bound give

$$\left| \iota \left(Q \overline{Q} \, \frac{\rho_{p_2}(\vartheta, 1)^2}{N_c^3 \|\Xi_{p,\vartheta}\|^2} \right) \right|_\ell \le e^{\beta_\ell + 2F_\ell \log \ell} .$$

The exponent 3 of N_c is obtained from the two factors of N_c^2 and the single factor of N_c in the strengthened projection denominator. Since $|\sigma(Q\bar{Q})| \le e^{2B_0}$, summing the finite-place bound gives

$$\mathrm{Ht}_{\le 3B_0 + 2\sum_\ell F_\ell \log \ell} \left(\frac{\rho_{p_2}(\vartheta, 1)^2}{N_c^3 \|\Xi_{p,\vartheta}\|^2} \right) .$$

Multiply by the estimate (17.26) and by the square of the inverse ratio of the ramified flat pair p_2 (height at most $2x^{999}$ from fit). Using (17.25), the product is $\rho_p(\vartheta, 1)^{-1} / (N_c^3 N(Q)^{2z-1})$. The integer factor N_c^3 costs nothing when multiplying back (its local absolute value is at most one), and the conductor power is a local unit. The saving is exactly minus the cost (17.1). Since $x^{999} \le B_0$ by the admissible magnitudes, we conclude

$$\mathrm{Ht}_{\le 6B_0 - \mathcal{U}}(\rho_p(\theta_\iota^j, 1)^{-1}). \tag{17.28}$$

This is `heightLe_inv_ratio`.

17.11 Combining the special values and closing the lower bound

We now have precisely the estimate for the long string in (17.16). For p_0 consider lower points z with $7m/2 < z \le e_0$. There are $k/2$ of them, all of rank-one type at j; this follows

from the specified weight of ν_0 and the numerical relations between r, k, t, m (the range verification occurs in `heightLe_inv_heckeValueProd_div`). Thus these $k/2$ inverse ratios have height at most $6B_0 - \mathcal{U}$ each, by (17.28). The remaining factors in the three products have $7m/2 + e_+ + e_- = 5m$ lower points in all; their underlying pairs are ramified at exactly M and their inverse ratios are bounded by the fit hypothesis with cost x^{999} each. All ratios enter squared, and the conductor powers in (17.8) are local units. Consequently

$$\mathrm{Ht}_{\le 10mx^{999}+k(6B_0-\mathcal{U})}\left(\frac{\Gamma_j(1)}{\Gamma_j(\theta_\iota^j)}\right). \tag{17.29}$$

This proves the deduction in `heightLe_inv_heckeValueProd_div` (`New/Rows/HeckeValueHeight.lean`). One should notice the count $k/2$ before squaring; counting k factors before the square would yield an incorrect leading coefficient.

Combine (17.29), the compressed determinant estimate (17.24), and identity (17.20). The total height of the one-lift inverse volume ratio is at most

$$3nrB_0 + n\big(10mx^{999} + k(6B_0 - \mathcal{U})\big) \le nk((\log b)B_0 - \mathcal{U}).$$

The scale inequality on the right is the one in `heightLe_liftVolume_ratio` (`New/Rows/LiftVolumeHeight.lean`): r/k is bounded, $mx^{999}/(kB_0)$ negligible and $\log b$ very large. Thus all the positive errors are absorbed by the $k(\log b)B_0$ allowance. Multiplying over $j = 1, \dots, m$ gives height bound $n(L - H)$ for the product of the ratios.

By Proposition 17.16 the full-chain ratio squared is the product of the squared single-lift ratios multiplied by the nth power of the inverse squared exchange factor. Its cost is m^2 squares, so its nth power has height at most $2nm^2tx^{999} \le 2nL$. We obtain

Proposition 17.21 (Reverse ratio)**.** *For every fit family,*

$$\mathrm{Ht}_{\le 2n(2L-H)}\left(\frac{Z(1)^2}{Z(\theta_\iota)^2}\right). \tag{17.30}$$

This is `Data.heightLe_volume_ratio_sq` (`New/Endgame/LongForms.lean`). Finally choose a fit family, apply Proposition 17.12 and square its ratio. The product of its ratio squared and that of (17.30) is identically one by injectivity. Lemma 17.4 gives

$$0 \le 2(V + 2d_{\mathrm{ind}} + nb) + 2n(2L - H).$$

As $b \le L$, rearrangement yields $n(H - 3L) \le V + 2d_{\mathrm{ind}}$, proving Theorem 17.1 from the explicitly identified inputs.

The arithmetic content of this argument is not an estimate of the size of individual L-values. It consists of finding comparison vectors for all good forms, extracting the rank-one ratios from their common periods, recovering adequate local integral tests from the geometrically constructed theta coefficients, and ensuring that projections of the tests have small determinant denominator. The same prime-power sum that appears in the divisibility cost in Theorem A is recovered, with precisely the opposite sign, in (17.28). The remaining issue, necessary for both occurrences of the height argument, is the existence of the simultaneous fit family. This calls for the congruence and local cohomological arguments in the next section.

Chapter 18

Period identities, finite matching, and the geometry of the lower bound

The proof of Theorem C in Chapter 17 uses period identities in two quite different ways. One transports an integral structure to the whole row space by comparing periods of complementary rank; the other computes the metric on a chain by comparing the pairing of consecutive theta lifts with a product of simpler pairings. The local character at the ramified places matters in both operations. We therefore develop the algebra of the residue matrix functions first, then explain the global period calculation and its relation to the integral tests. In particular, the matching identity for finite functions, the Fourier transform of the point data, and the support condition in a Rallis pairing can be analyzed without assuming any new nonvanishing theorem for automorphic forms. We retain the exact notation of Chapter 17 for the global variables, but use lower case a_0, b_0, n_0 for temporary matrix sizes. Statements marked *formal-source input* identify the assertions whose independently written proof is not completed in this monograph; the deductions from those assertions are provided explicitly.

18.1 Complementary frames and matched data

Let R be a finite product of quotients of discrete valuation rings by powers of the maximal ideal. For an N-dimensional free module, let $N = a_0 + b_0$, and choose an ordered splitting of the coordinates into b_0 and a_0. Write

$$k = [Y\ Z] \in \mathrm{GL}_N(R), \qquad k^{-1} = \begin{bmatrix} T \\ X \end{bmatrix},$$

so that $Y \in \mathrm{Mat}_{N \times b_0}(R)$, $X \in \mathrm{Mat}_{a_0 \times N}(R)$. Then

$$XY = 0, \quad TY = I_{b_0}, \quad XZ = I_{a_0}, \quad YT + ZX = I_N. \tag{18.1}$$

In particular Y and X belong to a split exact sequence $0 \to R^{b_0} \xrightarrow{Y} R^N \xrightarrow{X} R^{a_0} \to 0$. In the following such a pair is called compatible. This is a precise condition over a non-field: an injective matrix over the residue ring need not by itself define a direct summand. A matrix on which a left inverse exists does define a direct summand, as needed here.

Lemma 18.1 (Complementary row and column bases)**.** *An N by b_0 matrix Y with a left inverse extends to $[Y\ Z] \in \mathrm{GL}_N(R)$. Dually a surjection $X : R^N \to R^{a_0}$ extends to the*

last rows of an inverse matrix. Every compatible pair (X, Y) occurs as the corresponding submatrices of k and its inverse, and the determinant of such a k is determined by the pair. The choices of k giving a fixed pair form a torsor for the additive group $\mathrm{Mat}_{b_0 \times a_0}(R)$.

Proof. It suffices to treat one local factor of R, since a basis on each local factor gives a basis over the product (with the indicated ranks). A left inverse reduces modulo the maximal ideal to a left inverse, so the reduced columns are linearly independent over the residue field. Extend these columns to a basis of the residue vector space and lift the added columns: the determinant of the extended square matrix is a unit. The dual assertion follows by transposition. For a split exact sequence choose any section Z of X. The map $[Y\ Z]$ is an isomorphism by exactness: for any u, $u - ZXu$ is in the image of Y, with unique preimage. If Z' is another section, then $Z' - Z = YU$ for unique U. Thus

$$[Y\ Z'] = [Y\ Z] \begin{bmatrix} I & U \\ 0 & I \end{bmatrix},$$

proving all claims, including the assertion on determinants. □

This supplies the finite-ring instance of `IsValuationPair.exists_firstCols_iff` in `Lib/General/LinearAlgebra/PartnerAlgebra.lean`; the source has a more general valuation-pair formulation. For fixed Y, the compatible X's are a torsor under $\mathrm{GL}_{a_0}(R)$ acting by left multiplication; for fixed X, the compatible Y's form a torsor under $\mathrm{GL}_{b_0}(R)$ acting by right multiplication. These assertions follow either from (18.1) or from the uniqueness of an isomorphism of free modules inducing a given basis of the kernel or quotient.

Let $\chi : R^\times \to \mathbb{C}^\times$ be a character extended by zero off the units. In the text this extension is meant whenever $\chi(\det M)$ occurs for a square matrix M that is not assumed invertible; its inverse character χ^{-1} is also extended by zero. A *matched pair* (Φ, Ψ) of functions on row matrices X and column matrices Y has the following three properties: each function is zero off the frames that can occur in Lemma 18.1, and for every compatible pair furnished by k,

$$\Phi(X) = \chi(\det k)\Psi(Y). \tag{18.2}$$

This is `IsPartner` in the above file. In particular $\Phi(fX) = \chi(\det f)^{-1}\Phi(X)$ and $\Psi(Yq) = \chi(\det q)^{-1}\Psi(Y)$ for invertible f, q. It is important that both functions have inverse character covariance: indeed the compatible pair (fX, Yq) is given by $k\,\mathrm{diag}(q, f^{-1})$, whose determinant is $\det k \cdot \det q / \det f$.

Proposition 18.2 (Open-cell formula and all matched pairs)**.** *For $X = [X_1\ X_2]$, $Y = \binom{Y_1}{Y_2}$ in the coordinate splitting $R^{b_0} \oplus R^{a_0}$, let p be any function on a_0 by b_0 matrices. The functions*

$$\Phi(X) = \chi(\det X_2)^{-1} p(X_2^{-1} X_1), \qquad \Psi(Y) = \chi(\det Y_1)^{-1} p(-Y_2 Y_1^{-1})$$

are partners, with the respective formula understood as zero when the square block is not invertible. More generally every inverse-character-covariant function Ψ supported on left-invertible column frames admits a unique partner Φ.

Proof. For compatible X, Y supplied by k, the identities $XY = 0$ and

$$\det X_2 \det k = \det Y_1 \tag{18.3}$$

follow from (18.1); to verify the determinant identity, multiply the square block matrix with first block row $[I\ 0]$ and second block row X by k. The product is block upper triangular with lower right block I, and upper left block Y_1. Thus X_2 is invertible exactly when Y_1 is.

In that case $X_2^{-1}X_1 = -Y_2Y_1^{-1}$, and (18.3) gives (18.2). The support assertion is immediate from the invertible chosen block. For general Ψ, define $\Phi(X)$ by choosing compatible Y, k and setting the right hand side of (18.2). Changing Y to Yq multiplies the determinant by $\det q$ and Ψ by the inverse character. Changing k for a fixed pair does not change its determinant by Lemma 18.1. This proves well-definedness, uniqueness, and the required covariance and support. □

The first assertion is the precise formula `partner_block` in `New/Exchange/BlockPartners.lean`. The second, elementary completion is useful in understanding why the intertwining identity on partners yields an explicit matrix kernel rather than an unspecified relation between periods. A simple example over a field is $a_0 = b_0 = 1$. Then every nonzero column $(y_1, y_2)^T$ with $y_1 \neq 0$ corresponds to the row $(x_1, x_2) = (-y_2, y_1)$, up to scaling, and the parameter $-y_2/y_1 = x_1/x_2$ is the coordinate of the annihilator line on the same affine chart of a projective line. Since $\det k$ may be a nontrivial unit, the inverse-character factors cannot be dropped even in this example.

For a compatible pair let

$$\Delta_\chi(X, Y) = \frac{\chi(\det k)}{|\operatorname{GL}_{b_0}(R)|}, \tag{18.4}$$

and put it equal to zero otherwise. The determinant is well-defined by Lemma 18.1. Given invertible $g \in \operatorname{GL}_N(R)$, $f \in \operatorname{GL}_{a_0}(R)$, $q \in \operatorname{GL}_{b_0}(R)$, the new pair $(f^{-1}Xg, g^{-1}Yq)$ is obtained from $g^{-1}k \operatorname{diag}(q, f)$, proving

$$\Delta_\chi(X, Y) = \chi(\det g/(\det f \det q))\Delta_\chi(f^{-1}Xg, g^{-1}Yq). \tag{18.5}$$

For a covariant Ψ as in Proposition 18.2, summation of (18.4) against Ψ on column frames for a fixed X gives exactly its partner Φ. Thus the denominator in (18.4) is the orbit size, not a multiplicative loss in the integrality estimate.

Lemma 18.3 (Recovering a period point)**.** *Suppose L, L' are functions of row and column matrices, respectively, transforming by χ (not χ^{-1}) under their frame groups, and satisfying*

$$\sum_Y \Psi(Y)L'(Y) = c\sum_X \Phi(X)L(X) \tag{18.6}$$

for all partners. Then for each left-invertible Y,

$$L'(Y) = c\sum_X \Delta_\chi(X, Y)L(X).$$

There is a normalization $d(Y)$ of the kernel with the properties that $d(Y)\Delta_\chi(X, Y)$ is either zero or a root of unity, and its average under changing the column frame is the original kernel in the sense

$$\frac{1}{|\operatorname{GL}_{b_0}(R)|}\sum_q d(Yq)\Delta_\chi(X, Y) = \Delta_\chi(X, Y).$$

Proof. Let Y be fixed. Take Ψ supported on the right frame orbit of Y, with $\Psi(Yq) = \chi(\det q)^{-1}/|\operatorname{GL}_{b_0}(R)|$. The covariance makes the left side of (18.6) equal to $L'(Y)$. Its partner is $\Delta_\chi(\cdot, Y)$ by the orbit calculation above, proving the formula. For each orbit choose a single representative and let $d(Y)$ equal the group order at the representative and zero elsewhere. The action on each left-invertible frame is free, so the indicated average is one. On the support of a compatible pair, multiplying the kernel by d either annihilates it or leaves the character of a determinant, hence a root of unity. □

Here the replacement of the original kernel by the normalized kernel does *not* preserve the equality pointwise for a fixed Y outside the selected representative. Instead, one uses an equivariant period functional in the square argument, which allows the displayed averaging. This is exactly the role of the function d in `exists_partner_kernel`, `New/Exchange/SquareLawFourier.lean`. There is also a support condition in that lemma. Write $Y = \binom{Y_1}{Y_2}$ and $X = [X_1\ X_2]$. If Y_1 has a left inverse L_1, then X_2 has a left inverse. Indeed set $P = I - Y L_1 \operatorname{pr}_1$ and choose a right inverse Z of X. We have $PY = 0$ and $PZX = P$ by (18.1). Since $P\binom{0}{v} = \binom{0}{v}$, the matrix $\operatorname{pr}_2 PZ$ is a left inverse of X_2. The assertion with the indices reversed is identical. This verification is necessary for propagating tidiness to the next edge of the grid. It explains why over a residue ring the proof uses the existence of a left inverse, not just a rank inequality at one embedding.

These constructions also have an integral local version. Let F be a local field with ring of integers O, or a finite product of such fields, and $Y : F^{b_0} \to F^N$ injective. The module $\{z : Yz \in O^N\}$ is an O-lattice of rank b_0, and its image under Y is saturated in O^N. After choosing a basis for this lattice by a matrix $g \in \mathrm{GL}_{b_0}(F)$, the matrix Yg is primitive (its entries are integral and its columns extend to an integral basis). The right coset $g\,\mathrm{GL}_{b_0}(O)$ is uniquely determined; this follows because the image of O^{b_0} must be precisely that lattice. Consequently, in a sum over such cosets only one term can survive when the function is supported on primitive frames. This establishes the normalization used in `localSection_of_primitive` and clarifies the relation between matching over the local field and matching over the residue ring. The full algebraic extension of local-section functionals from primitive data to all locally constant compactly supported functions, when needed without finite support of the coset sum, is retained as a formal-source input `exists_common_coset_reps` and `matchingPairs.exists_linear_factor` in `New/Exchange/BlockPartners.lean`. Neither is used to replace a divergent sum by an unqualified analytic integral.

18.2 Fourier transformations of the matching identity

A finite ring R as above admits primitive additive characters: a character $\psi : (R, +) \to \mathbb{C}^\times$ is called primitive when its kernel contains no nonzero ideal. A multiplicative character is primitive in the stronger sense used here if, for every $x \neq 0$, there exists u such that $1 + ux$ is a unit and $\chi(1 + ux) \neq 1$. At squarefree level this amounts to taking a nontrivial multiplicative character on each residue-field factor. Indeed if $x \neq 0$ at one factor, one can prescribe $1 + ux$ there to be an arbitrary unit and take it to be 1 at all other factors. At higher conductor powers these hypotheses must be checked separately. In the compact code the primitive multiplicative condition at squarefree level follows from `isPrimitiveMulChar_of_isTameFor`, `New/Exchange/CostOfSquares.lean`; it is used only in the square-free triangle, not all the double-grid transport at level M^2Q.

For rectangular matrices A define the unnormalized finite Fourier transform by

$$(\mathcal{F}_\psi p)(A) = \sum_B p(B)\psi(\operatorname{tr}(AB)),$$

where B has the transpose shape to A. For a primitive character and a fixed nonzero matrix C the function $B \mapsto \psi(\operatorname{tr}(CB))$ is nontrivial by varying the entry of B paired with any nonzero entry of C. The sum of a nontrivial character over a finite group is zero (translate the sum by a point on which the character is not one). Thus

$$\sum_B \psi(\operatorname{tr}(CB)) = \begin{cases} |R|^{uv} & C = 0, \\ 0 & C \neq 0, \end{cases} \tag{18.7}$$

for C of size u by v. Expanding the double sum proves

$$\mathcal{F}_\psi^2 p(C) = |R|^{uv} p(-C), \qquad \sum_A |\mathcal{F}_\psi p(A)|^2 = |R|^{uv} \sum_B |p(B)|^2.$$

These are `matrixFourier_matrixFourier` and the finite-character orthogonality used in `Lib/Weil/GaussSum/LocalHomUnit.lean`. The present normalization, with no factor $|R|^{-uv/2}$, is essential in the following character Gauss sum.

Proposition 18.4 (Matrix determinant Gauss sum)**.** *Let d be a positive integer and let F be a function on d by d matrices with left covariance $F(gB) = \chi(\det g)F(B)$ for all $g \in \mathrm{GL}_d(R)$. For primitive ψ and χ put*

$$G_d(\chi, \psi) = \sum_{B \in \mathrm{Mat}_d(R)} \chi(\det B)\psi(\operatorname{tr} B).$$

Then $|G_d(\chi, \psi)|^2 = |R|^{d^2}$ and

$$\mathcal{F}_\psi F(A) = \chi(\det A)^{-1} G_d(\chi, \psi) F(I_d), \tag{18.8}$$

where the right hand side is zero for singular A.

Proof. A noninvertible square matrix B over the finite ring has a nonzero left kernel vector v^T: indeed its transpose as an endomorphism of a finite free module is not bijective, hence not injective. Let $v_i \neq 0$. By primitivity choose u with $1 + uv_i$ a unit and $\chi(1 + uv_i) \neq 1$. The matrix $g = I + ue_i v^T$ satisfies $gB = B$ and has determinant $1 + uv_i$, by the rank-one determinant identity. Covariance forces $F(B) = 0$. Thus $F(B) = \chi(\det B)F(I)$ for every B. For invertible A, the change of variable $C = BA$ in the Fourier transform, and cyclicity of trace, give (18.8). For singular A, there is a nonzero right kernel vector w; choose i, u as before and $g = I + uwe_i^T$. Since $Ag = A$, the change of variable $B \mapsto gB$ and covariance give $\mathcal{F}_\psi F(A) = \chi(\det g)\mathcal{F}_\psi F(A)$, again zero. Finally apply Plancherel to the function $B \mapsto \chi(\det B)$. Its squared norm is $|\mathrm{GL}_d(R)|$. The transform has squared norm $|G_d|^2 |\mathrm{GL}_d(R)|$ by the result just proved. Canceling the positive order of the group proves the magnitude identity. Equivalently,

$$\overline{G_d(\chi, \psi)}\, \mathcal{F}_\psi F(I) = |R|^{d^2} F(I).$$

□

This proves `matrix_gauss_sum` of `Lib/Weil/GaussSum/LocalHomUnit.lean`. Note that no identity involving the phase of the Gauss sum was needed. The distinction between magnitude and phase is crucial for the global scalar: phases may depend on the choice of character, whereas the square magnitude is a pure cardinality factor, of zero height cost at primes outside the twist conductor. The rank-one case is the usual Gauss sum of a primitive multiplicative character on a finite field. The matrix statement is equally true at a product of higher powers of finite residue rings under the stated primitivity condition; there is no hidden estimate by a power of the residue cardinality.

We next give the two exact finite summation identities behind the weighted square law. Take an ambient matrix coordinate decomposition of sizes (u, a_0, d), with the small side equal to a_0 and large side $u + d$. For $Y \in \mathrm{Mat}_{u \times a_0}(R)$, $V \in \mathrm{Mat}_{u \times d}(R)$, form the complementary frames

$$Y_* = \begin{pmatrix} Y \\ I_{a_0} \\ 0 \end{pmatrix}, \qquad X_* = \begin{pmatrix} I_u & -Y & V \\ 0 & 0 & I_d \end{pmatrix}. \tag{18.9}$$

Use the frame placement `swapFirstTwo` of the source: the small-side columns are the middle block, and the other two row blocks of the inverse are ordered as u, d. In this placement the endomorphism $k = \begin{pmatrix} I & Y & -V \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix}$ has the prescribed middle columns, and its inverse has the indicated first and third row blocks. Its determinant is one; this verifies the relative sign in the chosen ordering. Therefore by the orbit computation in Lemma 18.3,

$$|\operatorname{GL}_{a_0}(R)|L'(Y_*) = c|\operatorname{GL}_{u+d}(R)|L(X_*). \tag{18.10}$$

For example, choose a column datum supported on the orbit of Y_* with value 1 there, then its partner has the inverse-character values on the orbit of X_*; all the contributions cancel their determinant characters in the corresponding periods. The right side of (18.10) is independent of V, because a shear on the row frame of determinant one eliminates the upper right block V.

Given p on u by a_0 matrices define a new datum on $(u+d)$ by (a_0+d) matrices by

$$p^+\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{cases} p(-A)/|R|^{ud} & C = 0,\ D = I_d, \\ 0 & \text{otherwise.} \end{cases}$$

Multiplying (18.10) by $p(Y)$, summing over Y and averaging over V proves

$$|\operatorname{GL}_{a_0}(R)|\sum_Y p(Y)L'\begin{pmatrix} Y \\ I_{a_0} \\ 0 \end{pmatrix} = c|\operatorname{GL}_{u+d}(R)|\sum_X p^+(X)L\begin{pmatrix} I_u & X_{11} & X_{12} \\ 0 & X_{21} & X_{22} \end{pmatrix}. \tag{18.11}$$

This proves `SquareLaw.sum_eq_nextDatum` in `New/Exchange/SquareLawFourier.lean`, with the basis reorderings made explicit. The matrix X_{11} has size u by a_0, etc. Formula (18.11) remains valid without primitivity of χ, because it uses only covariance and the matched-pair relation.

Proposition 18.5 (Fourier version of the finite square)**.** *Assume $a_0 = u + d$ and that ψ, χ are primitive. Given a function p on matrices of shape $(u+d)$ by u, there is a scalar γ, independent of p, c, L, L', with $|\gamma|^2|R|^{d(d+2u)} = 1$, for which*

$$\begin{aligned} &\sum_{Y\in \mathrm{Mat}_{u\times a_0}(R)} \mathcal{F}_\psi(\bar p)(Y)L'\begin{pmatrix} Y \\ I_{a_0} \\ 0 \end{pmatrix} \\ &\quad = \gamma c \sum_{X\in \mathrm{Mat}_{(u+d)\times(a_0+d)}(R)} \mathcal{F}_\psi(\overline{p^+})(X)L\begin{pmatrix} I_u & X_{11} & X_{12} \\ 0 & X_{21} & X_{22} \end{pmatrix}. \end{aligned}$$

Here p^+ is the extension just defined, with the row and column sizes of p (transposed from those of Y).

Proof. Put $F = \bar p$. Fourier-transform $\overline{p^+}$ at a matrix $X = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$; the summation variable has shape (a_0+d) by $(u+d)$. It is restricted to matrices $\begin{pmatrix} S & T \\ 0 & I_d \end{pmatrix}$ with weight $F(-S)/|R|^{a_0 d}$. The trace pairing gives

$$\operatorname{tr}\left(\begin{pmatrix} A & B \\ C & D \end{pmatrix}\begin{pmatrix} S & T \\ 0 & I_d \end{pmatrix}\right) = \operatorname{tr}(AS) + \operatorname{tr}(CT) + \operatorname{tr} D.$$

By (18.7) the sum over T vanishes unless $C = 0$, when it cancels the factor $|R|^{-a_0 d}$. Hence

$$\mathcal{F}_\psi(\overline{p^+})(X) = \mathbf{1}_{C=0}\,\psi(\operatorname{tr} D)\,\mathcal{F}_\psi F(-A).$$

For fixed A, B, the function $D \mapsto L\begin{pmatrix} I & A & B \\ 0 & 0 & D \end{pmatrix}$ transforms by the determinant character under left multiplication of D; indeed left multiplication by $\operatorname{diag}(I, g)$ on the full frame has exactly this action. Proposition 18.4 gives the sum over D, with additive character $\psi(\operatorname{tr} D)$, as $G_d(\chi, \psi)$ times its value at the identity. That value is independent of B by the determinant-one shear already used in (18.10). The sum over B therefore supplies $|R|^{ud}$. The right hand sum in the proposition, before multiplication by γc, is precisely

$$|R|^{ud} G_d(\chi, \psi) \sum_A \mathcal{F}_\psi F(-A) L \begin{pmatrix} I & A & 0 \\ 0 & 0 & I \end{pmatrix}.$$

Apply (18.11) with p there equal to $\mathcal{F}_\psi F$; the two group orders cancel since $a_0 = u + d$. Averaging over the free upper right block in the extension merely gives the last sum. The desired scalar is

$$\gamma = \frac{1}{|R|^{ud} G_d(\chi, \psi)} = \frac{\overline{G_d(\chi, \psi)}}{|R|^{(u+d)d}},$$

which has the asserted squared magnitude by Proposition 18.4. □

This proves the identity `SquareLaw.sum_fourier_eq_nextDatum` of `New/Exchange/SquareLawFourier.lean`. It is considerably stronger than a triangle inequality on Fourier coefficients: the determinant Gauss sum kills singular lower blocks and gives the exact scalar. In the triangle of rank increments t, $d = t$, and u is the rank of the smaller common block; hence the exponent $t(t + 2u)$ is the difference between the products of the adjacent dimensions that appears when two kernels are reversed. The normalization of the semilinear operation of reversal in `exists_conj_modKernel_eq_swap` (`Lib/SiegelWeil/RallisInner/SharpData.lean`) gives a factor of squared modulus $|R|^{-n_1 n_2}$ for a single reversed edge of ranks n_1, n_2. That identity of genuine Weil kernels is retained as a formal-source input, but the Fourier calculation proves why the quotient of these factors in the square is a pure cardinality power. For $R = \mathbb{F}_q$, the exponent of q in the factor of the preceding proposition is exactly $-d(d+2u)$ in squared magnitude, including off-diagonal blocks. Omitting the Fourier factor at off-diagonal rectangles would change the relative metric in the main determinant calculation.

18.3 The graph term in doubling

To identify the analytic input in the Gram factors of Proposition 17.8, write S^- for the Hermitian space with form opposite to that of S, of dimension r, and T for a target of dimension $n > 2r$. The doubled source is the split Hermitian space $S \oplus S^-$. Its rational Lagrangians are the indices in the Siegel–Weil sum. They are not all graphs. Dropping the nongraph terms without explanation would be a serious omission: the finite support of the datum is essential to this simplification.

Lemma 18.6 (Graphs and the rank test)**.** *Let J be the matrix of a nonsingular Hermitian form on $V = E^r$, and let $L \subset V \oplus V^-$ be a Lagrangian. Represent it by the full-column-rank matrix $\binom{K}{K'}$. The equation of isotropy is $K^* J K = K'^* J K'$. If K' is invertible, then L is the graph of a unique isometry of V, namely $K K'^{-1}$. If V is definite at an archimedean place, every L is of this type.*

Proof. The determinant of K^*JK equals that of K'^*JK' and is nonzero if K' is invertible; hence K is invertible. With $\delta = KK'^{-1}$ the isotropy equation gives $\delta^*J\delta = J$, and the vector in L with prescribed second component v must be $(\delta v, v)$. For the last claim, if $K'c = 0$ for some nonzero c, isotropy implies $(Kc)^*J(Kc) = 0$. Definiteness implies $Kc = 0$, contradicting the full rank of $\binom{K}{K'}$. Uniqueness follows from the prescribed second projection. □

The Lean graph convention also permits a minus sign in the second component, which has the effect of replacing δ by $-\delta$. There is no second copy of the set of graphs caused by this convention. In indefinite signature the last argument has no force; for example an isotropic vector can occur with second component zero. For the complementary-rank exchange the following more general rank test is needed.

Formal-source input (zero terms). In `Lib/SiegelWeil/RallisInner/SharpData.lean`, `IsClean` is the alternative of definiteness or a finite named datum supported on matrices admitting a left inverse at a chosen split place. The theorem `exists_graph_eq_of_zeroTerm_ne_zero` states that for a clean datum a nonzero zero term of the double datum is attached to a graph. Its essential new assertion, beyond Lemma 18.6, is the vanishing when $K'c = 0$ and $Kc \neq 0$: invariance of the zero term under the corresponding Heisenberg translations enforces a Fourier support constraint incompatible with the primitive support at the chosen place. This analytic invariance and its compatibility with naming a finite datum are retained as formal-source input. The algebraic consequence is now identical to the proof of the lemma: K' is invertible, and the graph is unique. The same module's `zeroTerm_graphSplitting` identifies the surviving term, for datum f' on the first copy and f on the second:

$$\mathrm{ev}_{L_\delta}(\mathrm{double}(f', f)) = \langle f', \omega(\delta^{-1}, 1)f\rangle.$$

Here the double uses conjugation on the opposite source, as in that declaration. This specifies the metric normalization of this formula, not merely a proportionality.

Formal-source input (Siegel–Weil). The identity used with these graph terms is precisely `siegel_weil` of `Lib/SiegelWeil/Formula/SiegelWeilFormula.lean`. In the split source of rank $2r$ and with $n > 2r$, its mean over the target quotient is

$$\frac{1}{\mathrm{vol}[U(T)]}\int_{[U(T)]}\theta_\Phi(h)\,dh = \sum_L \mathrm{ev}_L(\Phi),$$

with summable zero terms. The absolute bounds needed also for variable Φ are developed in `Lib/SiegelWeil/Formula/SiegelSets.lean` and `Lib/SiegelWeil/Formula/ZeroTermHeight.lean`, including `exists_summable_zeroTerm_bound` and `theta_normalisedTorus_le`. The latter is a bound along a normalized expanding torus by an integrable function. The majorants, the independence of the evaluation from completing a Lagrangian to a splitting, and the formula itself are not reproved here.

Proposition 18.7 (Unfolding, conditional on the analytic inputs)**.** *Use Haar measures as in the preceding mean identity, and let $G = U(S)(\mathbb{A})$, $\Gamma = U(S)(E^+)$, with right-regular action $R(g)y(x) = y(xg)$. For rapidly decreasing automorphic functions y, y' and a clean f, the graph-term identity and Siegel–Weil identity give*

$$\langle \Theta_{f'}y', \Theta_f y\rangle = \mathrm{vol}[U(T)]\int_G \langle f', \omega(g,1)f\rangle\langle y', R(g)y\rangle\,dg$$

provided the integrals and sums may be interchanged absolutely.

Proof. Choose compatible measurable fundamental domains for the source quotient in each of the two lift integrals. Multiply the conjugate of the first theta lift and the second theta lift, and integrate over the target quotient. The inner target integral is the split double theta integral. Apply the mean formula, and discard nongraph terms by the clean-data criterion. Index the graphs by $\delta \in \Gamma$; the graph term for representatives x, x' of the two source variables is the matrix coefficient at $x'^{-1}\delta x$, after inversion of the indexing isometry if needed. Thus the result, divided by $\mathrm{vol}[U(T)]$, is

$$\int_{\Gamma\backslash G}\int_{\Gamma\backslash G} \overline{y'(x')}y(x)\sum_{\delta\in\Gamma} M(x'^{-1}\delta x)\,dx\,dx', \qquad M(g)=\langle f', \omega(g,1)f\rangle.$$

Left invariance of y under Γ identifies $\sum_\delta \int_{\Gamma\backslash G}$ with an integral over G. Substitute $x = x'g$ therein, using left Haar invariance, and interchange these last two integrals. The integral in x' is $\langle y', R(g)y\rangle$, as required. □

The convergence condition is not optional. In `SharpData.lean`, `inner_product_formula` establishes the integrability of exactly this coefficient product and proves this identity, with rank hypothesis $n > 2r$ and clean f; its proof uses the rational norm bound `exists_rat_norm_le_height` for the pair kernel, the rapid bounds on y, y', and the zero-term summability in the Siegel–Weil formula. We take those analytic estimates and all measure identifications as formal-source input. The unfolding argument above isolates the part that is independent of the estimates. It also explains why the final integral is over the *group* of the source, not the source quotient. This distinction will reappear as a sum over infinitely many cosets in the Euler expansion; replacing it by the finite automorphic class set would change the formula.

18.4 Sharp support, level factors, and the common scalar

Write $R = \mathcal{O}_E/\mathfrak{m}$ for a finite ring assembled from the selected split places, with level one at each prime, and let

$$Z = \mathrm{Mat}_{n,r}(R), \qquad F : Z \longrightarrow \mathbb{C}, qquad \mathrm{supp}\, F \subseteq \{Y : LY = I_r\},$$

where L is fixed of size $r \times n$. This is a useful particular case of the primitive support of the preceding subsection. Our local vacuum conventions are those of the named finite datum in the compact source.

Lemma 18.8 (Support under column action)**.** *If $A \in \mathrm{GL}_r(R)$ and $F(Y)F(YA) \neq 0$ for any Y, then $A = I$. Over a local DVR with fraction field K, if Y and YA are both primitive integral n-by-r matrices for $A \in \mathrm{GL}_r(K)$, then $A \in \mathrm{GL}_r(\mathcal{O}_K)$.*

Proof. In the first case $LY = LYA = I_r$, so $A = I_r$. In the second case there are integral left inverses P, P' of Y, YA, respectively. Then $A = P(YA)$ is integral, and $A^{-1} = P'Y$ is integral too. □

Consequently the overlap of two such compact finite data in the matrix coefficient for the source vanishes outside the principal level subgroup at the selected places. On that subgroup the right translation and determinant twists reduce to the identity. Counting using the normalized Haar measure on integral matrices gives

$$\|f\|_{\mathrm{fin}}^2 = \frac{\sum_{Y\in Z}|F(Y)|^2}{|R|^{nr}}, \qquad c = \frac{\|f\|_{\mathrm{fin}}^2}{|\mathrm{GL}_r(R)|}.$$

Indeed, the reduction of the integral matrix lattice has equally measured fibers, whereas the principal kernel in the integral group has measure the reciprocal of the order of the reduction group when the maximal compact group has measure one. There are two independent denominators. They have quite different origins, and conflating them removes a large dimension-dependent factor.

Formal-source input (sharpness and Euler expansion). These finite arguments are the support and counting core of `isSharp_modDatum` in `Lib/SiegelWeil/RallisInner/SharpData.lean`; the named model action, the Hilbert norm identity for that model, and the passage from split residue factors to genuine Weil coefficients are taken from that declaration. Its hypothesis includes definiteness or nonempty selected places, so that clean support is available. The exact coset expansion and injectivity needed in the application are `rallis_inner_product_euler` in `Lib/Weil/CompactLift/RallisFormula.lean`: for a fixed real edge K and $n > 2r$ there is a common $A \neq 0$ such that for weighted characters a, b, sharp finite datum with parameter c, and a form y of level J, the lift $\mathcal{L}$ satisfies

$$\mathcal{L}y = 0 \implies y = 0, \qquad \langle \mathcal{L}y', \mathcal{L}y \rangle = \sum_{\gamma} A c \,\mathrm{cosetTerm}(n, b, S_0; y', y; \gamma).$$

The sum is a convergent `HasSum`. Here γ indexes cosets in the adelic group integral and S_0 denotes the selected bad places in the sharpness condition. In particular A does not depend on the finite sharp datum or the varying characters for that real edge. The injectivity assertion is not a consequence of the positive norm of the identity term alone: other coset terms can have phases of either sign. The nonzero pairing of the real edge and the integral nondegeneracy argument in this declaration are used as formal-source input.

A common source of confusion is that a finite matrix norm and a Petersson norm use different measures. The effect can be seen without opening the entire analytic proof. At the auxiliary edge of rank $r \longrightarrow r + t$, let $\mathcal{Q}$ be the selected twisting places and M the fixed matching places. Set

$$N(\mathcal{Q}) = \prod_{w \in \mathcal{Q}} q_w, \quad \gamma_{\mathcal{Q}} = \prod_{w \in \mathcal{Q}} |\mathrm{GL}_r(k_w)|, \quad R = \prod_{w \in M \cup \mathcal{Q}} k_w.$$

At M, the auxiliary indicator fixes a standard inclusion matrix; at $\mathcal{Q}$ it fixes the designated r rows to the identity, leaving tr free entries. The auxiliary datum multiplies this indicator by $\gamma_{\mathcal{Q}}$.

Proposition 18.9 (Auxiliary count)**.** *With this datum F, the sharp parameter is*

$$c = \frac{\gamma_{\mathcal{Q}}^2 N(\mathcal{Q})^{tr}}{|\mathrm{GL}_r(R)|\,|R|^{(r+t)r}}.$$

Proof. For $w \in M$ there is one allowed matrix; for $w \in \mathcal{Q}$ there are q_w^{tr} matrices. Since the factors form a Cartesian product, the number of allowed data is $N(\mathcal{Q})^{tr}$. Every nonzero value of F equals $\gamma_{\mathcal{Q}}$. Apply the two denominators computed above. □

The finite identity is `AuxPair.levelIndicator_support_and_normSq` in `New/Exchange/CharacterLiftNorm.lean`. To compare algebraic Gram forms, multiply the coefficient by the level index $[K_0 : J']$ on the target. The remaining common real factor contains the cardinality of the Weyl index at that target, and is independent of the character; the precise

conversion is contained in that module's `Data.auxLift_injective_and_hasSum_inner`. The result, in the notation actually used later, is

$$d = [K_0 : J'] \frac{\gamma_{\mathcal{Q}}^2 N(\mathcal{Q})^{tr}}{|\mathrm{GL}_r(R)|},$$
$$R_j(\gamma) = \frac{\mathrm{cosetTerm}(r+t, b, S_0; x, y; \gamma)}{|R|^{(r+t)r}},$$

and the sum of $AdR_j(\gamma)$ is the auxiliary algebraic Gram pairing, for one $A \neq 0$ common across the twists. These are exactly `AuxPair.gramFactor` and `Data.rallisTerm`. In particular one cannot move the level index into the summand and then also multiply by it when converting from Petersson to algebraic norms.

There is a useful one-prime example. Let $r = 1$, $t = 3$, $R = \mathbb{F}_q$, and let the only selected place be a twisting place. The allowed column vectors are $(1, z_1, z_2, z_3)^{\mathrm{t}}$; there are q^3, with $F = q - 1$. Then

$$c = \frac{(q-1)^2 q^3}{(q-1)q^4} = (q-1)/q.$$

The finite norm alone is $(q-1)^2/q$, which differs by $q - 1$. This illustrates the independent GL-group factor. For a fixed matching place the corresponding free-entry count is absent, but both denominators remain. At higher rank the GL-group cardinality changes much faster with rank than one could infer from any rank-one formula.

18.5 A rank-one Euler factor test

The auxiliary count determines the dependence on the support of a finite datum, but does not explain the appearance of the desired Hecke L-value. For the latter, consider $r = 1$, $n = 2k$ and $k > 1$. At a split unramified prime of the real base, the local unitary group is identified with $E_w^{\times}$ and its compact-unit quotient is indexed by integers. Let q be the common residue size, let $a = q^{-k}$, and let λ have modulus one. Then

$$\sum_{j \in \mathbb{Z}} a^{|j|} \lambda^j = 1 + \frac{a\lambda}{1 - a\lambda} + \frac{a\lambda^{-1}}{1 - a\lambda^{-1}} = \frac{1 - a^2}{(1 - a\lambda)(1 - a\lambda^{-1})}.$$

The positive and negative valuations give the two Euler factors in the denominator. The numerator corrects for counting the valuation zero twice. At nonsplit places the unitary line group is compact and has a single coset modulo its maximal compact group. This simple check explains the split and nonsplit factors of the line calculation. The actual identification of its character factors, and the unramified vacuum coefficient equal to the required character times $q^{-k|j|}$ in this rank, is the formal-source input `exists_hasSum_vacuumCoeff_at_place` in `Lib/SiegelWeil/RallisInner/SharpData.lean`. The pair of factors $1 - \varepsilon_w q^{-k}$ at a split conjugate pair has opposite signs $\varepsilon_{\bar{w}} = -\varepsilon_w$ and thus product $1 - q^{-2k}$.

For completeness, in this absolutely convergent range Euler products with coefficients of norm at most one cannot vanish. If z_w has norm at most q_w^{-k}, the expansion $\log(1 - z_w) = -\sum_{l \geq 1} z_w^l / l$ has absolute sum bounded by $|z_w|/(1 - |z_w|)$. The sum of the latter over places converges: over each rational prime the number of places is bounded by the field degree and the residue size is at least the rational prime. Therefore exponentiating the sum of logarithms gives both a nonzero product and a convergent product of reciprocals. This argument does not apply at $k = 1$ without hypotheses and further analysis.

Formal-source input (the rank-one norm identity). In `New/Exchange/CharacterLiftNorm.lean`, `LValueData.inner_trivialLift_self_eq_L` proves, for $n = 2k$ with $k > 1$,

$$\langle f_\chi, f_\chi \rangle = C \prod_{w \in M} (q_w - 1)\ H_\chi\ d_\chi L^S(k, \mu^{-1}) N(\mathcal{Q})^{n-1}.$$

Here f_χ is the trivial-source auxiliary lift, H_χ is the hook-number factor for its Weyl lattice, d_χ the Weyl-index cardinality, S the omitted places, and μ the weighted character occurring in the declaration. The factor $C \neq 0$ is given there by the product of two quotient volumes divided by the real-zeta Euler product at k. This is a norm formula for the particular real and finite datum, not an identity for arbitrary harmonic vectors. The local coset calculation, the rank twist in the character and the comparison of the real Fock and algebraic Gram forms are the formal ingredients of that theorem. The finite count above and the elementary geometric-series calculation explain the shape of its conductor power and Euler factors. It is especially important that the displayed identity involves $L^S(k, \mu^{-1})$ with the character prescribed by the weighted lift, rather than its absolute square: the base change in the Euler factor computation has already paired the appropriate local factors.

18.6 Hyperbolic-plane removal and the signed product

We now pass from the absolutely convergent range of the Rallis calculation to the distinct argument used to compare theta integrals of complementary rank. Adding a hyperbolic plane to both source and target changes the source and target dimensions by two each. To understand the scalar in this step, introduce

$$P_\nu(e) = N(S)^{e^2} \prod_{j=1}^{e} |L^S(j, \nu)|^2 \quad (e \geq 0), \qquad \Pi_\nu(e) = \begin{cases} P_\nu(e) & e \geq 0, \\ P_\nu(-e)^{-1} & e < 0. \end{cases}$$

The second definition agrees with the source's `lValueProdInt` of `Lib/SiegelWeil/Tower/LargeSideConstantTerm.lean`. The positivity of the factors under the actual weight hypotheses is the result `lValueProdInt_pos`; its proof uses the weighted-character nonvanishing theorem for the involved values, not general nonvanishing on a hypothesized critical line. In particular we work here on a positive part of the theory and do not assume GRH to define the reciprocal.

Lemma 18.10 (The transition quotient)**.** *For $e \in \mathbb{Z}$, let $b = e + 1$ if $e \geq 0$, and $b = -e$ if $e < 0$. Then*

$$\Pi_\nu(e+1)/\Pi_\nu(e) = N(S)^{2b-1} |L^S(b, \nu)|^2.$$

Proof. For $e \geq 0$ divide the two positive products; the exponent increment is $(e+1)^2 - e^2 = 2e+1$. For $e < 0$, put $b = -e \geq 1$. The inverse product at $e+1$ is the inverse of the product at $b-1$, including $b = 1$. Dividing by the inverse at b gives $P_\nu(b)/P_\nu(b-1)$, which has exponent $b^2 - (b-1)^2 = 2b - 1$. □

In particular $\Pi_\nu(-1) = 1/(N(S)|L^S(1, \nu)|^2)$ and both transitions from -1 to 0 and from 0 to 1 have quotient $N(S)|L^S(1, \nu)|^2$. This gives a check on the signed-exponent convention at the equal-rank boundary. If one used the unsigned product at negative integers, the first transition would have reciprocal scalar, breaking the exchange relation. More generally the quotients telescope without any case-dependent residual factor.

Let $W' = \mathbb{H} \oplus W$ and $V' = \mathbb{H} \oplus V$ with base-adapted splittings. For $n = \dim W$ even and $r = \dim V$, let $e = r - n/2$. At the larger pair the index is $(r+2) - (n+2)/2 = e + 1$. The constant term in a line radical is governed by the isotropic cone of V'. Here is the elementary mechanism of the constant-term reduction.

Lemma 18.11 (Orthogonality on the line radical)**.** *Let A be a compact abelian group with normalized Haar measure. If $\eta : A \to \mathbb{C}^\times$ is a continuous unitary character, then its mean is one for the trivial character and zero otherwise.*

Proof. The trivial case is the mass of A. Otherwise choose a_0 with $\eta(a_0) \neq 1$ and translate the integral by a_0 to obtain $I = \eta(a_0)I$. □

In the adapted mixed model, the line radical acts on the theta summand indexed by a rational vector $x \in V'$ by the character attached to the scalar $\langle x, x\rangle$. Integrating over the compact quotient of the line radical kills every summand with $\langle x, x\rangle \neq 0$. The model action and the interchange with the theta summation are formal-source inputs `constTermAt_eq_integral` and `hasSum_constTermAt`, in `Lib/SiegelWeil/Tower/ConstantTerm.lean`. The surviving rational cone decomposes into the vertex and one nonzero orbit: by Witt extension, any two nonzero isotropic vectors are conjugate under the rational unitary group. Indeed, the map of their one-dimensional subspaces taking one vector to the other is an isometry for the restricted zero form, so Witt extension provides the isometry of the ambient nondegenerate space. The stabilizer of the isotropic line has Levi $E^\times \times U(V)$; its action on the chosen vector is through the first factor, and the stabilizer of the vector is the kernel of that scaling.

The exponent $r + 1$ in the measure from unfolding this orbit can be checked in a Witt basis for V'. The radical is a two-step group with quotient an E-vector space of dimension r and central coordinate in a one-dimensional skew-Hermitian E^+-space. Conjugation by the line-scaling element $\mathrm{diag}(a, 1, \bar{a}^{-1})$ multiplies the first block by a scalar of E-absolute value $|a|_E$ and the second by $a\bar{a}$. This has Jacobian $|a|_E^r |a|_E = |a|_E^{r+1}$ under compatible normalized absolute values. Unfolding the nonzero orbit, when the vertex contribution vanishes, therefore introduces a multiplicative line integral over ideles modulo $E^\times$, an integral over $[U(V)]$, and an integral over the maximal compact of the larger space. The source supplies this precise normalized unfolding, including integrability and the nonzero measure constant, in `exists_unfold_constant_term` of `Lib/SiegelWeil/Tower/ConstantTerm.lean`.

Formal-source input (the normalized one-plane identity). The full formula `exists_constant_term_thetaLarge`, in `Lib/SiegelWeil/Tower/LargeSideConstantTerm.lean`, chooses compatible rapid datum and local row datum, and compares the constant term on W', V' with the smaller integral on W, V evaluated on the real cone-integral datum. It gives a scalar $\gamma \neq 0$, for one common constant C, with

$$|\gamma|^2 \Pi_\nu(e) = C \Pi_\nu(e + 1).$$

The equality holds on the indicated Levi variables when the line-scaling variable is one, or the smaller ranks are equal. This proviso matters in using the formula; there is no assertion here of full arbitrary-variable equality merely from a constant term. The local integrations use the adapted frames and the line datum constructed by `exists_lineTerm_localDatum`; the character factors are those of the declaration, and the real cone integral is also prescribed there. In particular C is independent of the twist. By Lemma 18.10 this scalar has squared norm $CN(S)^{2b-1}|L^S(b, \nu)|^2$. The single $N(S)$-power accompanying $e = 0$ or $e = -1$ and the growth of the exponents thereafter arise from the full conductor and line Gauss-sum calculations, not a guess based on dimensions.

18.7 What makes an exchange an identity on all partners?

The preceding line calculation proves neither the complementary-rank identity nor its square version in isolation. There are two more requirements. One must arrange a nonzero common

starting ratio at equal rank and prescribe the admissible real datum after removing planes. One must also show that two theta integrals, whose difference has appropriate vanishing Fourier coefficients, agree on the full space of matched data. A function on a group can have zero constant term without being zero, even if it is automorphic.

For a space X of even rank n the input ratio is

$$\ell_\nu(v) = \frac{|\lambda_\nu(v)|^2}{\Pi_\nu(n/2)}, \qquad \lambda_\nu(v) = \theta_{\mathrm{large}}(X, X; \nu, v, F_{\mathrm{std}})(1).$$

It is `equalRankThetaRatio` of `New/Exchange/LowRankIdentity.lean`; F_{std} is the indicated standard-column datum. The theorem `exists_equalRankThetaRatio_value` of `New/Exchange/EqualRankValue.lean` constructs a rapid real datum on one hyperbolic enlargement of the input W of rank $2k + 2q$, with q its number of negative signs. For a character weight componentwise at least k, its equal-rank ratio is a common nonzero number ℓ, independent of ν. This is a formal-source input: it requires the nonvanishing for a suitable real datum and consistency of the local comparisons as planes are removed, not simply the positivity of Π_ν.

For the chosen matched rank pairs W, V and U, W, `exchange_of_partners` of that module takes rapid test data z, y, with z measured at the larger equal-rank pair, and imposes the signature equality of complementary ranks. Its conclusions are: there are a common constant C and a single rapid datum v, and for each split tame character a scalar c, such that

$$C|c|^2 = \Pi_\nu(\dim V - \tfrac{1}{2}\dim W)\ell_\nu(z), \tag{18.12}$$

and the two global theta integrals on partners agree after scaling by c. With the equal-rank datum chosen above, the right-hand side is positive, so C and c cannot vanish. Dividing (18.12) for two twists therefore cancels all real constants:

$$\frac{|c_\nu|^2}{|c_{\nu_0}|^2} = \frac{\Pi_\nu(\dim V - \frac{1}{2}\dim W)}{\Pi_{\nu_0}(\dim V - \frac{1}{2}\dim W)}.$$

This is the precise scalar feature of the complementary-rank exchange in Proposition 17.8.

Formal-source inputs in the passage to matched data. For the real data, the cone integration operator must reach prescribed determinant vectors. The needed surjectivity is `surjOn_cone_integral_detVectors`, `New/Exchange/DeterminantVectors.lean`. On the finite side the difference of the integrals descends through the linear space of matches and is equivariant, by `exists_isEquivariantLift` in `New/Exchange/ExchangeDifference.lean`. Then `IsEquivariantLift.apply_stdColDatum_eq_zero` in `New/Exchange/StandardCoefficients.lean` promotes vanishing on the specified standard coefficients: its hypothesis is a family of centre-coefficient vanishings at adapted frames, not just a single constant-term vanishing; its conclusion is vanishing on all the integral standard-column data. These declarations and their singular-coefficient and approximation arguments are retained as formal-source input. The finite partner identities and the signed-product algebra above are necessary for their application; they do not replace them. This distinction is especially important when the equal-rank target is no longer definite.

The upgrade from an exchanged period to a seesaw square in `New/Exchange/OneSquare.lean`, `one_square`, uses continuous linearity in a real rapid datum. We give two elementary points behind that upgrade.

Lemma 18.12 (A rank-one span closed under pointwise limits)**.** *Let I be any set and $h : I \to \mathbb{C}$. If a net of functions of the form $c_i h$ converges pointwise to f, then $f = ch$ for some scalar c.*

Proof. If $h = 0$, all the functions are zero. Otherwise pick $j_0 \in I$ with $h(j_0) \neq 0$. Pointwise convergence implies $c_i \to f(j_0)/h(j_0)$. Evaluating at any other $j \in I$ and passing to the limit proves the assertion. The same argument applies to the convergence of sequences or to the closure in the product topology. □

There is no finiteness assumption on I in this lemma; one may group many finite twisting families into one family of indices. In the seesaw argument, a real datum is first approximated by linear combinations of elementary tensors. The scalar direction for such approximants must be common to every index. A different scalar for each index would leave no useful rank-one assertion. The lemma shows why the identity then passes to the limit; it does not assert that the limiting scalar is nonzero.

Lemma 18.13 (Uniform bound for the integrated kernel)**.** *Suppose Y, Y', A have finite quotient measure and nonnegative height functions at least one. Let $K_v(x, x', z)$ be a measurable kernel depending linearly on a real datum v, with*

$$|K_v(x, x', z)| \leq C\|v\|_N H_Y(x)^B H_{Y'}(x')^B$$

uniformly in $z \in A$. If y and y' are rapid, then multiplication by $y(x)y'(x')$ and by any unit-modulus character factor gives an integrable kernel, with L^1 norm at most a constant times $\|v\|_N$.

Proof. There exist constants a, a' with $|y(x)| \leq aH_Y(x)^{-B}$ and $|y'(x')| \leq a'H_{Y'}(x')^{-B}$. Thus the magnitude of the integrand is bounded by $Caa'\|v\|_N$, independent of all three quotient variables. Integrate this constant over the product of the three finite measure quotients. This also proves continuity of the resulting integral in the stated seminorm. □

The actual kernel bound and the passage from the height of a pair to the product of individual heights are formal-source inputs in `exists_integrable_kernel_bound` of `OneSquare.lean`. The proposition `one_square` combines this bound, the density of elementary tensors in the real-vector space with the prescribed weighted norms, the seesaw identity, and the partner comparison. Its finite consequence is precisely a square law with factor κ/c. The auxiliary scalar κ may initially be zero; c is the nonzero scalar of (18.12). The nonzero square scalars in the weighted-lift triangle are proved from injectivity and the properly chosen kernels afterward. The arithmetic estimates can divide by c, but not automatically by κ.

18.8 Level orders in the paths

The two-grid path from the starting periods to the terminal auxiliary pair is not an arbitrary path in a rank diagram. A move changes real data, local matrices and level characters. Proposition 17.11 furnishes the correctly synchronized path. It would be insufficient to find a formal sequence of dimensions of the same cost: each move must take a standard matched datum to another one in the family, and its resulting lift must be related by the square law. The preservation of the tidiness condition in `IsTidy.exists_square_move`, `New/Exchange/GridMove.lean`, remains formal-source input. The local square calculations above give the finite, cardinality-sensitive part of such moves.

The index bookkeeping of the triangle and the final comparison level can nevertheless be checked explicitly. For a block parabolic $P(n_0, \ldots, n_m)$ of $\mathrm{GL}_n(\mathbb{F}_q)$, where $n = \sum_i n_i$, block upper-triangular order gives

$$|P(n_0, \ldots, n_m)| = q^{\sum_{i<j} n_i n_j} \prod_i |\mathrm{GL}_{n_i}(\mathbb{F}_q)|.$$

Indeed the diagonal blocks must be invertible, all strictly upper blocks are arbitrary, and every such matrix is invertible. This holds also if some blocks have dimension zero, with the convention $|\mathrm{GL}_0| = 1$. The same counting argument as in the proof of Proposition 18.4 gives $|\mathrm{GL}_n(\mathbb{F}_q)| = \prod_{i=0}^{n-1}(q^n - q^i)$ by selecting independent columns.

For blocks $r, t, \ldots, t$ with m copies of t set

$$k(m) = mrt + \binom{m}{2} t^2.$$

This is the number of entries strictly above the diagonal blocks. Consequently $k(1) = rt$ and

$$m(k(1) - rt) - k(m) = -mrt - \binom{m}{2} t^2.$$

It is permissible to use a negative integer exponent in an identity of level factors, since residue-ring orders are nonzero integers. Calling this exponent a negative group order would be an error: it arises from expressing a quotient of known positive cardinalities.

To illustrate the count, take $r = t = 1$, $m = 2$ over $\mathbb{F}_3$. The parabolic has order $3^3 \cdot 2^3 = 216$, while $\mathrm{GL}_3(\mathbb{F}_3)$ has order $(27-1)(27-3)(27-9) = 11232$. Its index is 52, the number of full flags of lines inside planes: there are 13 lines and four planes containing any fixed line. At rank two the parabolic order is $3 \cdot 2 \cdot 2 = 12$ and its index is four. This is why naive powers of the first-step index cannot serve as the last-step index.

The full adelic index at a level given by inverse images of subgroups modulo split primes is the product of their finite indices, with the fixed level factors separated off. To see this, use surjectivity of reduction at each of the specified split primes and the Cartesian product of local maximal compacts; the quotient by the inverse image of the product subgroup has cardinality the product of the finite quotient cardinalities. In the source let $w(j)$ denote the `triangleOrderfalse` at stage j and J_m the level `lastResLevel`. Then `Data.stdLevelIndex_lastResLevel_eq` in `New/Exchange/ChainNorms.lean` proves

$$[K_0 : J_m] = C\,|R_0|^{m(k(1)-rt)-k(m)} d^m \frac{w(m)}{w(0)w(1)^m},$$

where R_0 is the residue ring at the initial varying modulus, d is the first auxiliary `gramFactor`, and C is independent of the twist. Here the symbol $k(j)$ is the number of strictly upper-block entries calculated above. This is precisely the elementary block count used in that declaration (via `card_parabolic_blockNumber`); the identification of the actual inverse-image level subgroups in the triangle is retained from the source. An advantage of this expression is the cancellation of the triangle-order quotient when it is combined with the weighted-chain norm. The powers of d and $|R_0|$ must still be tracked separately.

18.9 Why a false lift with one common root contributes no period

The purpose of comparison forms in Section 17.8 is not just to detect the prescribed eigensystem. They must be insensitive to certain other eigenspaces even when one root of their Hecke polynomial coincides with a root of a genuine lift. We explain an important mechanism that makes this possible. On an adelic quotient an arbitrary Hecke annihilator does not imply a zero function. In a sufficiently small argument of the Hecke polynomial, however, the operator cannot kill a nonzero bounded spherical function.

Lemma 18.14 (The subspace binomial identity)**.** *Let $a_{m,d}$ denote the number of d-dimensional subspaces of $\mathbb{F}_q^m$. Then*

$$\sum_{d=0}^{m} a_{m,d} q^{d(d-1)/2} T^d = \prod_{j=0}^{m-1} (1 + q^j T).$$

Proof. Fix a hyperplane in $\mathbb{F}_q^m$. Subspaces of dimension d contained in it are counted by $a_{m-1,d}$. A subspace not contained in it intersects the hyperplane in a $(d-1)$-dimensional subspace U. In the quotient by U it is a line with nonzero projection onto the one-dimensional quotient by the hyperplane. Choose its unique representative projecting to 1; its remaining coordinates have q^{m-d} possibilities. Hence

$$a_{m,d} = a_{m-1,d} + q^{m-d} a_{m-1,d-1}.$$

Multiplying by $q^{d(d-1)/2} T^d$ and summing, the second term is $q^{m-1}T$ times the polynomial in dimension $m-1$. Induct from the polynomial one in dimension zero. □

Proposition 18.15 (A Hecke injectivity bound for bounded functions)**.** *Let F be a bounded spherical function on an adelic unitary group of rank m at a split place w, with $q = q_w$. If the unnormalised Hecke polynomial operator at $X \in \mathbb{C}$ kills F and*

$$q^{m+1}|X| \le 1,$$

then $F = 0$.

Proof. Use the split Hecke-coset expression for spherical vectors, with the normalization of the source: the terms are indexed by all subspaces U of the m-dimensional residue vector space, and have coefficients $(-X)^{\dim U} q^{\dim U(\dim U - 1)/2}$. Each term is the appropriate translate of F. The term of dimension zero has coefficient one; its translate has the same supremum norm as F (and is the unshifted term with the identity representative). All other translates have supremum norm at most $\|F\|_\infty$. Thus

$$\|F\|_\infty \le \left(\prod_{j=0}^{m-1} (1 + q^j |X|) - 1 \right) \|F\|_\infty,$$

by Lemma 18.14. Using $1 + u \le e^u$ for $u \ge 0$,

$$\prod_{j=0}^{m-1} (1 + q^j |X|) \le \exp\left(|X| \frac{q^m - 1}{q - 1} \right) \le \exp\left(\frac{1}{q(q-1)} \right) \le e^{1/2} < 2.$$

The coefficient on the right is strictly less than one. A finite nonnegative supremum norm satisfying this inequality must be zero. □

Apart from the spherical coset description (`heckePolyOp_eq_sum`), this is a complete proof of `eq_zero_of_heckePolyOp_eq_zero` in `New/SiegelWeil/DescentsVanishPlanes.lean`. It is useful to view the argument as a Neumann-series estimate in a space of bounded functions. It does not require decomposing the bounded functions into automorphic representations, a Ramanujan bound, or the existence of a discrete spectrum on an indefinite group. It requires boundedness, which must be supplied independently.

In descending through a theta tower, write r for an initial rank, $r' = r + 2p$ for the rank after adding p hyperbolic planes, and $n > r'$ for the rank of the fixed definite target. The relevant Hecke root on the latter has magnitude $q^{(r+2-2n)/2}$. In the transfer from rank n

to rank r' the polynomial argument is rescaled in magnitude by $q^{(n-r')/2}$; the additional trivial-character factors must be nonzero. The rescaled argument satisfies

$$|X'| = q^{1-n/2-p}, \qquad q^{r'+1}|X'| = q^{r+p+2-n/2}.$$

Thus the integer hypothesis $n > 2r + 2p + 3$ gives exactly the small-argument condition in Proposition 18.15. Indeed the other, trivial Hecke factors have argument norms bounded by $q^{l+r'/2}|X|$ for $0 \le l < n - r'$, and the largest such bound equals $q^{-p} < 1$ when $p > 0$. The exact transfer and its conjugate-place adjunction are `hecke_transfer_of_equivariant` and `heckePolyOp_adjoint` in `Lib/Weil/Unramified/HeckeTransfer.lean` and `Lib/Weil/Unramified/SphericalHecke.lean`, respectively. These operator identities are formal-source inputs. Their normalization, and in particular the half-integral dimension shift in the argument, is essential to the strict bound.

Formal-source input (the tower vanishing). The initial zero auxiliary descent and reflection-sum annihilators imply general descent vanishing at that rank by `AuxPair.descentsVanish_of_reflectionSum`, `New/SiegelWeil/AuxPairVanishing.lean`, provided the local twists at the twisting places are nontrivial. This is a statement about all vacuum-off-level data, not only the one auxiliary datum. The induction theorem `descentsVanish_add_planes`, `New/SiegelWeil/DescentsVanishPlanes.lean`, then gives vanishing for data at the enlarged source: the inductive constant term vanishes; the archimedean and adelic estimates supply boundedness, and the Hecke operator is transported to the larger source in the preceding normalization. The case of adding zero planes needs no contraction estimate. For the other cases the full boundedness argument and the transport through each adapted splitting remain formal-source input. The just-proved proposition supplies the last step once boundedness and the Hecke equation are established.

From descended functions to comparison rows. The theorem `descentsVanish_hypSum` in `New/Rows/MemberOrthogonality.lean` applies the tower result to each conjugate matrix coefficient of an algebraic form killed by the auxiliary descent, the reflection sums, and the selected Hecke polynomial. It requires the inequality $2r + 2p + 3 < n$. The theorem `stdClassSum_period_eq_zero` in that module converts this vanishing into zero pairing with the comparison periods arising from the divided lattice. Its additional level, embedding and equivariance hypotheses are needed in this conversion; these two declarations are formal-source inputs. Notice the logical difference between sharing a single test root and possessing the entire Hecke eigensystem of a genuine lift. The former is simply enough to impose the operator equation. A false lift in its nullspace must be removed by the auxiliary descent condition and the nontrivial twists. It is not enough to compare roots as numbers.

18.10 Two finite ingredients of integral comparison

The rank and period arguments have been over $\mathbb{C}$; the construction of comparison forms also depends on p-adic integrality. We isolate two algebraic estimates behind Proposition 17.17. For the first, let $e : [a] \hookrightarrow [N]$ be an injection, write i_e for its column inclusion matrix, let $f : R \to S$ be a ring map, and let w be a matrix on which right multiplication by N-by-N matrices over S is defined. Put

$$H = \{g \in \mathrm{GL}_N(R) : g i_{e^c} = i_{e^c}, \quad w f(g) = w\}.$$

Use a determinant character χ extended by zero to nonunits of R and define

$$C(s) = 1_{wf(s)=wi_e}\, \chi(\det(i_e^{\mathsf{t}} s)), \qquad s \in \mathrm{Mat}_{N,a}(R).$$

Lemma 18.16 (A row-stabilizer cut)**.** *Suppose F is a function on matrices invariant under left multiplication by H. For any square matrix z and appropriately sized y,*

$$C(zi_e)1_{zi_{e^c}=i_{e^c}}F(y)=\begin{cases}\chi(\det z)F(z^{-1}y), & z\in H,\\ 0, & z\notin H.\end{cases}$$

Proof. Outside the fixed-column condition both sides vanish. Otherwise order the coordinates with the image of e first. Then $z=\begin{pmatrix}A & 0\\ B & I\end{pmatrix}$; $i_e^{\mathsf{t}}zi_e=A$, so $\det z=\det A$. Moreover

$$z=I+(z-I)i_ei_e^{\mathsf{t}},$$

whence $wf(z)=w$ is equivalent to $wf(zi_e)=wi_e$ (inclusion matrices are defined over the prime ring and unchanged by f). The determinant factor is zero unless z is invertible, in which case the two conditions defining H are now exactly those on the left. For $z\in H$, $z^{-1}\in H$ too, so invariance gives $F(z^{-1}y)=F(y)$. □

This proves the finite identity `columnDatum_mul_ite_eq` in `New/Rows/PartnerCut.lean`. It is the ingredient ensuring that a cut along selected columns enforces the required projector on integral data rather than introducing new, unsynchronized level characters. The group of fixed complementary columns by itself is too large; keeping the equation involving w is essential.

For the denominator estimate write $|\cdot|_p$ for an absolute value normalized by $|p|_p=p^{-1}$. The Schur lattice of a Young diagram has weights read from tableaux, and each entry corresponds to one occurrence of the associated standard basis coordinate. The essential combinatorial restriction is strict increase down every column of a semistandard tableau. If all entries belong to an alphabet of size l and the diagram has b nonempty columns, then at most b occurrences of any one alphabet symbol can occur. Consequently at most $b|S|$ entries can be drawn from a specified subset S of the alphabet.

Lemma 18.17 (Cost of bad columns)**.** *Let A be an l-by-l matrix over a p-adic field equipped with its Weyl lattice for a Young diagram as above. Suppose multiplying columns in a set S by an integer $M\neq 0$ makes A integral, all other columns already being integral. Then, with the lattice sup norm and the Schur action ρ,*

$$\|\rho(A)v\|_p\le |M|_p^{-b|S|}\|v\|_p.$$

Proof. Set $D=\operatorname{diag}(d_i)$ with $d_i=M$ on S and 1 outside. Then AD is integral. The integral polynomial Schur action preserves the integral lattice on all integral matrices, so has norm at most one for the corresponding lattice norm. The diagonal action $\rho(D^{-1})$ acts on the tableau weight space by M^{-u} where u counts entries from S. Since $u\le b|S|$ and $|M|_p\le 1$, its norm is at most $|M|_p^{-b|S|}$. Apply multiplicativity to $A=(AD)D^{-1}$. □

One must identify the bad columns before using the lemma. In the divided lattice arising from gluing, `norm_schur_inv_cols_le_weylGauge` of `New/Rows/CartanIntegrality.lean` chooses an integral change of basis on the columns whose reduction is specified modulo the squared chosen modulus. In that basis all but the indicated set of inverse columns are integral. The bound then has one factor per archimedean weight embedding, interpreted in the chosen p-adic embedding. The integral basis change and its effect on the Weyl gauge are formal-source input; the estimate after this change is the argument of Lemma 18.17. The full `weylGauge_cartanProduct_le` in the same file separately accounts for the positive and

negative columns of the gluing and for the conjugate embedding. In the source's notation the resulting exponent is

$$mloss = \sum_{\tau} \left(|a| \sum_{h<|N|} \kappa_{\tau}^{+}(h) + |m| \sum_{h<|m|} \kappa_{\tau}^{-}(h) \right).$$

The functions $\kappa^{\pm}$ count columns of the indicated heights, and the term $|a|$ is the number of divided coordinates. This estimate avoids charging every box of every diagram for a denominator occurring in just a few matrix columns. For instance the shape $(3,2)$ has three columns; a datum with only one bad column costs at most $|M|_p^{-3}$, rather than the total-degree bound $|M|_p^{-5}$. When several factors of this kind are multiplied in an integral Cartan product, their exponents add by multiplicativity and the ultrametric inequality. That saving is important in the subsequent determinant bound.

18.11 Limits of the reconstruction

The expanded argument now exhibits the finite matrix and character identities, the exact origin of the two sharp norm denominators, the algebraic part of the Rallis unfolding, the transition of signed critical-value products, the contraction for a bounded Hecke annihilator, and two comparison-lattice estimates. In particular the scalar in the exchange cannot simply be treated as an arbitrary norm of a theta integral, and the Gram factor cannot simply be treated as the norm of one finite indicator. In each case the normalization is what controls the twist-dependent arithmetic variation.

The analytic Siegel–Weil formula and its uniform estimates, genuine Weil model identifications, the existence and nonvanishing of the equal-rank test datum, the full exchange on partners with real determinant-vector surjectivity, propagation along the two-grid path, the sharp Euler pairing and its injectivity, the general tower boundedness and period conversion, and the geometric construction of the glued comparison lattice have been kept as formal-source inputs at their stated declarations. Nothing in the finite calculations substitutes a standard theorem with superficially similar hypotheses for these declarations. Their use in the rest of Theorem C continues to have the logical status specified in Section 17.5.

Chapter 19

The exchange construction, the two grids, and comparison lattices

There are three rather different issues hidden in the comparison of twisted and untwisted row volumes. The first is the adequacy of the finite matching relation: one must know that an identity tested on standard column functions implies an identity for general partners. The second is the construction of the two paths of theta lifts with precisely matching signatures, parameters, and ramified data. The third is the integrality of a comparison form which sees both the small definite group and the large auxiliary group. It is inappropriate to infer this last integrality from a rational seesaw identity. A lattice on which the two constituents really glue is needed.

We expand these parts of Chapter 17 and the companion Chapter 18. The results below give conventional proofs of the local matching factorization, the stable-space argument with determinant Gauss sums, the real hyperbolic enlargement of determinant vectors, the terminal branching interlacing, and the construction of the divided lattice and its denominator estimate. We then return to how these calculations enter the volume formula and the rank-one value estimates. Some additional identities of the genuine global theta model are explicitly retained as inputs. In particular the analytic existence and control of the Fourier and constant terms in that model must not be confused with the finite Fourier calculations that are proved here.

19.1 Local integrals of complementary frames

Let F be a finite product of non-Archimedean local fields, O its product of valuation rings, and ϖ a simultaneous uniformizer. Write $G_d = \mathrm{GL}_d(F)$ and $K_d = \mathrm{GL}_d(O)$. A primitive column matrix is an injective matrix over O admitting a left inverse over O; for row matrices use the transposed definition. These sets are open and compact: they are images of K_N under projection to the appropriate first columns or last rows. They are open since invertibility of a square matrix over O holds exactly when its reduction is invertible. A locally constant function supported on the primitive set is thus a Schwartz–Bruhat function. More generally our Schwartz–Bruhat spaces are the locally constant compactly supported functions on the full matrix spaces whose support consists of full-rank matrices. This last convention is important: support can approach a singular matrix neither from within nor from outside.

Lemma 19.1 (Two-sided choice of coset representatives)**.** *There is a set of representatives $\mathcal{R}_d$ for the right-K_d cosets which also represents the left-K_d cosets. In particular both $\mathcal{R}_d$ and $\{r^{-1} : r \in \mathcal{R}_d\}$ represent the right-K_d cosets.*

Proof. The assertion can be proved inside each double coset $K_d g K_d$. Its right and left coset decompositions both have finite length, since the intersection of two compact open subgroups has finite index in each. A left Haar measure on G_d is also right invariant: the density $|\det A|^{-d}\,dA$ on the invertible matrices has that property, by the determinant Jacobian of left or right matrix multiplication. With K_d of volume one, the lengths of the left and right decompositions therefore agree. Each of the left cosets intersects each of the right cosets in the double coset: for $K_d gk$ and $k'gK_d$ take $k'gk$. Match the two finite sets by a bijection, choosing an element in each corresponding intersection. The union of these choices over the double cosets is the required set. For a product of fields the same argument applies using product Haar measures. □

In the source this is the freedom of choosing representatives in `exists_common_coset_reps`, `New/Exchange/BlockPartners.lean`. For an injective $Y : F^d \to F^N$, the lattice $\Lambda_Y = \{u : Yu \in O^N\}$ has rank d. Over each valuation ring it is free, by elementary divisors, and $Y\Lambda_Y$ is a direct summand of O^N: it is saturated because F is a field in each factor. For a matrix r of a basis of Λ_Y, Yr is primitive. Conversely any primitive Yr' has $r'O^d = \Lambda_Y$, because a primitive matrix has an integral left inverse. Exactly one $r \in \mathcal{R}_d$ has this property. Exactly one $r \in \mathcal{R}_d$ has Yr^{-1} primitive as well.

For any character $\chi : F^\times \to \mathbb{C}^\times$ tame on $O^\times$, define the local section of a column function by

$$\ell_Y(\Psi) = \sum_{r \in \mathcal{R}_d} \chi(\det r)\Psi(Yr).$$

For primitive-supported functions this sum has at most one nonzero term, with the representative independent of Ψ. For a Schwartz–Bruhat function with support contained in the injective matrices it has finite support: the linear embedding $g \mapsto Yg$ has a continuous left inverse, and the inverse image of the support is compact in G_d. Its image in the discrete coset set is finite. There is also an algebraic definition even for arbitrary functions: extend the ordinary summation functional on finitely supported sequences to all sequences by extending a basis, and compose with the evaluation map in the displayed formula. This extension has no assertion of continuity. It is not used to define a convergent integral for a nonintegrable datum.

For row functions define $\ell_X(\Phi) = \sum_r \chi(\det r)\Phi(r^T X)$. Use the same notation for the finite sum or the algebraic extension. Fix complementary dimensions $N = a + b$ and the splitting of indices into first b and last a. For $k \in G_N$, let Y_k be the first b columns of k, and X_k the last a rows of its inverse. Thus $X_k Y_k = 0$. A pair of Schwartz–Bruhat functions (Φ, Ψ) is matched in the larger sense if

$$\ell_{X_k}(\Phi) = \chi(\det k)\ell_{Y_k}(\Psi) \qquad (k \in G_N). \tag{19.1}$$

The definition does not demand that the functions be supported on primitive matrices. It is precisely `matchingPairs`.

Proposition 19.2 (Factoring a matched row functional)**.** *Let V be a complex vector space and let A be a linear map from the row Schwartz–Bruhat space to V satisfying*

$$\chi(\det g)A(\Phi(g\,\cdot)) = A(\Phi) \qquad (g \in G_a). \tag{19.2}$$

There is a linear map B on the column Schwartz–Bruhat space such that for all pairs in (19.1)*, $B(\Psi) = A(\Phi)$.*

Proof. For each primitive X choose $k_X \in K_N$ with $X = X_{k_X}$ by completing the rows to a basis. Define a linear map of column functions into the space of all row functions by

$$T\Psi(X) = \begin{cases} \chi(\det k_X)\ell_{Y_{k_X}}(\Psi), & X \text{ primitive}, \\ 0, & \text{otherwise}. \end{cases}$$

Extend A algebraically to the space of all row functions. We will show that $B = A \circ T$ works and that the value on any matched pair is independent of both extensions.

For each $r \in \mathcal{R}_a$, let $f_r(X) = 1_{\mathrm{prim}}(X)\Phi(r^T X)$. This is a row Schwartz–Bruhat function; on the primitive set its support is compact, and multiplication by the primitive-set indicator is locally constant. Write

$$\Phi(X) = \sum_r f_r(r^{-T}X).$$

The equality follows because, for any full-rank X, exactly one r has $r^{-T}X$ primitive, and at a rank-deficient X both sides vanish. In fact only finitely many terms of this equality are nonzero as functions. The full-rank support of Φ is compact, and the sets on which $r^{-T}X$ is primitive are pairwise disjoint open sets covering it. A finite subcover proves the claim. By (19.2), applied to $g = r^{-T}$, one therefore has

$$A(\Phi) = \sum_r \chi(\det r)A(f_r).$$

For primitive X, the finite sum $\sum_r \chi(\det r)f_r(X)$ is $\ell_X(\Phi)$, so by matching it equals $T\Psi(X)$. Off the primitive set it is zero. Hence $T\Psi$ itself is Schwartz–Bruhat in the matched case, and its image under A equals $A(\Phi)$. This proves the proposition without any continuity of A or the extensions. □

This proves `matchingPairs.exists_linear_factor` in `BlockPartners.lean`. It is an algebraic factorization, not a norm estimate.

We check that the integral partners of the earlier chapter belong to this larger relation. Suppose both functions satisfy the inverse-character frame covariances and are supported on primitive frames, and their integral pairing is given by (18.2). Given an arbitrary $k \in G_N$, choose $g \in G_a$, $u \in G_b$ such that gX_k and $Y_k u$ are primitive. The column image is exactly the integral kernel of the row matrix: indeed this is true over F factor by factor, and both are saturated. Complete the column frame to a basis by choosing a right inverse of gX_k; this gives $k' \in K_N$ with $X_{k'} = gX_k$, $Y_{k'} = Y_k u$. Relative to k, the determinant changes by the factor $\det u/\det g$, as in Chapter 18. Thus

$$\chi(\det g)\Phi(gX_k) = \chi(\det k)\chi(\det u)\Psi(Y_k u).$$

Taking u the column representative and g the transpose of the row representative proves (19.1). This also proves the content of `IsPartner.mem_matchingPairs`, `PartnersMatch.lean`: local constancy of one integral partner implies local constancy of the other, since on K_N the relation uses the continuous frame projection and the locally constant determinant character, and the frame projections are open maps onto the primitive loci. Zero extension is locally constant on the clopen primitive loci and their complements.

Equivariant passage to the difference of periods. For a pair of real data v, v' and a scalar c, let the two target-valued linear maps defined by genuinely convergent theta integrals be A_v and $A'_{v'}$. The determinant covariance required by (19.2) must be derived from the Weil action and the change of variable on the source group. Suppose it holds. Proposition 19.2 gives a linear difference

$$\mathcal{D}(\Psi) = B_v(\Psi) - cA'_{v'}(\Psi)$$

on the space M of matched column data, with $\mathcal{D}(\Psi) = A_v(\Phi) - cA'_{v'}(\Psi)$ for matched pairs. A translate by $g \in G_N$ preserves the matched relation: the row function becomes $\chi(\det g)\Phi(\cdot\, g)$ and the column function becomes $\Psi(g^{-1}\cdot)$. To check this, substitute $g^{-1}k$ in (19.1); the determinant factor is exactly $\chi(\det g)^{-1}$. The same statement with the local Weil determinant

multipliers on the two target integrals gives equivariance of the difference. The cancellation of those multipliers is a consequence of complementary dimensions and equal combined rank twist, not an arbitrary identification of the two Weil representations. The automorphy and centre-coefficient properties of this difference are the properties collected in `IsEquivariantLift` by `exists_isEquivariantLift`, `ExchangeDifference.lean`. The factorization and matching parts have now been proved. The precise local naming and global automorphy laws of the Weil representation used to obtain the full `IsEquivariantLift` are retained as inputs.

19.2 Generating all partner data by stable spaces

The following local argument supplies a particularly useful converse. Let $s : O^N \to O^b$ and $x : O^b \to O^N$ satisfy $sx = I_b$. Define the standard column function

$$S_s(Y) = 1_{\mathrm{Mat}_{N,b}(O)}(Y)\,\mathrm{inv}_O\,\chi(\det(sY)),$$

where $\mathrm{inv}_O\,\chi(a)$ is $\chi(a)^{-1}$ for $a \in O^\times$ and zero otherwise. Thus its support consists of primitive columns. This is `stdColDatum` in the source. Assume P is a subspace of the column Schwartz–Bruhat space stable under $f(Y) \mapsto f(gY)$ for $g \in G_N$. We give two results corresponding to `partner_mem_of_stdColDatum` and `IntegralFrame.partner_mem_of_unimodular_vanishing` in `New/Exchange/StableSpaces.lean`.

Lemma 19.3 (Positive standard function generates all partners)**.** *Suppose $p \in P$ has the form $p(Y) = r(Y)S_s(Y)$ for a nonnegative integer-valued function r, with $r(x) > 0$. Then every locally constant column function Ψ belonging to integral partners of character χ is in P.*

Proof. Because Ψ is locally constant, is supported on the primitive set, and obeys right frame covariance $\Psi(Yq) = \chi(\det q)^{-1}\Psi(Y)$, compactness of $\mathrm{Mat}_{N,b}(O)$ supplies an integer $c > 0$ such that

$$\Psi(Y + \varpi^c Z) = \Psi(Y) \qquad (Y, Z \in \mathrm{Mat}_{N,b}(O)).$$

For products of local fields choose the maximum of the finitely many depths. Define an invertible matrix

$$d = xs + \varpi^{-c}(I_N - xs), \qquad d^{-1} = xs + \varpi^c(I_N - xs).$$

The formulas follow since xs and $I_N - xs$ are orthogonal idempotents. In particular $sd = s$, $dx = x$, and d^{-1} has integral entries. Consider

$$f(Y) = \int_{K_N} \Psi(k^{-1}x)p(dkY)\,dk,$$

with K_N of Haar measure one. This integral is a finite linear combination of translates of p: the functions of k in the integrand are locally constant, and $k \mapsto p(dk\cdot)$ has open stabilizer because p is a locally constant compactly supported function. The last fact follows by taking a common residue depth on its compact support and a compact integral box containing that support. Thus $f \in P$.

If $S_s(dkY)$ is nonzero, write $m = skY \in K_b$. Then $dkY \in O^{N\times b}$, whence kY is integral by integrality of d^{-1}. Moreover Y is primitive, as m is invertible. The complementary part satisfies $(I - xs)kY \in \varpi^c O^{N\times b}$, so $kY = xm + \varpi^c Z$. By local constancy and covariance,

$$\Psi(Y) = \Psi(k^{-1}xm) = \chi(\det m)^{-1}\Psi(k^{-1}x).$$

Consequently, on the primitive set,

$$f(Y) = \Psi(Y) \int_{K_N} r(dkY) 1_{\{S_s(dkY)\neq 0\}}\, dk.$$

The integral is independent of the primitive frame Y, since K_N acts transitively on such frames. At $Y = x$, it is positive: $dx = x$ and $r(x) > 0$; the integrand is locally constant near $k = 1$ on the set where S_s is nonzero, and hence is positive on an open neighborhood of positive Haar measure. On nonprimitive columns both f and Ψ are zero. Dividing by this positive constant proves $\Psi \in P$. □

We next show why it suffices to annihilate the unimodular centre coefficients in the exchange difference. Let

$$[K\ I\ L] \in \mathrm{GL}_N(O), \qquad \begin{bmatrix} K' \\ I' \\ L' \end{bmatrix} = [K\ I\ L]^{-1}$$

with both K and L of b columns. Let $\bar{O} = O/\varpi O$, write η for the residue character equal to χ^{-1} on $O^\times$ (extended by zero on nonunits), and assume η is primitive as a multiplicative character of $\bar{O}$ in the sense of Chapter 18. Let ψ be an additive character of F with conductor exactly O. The pairing $\psi(\varpi^{-1}uv)$ on $\bar{O}$ is primitive.

For $z \in \mathrm{Mat}_b(F)$ put

$$S(z)(Y) = S_{K'}(Y + LzK'Y).$$

The frame identities include $K'L = 0$, $K'K = I_b$, $L'L = I_b$, and $I'L = L'K = I'K = 0$. Thus $S(z)$ is supported on full-rank columns and $S(z+T) = S(z)$ for integral T. The latter follows since whenever $K'Y$ is integral the extra term $LTK'Y$ is integral, while $K'(Y+LzK'Y) = K'Y$. When $K'Y$ is not an integral invertible matrix both sides vanish.

Proposition 19.4 (Unimodular-coefficient synthesis)**.** *Suppose M contains all the $S(z)$ and P is as in Lemma 19.3. Assume the following implication: a vector $f \in M$ belongs to P whenever all linear functionals Λ on M satisfying, for some $T \in \mathrm{Mat}_b(F)$,*

$$\Lambda(S(z)) = \psi(\mathrm{tr}(Tz))\Lambda(S(0)), \qquad T \in \mathrm{GL}_b(O) \implies \Lambda(S(0)) = 0 \tag{19.3}$$

vanish on f. Then every locally constant column function of integral partners is in P.

Proof. For each $r \in \mathrm{Mat}_b(\bar{O})$ choose an integral lift $\widetilde{r}$, with $\widetilde{I}_b = I_b$. Let

$$q = \sum_r \eta(\det r) S(\varpi^{-1}\widetilde{r}) \in M.$$

We verify that every functional in (19.3) kills q. If $\Lambda(S(0)) = 0$ this is immediate. Otherwise the invariance under integral translations of z gives $\psi(\mathrm{tr}(TZ)) = 1$ for all integral Z. Taking matrices with a single nonzero entry and using the conductor, T has integral entries. Then

$$\Lambda(q) = \Lambda(S(0)) \sum_r \eta(\det r)\psi(\varpi^{-1}\,\mathrm{tr}(T\widetilde{r})).$$

The sum is a matrix determinant Gauss sum. Proposition 18.4 shows that it is zero unless the reduction of T is invertible. In that latter case $T \in \mathrm{GL}_b(O)$, contradicting $\Lambda(S(0)) \neq 0$. Therefore $q \in P$.

Since P is stable under G_N it also contains $p(Y) = q(-\varpi^{-1}Y)$. We compute this function. In the frame coordinates write $Y_K = K'Y$, $Y_I = I'Y$, $Y_L = L'Y$. Set $m = -\varpi^{-1}Y_K$. A term indexed by r can be nonzero only if

$$m \in K_b, \qquad Y_I \in \varpi O^{(N-2b)\times b}, \qquad Y_L \equiv \widetilde{r}m \pmod{\varpi}, \quad Y_L \in O^{b\times b}.$$

Its coefficient is $\eta(\det r)\chi(\det m)^{-1}$. If r is invertible then Y_L is invertible; the conditions determine $r = \overline{Y_L}\,\bar{m}^{-1}$ uniquely. By the definition of η the coefficient is exactly $\chi(\det Y_L)^{-1}$. Hence $p(Y)$ is $S_{L'}(Y)$ times a zero-or-one function. At the integral frame $x_0 = L - \varpi K$, the conditions hold with $m = I_b$ and $r = I_b$, and $L'x_0 = I_b$. In particular the nonnegative multiplier of $S_{L'}$ is positive at x_0. Apply Lemma 19.3 with $s = L'$ and $x = x_0$. □

The last proof explains the mechanism that is otherwise easy to miss in the formal argument: an additive Fourier test at an invertible index can be converted by a multiplicative determinant Gauss sum into a positive cut-off standard column function, which is cyclic for the stable space of partner data. The sign of the inverse character matters in the calculation of $p(Y)$. If it were changed to the original character, the two factors of $\chi(\det m)$ would not cancel.

For the genuine automorphic application take P to be the stable kernel of the difference of the two periods, and Λ to be a centre Fourier coefficient of its value on a datum. The character eigen-relation in (19.3) is the equivariance for translations by the indicated radical. The proof of `IsEquivariantLift.apply_stdColDatum_eq_zero`, `New/Exchange/StandardCoefficients.lean`, establishes the needed hypothesis about *all* such functionals from two geometric vanishing conditions: vanishing at the unimodular Gram matrix of the smaller member in an adapted splitting, and vanishing at invertible indices not represented by its archimedean signature. The translation of an arbitrary primitive rational Hermitian index to these conditions involves a second splitting, weak approximation, the Fourier completeness on the radical quotient, and a local orthonormalization of integral Hermitian matrices. Proposition 19.4 settles the residue synthesis in that argument. The geometric vanishing and its promotion to the full automorphic coefficient statement are distinct inputs, to be verified in the exact mixed model.

19.3 Construction of the real data on an enlarged Hermitian space

The exchange theorem begins at a hyperbolic enlargement and descends by removing planes. There are two real analytic problems. One must construct data with a prescribed slice, and the removal of a plane must be sufficiently surjective on determinant vectors. Replacing these requirements by existence of some nonzero Gaussian integral would not suffice: the resulting datum would generally depend on the twist.

Lemma 19.5 (Smooth local Witt extension)**.** *Let $J = J^*$ be nondegenerate, $x \in \mathrm{Mat}_{n,b}(\mathbb{C})$ with x^*Jx invertible. There is a smooth map u from a neighborhood of x to $U(J)$ such that $u(x) = 1$ and $u(y)x\mathbb{C}^b = y\mathbb{C}^b$ for all y in that neighborhood. It can be taken to be a globally smooth map into $U(J)$ with this property only near x.*

Proof. Put $A^\sharp = J^{-1}A^*J$ and

$$P_y = y(y^*Jy)^{-1}y^*J.$$

For y close to x, P_y is a $\sharp$-self-adjoint idempotent with image $y\mathbb{C}^b$ and kernel its J-orthogonal complement. Put $S = P_yP_x + (1 - P_y)(1 - P_x)$; then $S(x) = 1$ and $SP_x = P_yS$. Thus S is invertible near x and maps the two summands of P_x to those of P_y. By taking the $\sharp$-adjoint, $P_xS^\sharp = S^\sharp P_y$, so $Q = S^\sharp S$ commutes with P_x. For y sufficiently near x we can form

$$Q^{-1/2} = \sum_{n\geq 0} \binom{-1/2}{n} (Q - 1)^n.$$

It is $\sharp$-self-adjoint since all coefficients are real, commutes with P_x, and depends smoothly on y. Then $u(y) = SQ^{-1/2}$ has $u(y)^\sharp u(y) = 1$ and maps the desired subspace onto $y\mathbb{C}^b$. The invertible b by b change of frame is $C = (x^*Jx)^{-1}x^*Ju(y)^{-1}y$, for which $u(y)xC = y$.

To make the map global, shrink the neighborhood so that the matrix logarithm $L(y) = \log u(y)$ is given by its series at 1. From $u(y)^\sharp = u(y)^{-1}$ one obtains $L(y)^\sharp = -L(y)$ by the functional calculus near the identity. Take a real-valued smooth bump ρ equal to one in a smaller neighborhood, supported where L is defined. Extend ρL by zero and use $\exp(\rho L)$. Its $\sharp$-adjoint is its inverse, and it agrees with u near x. □

This is the smooth Witt theorem `exists_contDiff_witt` of `New/Exchange/RealPair.lean`, with a direct matrix proof instead of the inverse-function-theorem argument of that file. It does not require J to be positive definite. The global extension is useful because compactly supported mixed-model data will be specified by a global smooth formula that uses $u(y)$ only near x.

We recall the relevant elementary analytic argument for the real Fock model. Choose signed coordinates so that the holomorphic variable in a positive coordinate is z, and in a negative coordinate is $\bar{z}$. Taking the product of these choices is a real isometry, so the weighted moments are unchanged. In unscaled coordinates, write a Fock function as $v(z) = e^{-|z|^2/2}f(z)$ with f entire, and normalize Lebesgue measure by $d\mu(z) = \pi^{-d}d\,\mathrm{vol}(z)$ on $\mathbb{C}^d$. The coherent kernel is $K(z,w) = \exp(z\cdot\overline{w} - |z|^2/2 - |w|^2/2)$ and has $|K(z,w)| = e^{-|z-w|^2/2}$. For a real radius or norm, replacing the Euclidean norm by the sup norm changes only constants depending on dimension. Define $m_N(v) = \int(1+|z|)^N|v(z)|\,d\mu(z)$. The source's notion `IsRapidDatum` asks that all these moments be finite; it does not impose a particular algebraic polynomial expression on v.

Lemma 19.6 (Simultaneous polynomial approximation)**.** *A Fock vector with $m_N(v) < \infty$ for all N can be approximated in every one of the norms m_N by one sequence of Gaussian times holomorphic polynomials.*

Proof. The reproducing identity is $v(z) = \int K(z,w)v(w)\,d\mu(w)$. For Gaussian monomials it is proved by integration in polar coordinates; the full Fock reproducing identity follows by the orthogonal Gaussian monomial basis. Absolute convergence for a rapid vector also follows from the kernel bound. Replace the factor $\exp(z\cdot\bar{w})$ by its scalar Taylor polynomial of degree l, obtaining $v_l(z)$ as the integral with this truncated kernel. The coefficients are Gaussian moments of $v(w)$, so v_l is a Gaussian times a polynomial. For all l the norm of the kernel is bounded by

$$\exp(-(|z|^2+|w|^2)/2 + |z||w|) = e^{-(|z|-|w|)^2/2}.$$

For fixed w the integral in z of this bound times $(1+|z|)^N$ is at most $C_{d,N}(1+|w|)^{N+2d}$: use polar coordinates $r = |z|$, bound $1 + r \le (1+|w|)(1+|r-|w||)$, and integrate the remaining polynomial times the Gaussian over the real line. The bound is integrable against $|v(w)|d\mu(w)$. Since the Taylor polynomials converge pointwise to the exponential, dominated convergence for the product measure in (z,w) shows $m_N(v - v_l) \to 0$. The polynomial degree l is independent of N, so the assertion is simultaneous. □

This reconstructs `exists_tendsto_weightedNorm` of `Lib/Weil/Fock/MixedFockModel.lean`. The argument gives density in all moments, which is stronger than density in the Hilbert norm. This strengthening matters in the square law: matrix elements of the kernel may grow polynomially with adelic height, so continuity merely in the Hilbert norm cannot be invoked.

We spell out the regularity estimate used in constructing smooth slices. In an adapted hyperbolic splitting the real oscillator realization can be written as the tensor product of a

Schrödinger coordinate and a lower Fock space; change of the Schrödinger polarization is by a normalized Fourier transform. A normalized coherent vector at z of the lower Fock space is denoted ω_z; it is the unitary Heisenberg translate of the vacuum and satisfies

$$m_N(\omega_z) \le C_N(1+|z|)^N.$$

This follows directly from the absolute Gaussian profile, by the inequality $(1+|x+z|) \le (1+|x|)(1+|z|)$. The signed Fock action of a unitary matrix U on the coherent expansion sends z to the corresponding linear coordinate transform and introduces its vacuum multiplier of absolute value one. For a smooth compactly supported multiplier t and a smooth family $U(a)$, the mixed-model section $t(a)\omega(U(a))v$ is constructed by the coherent expansion, Fourier transform in a, and integration of the resulting coherent translates in the extra coordinates. The following estimate verifies the existence in the moment topology for smooth data. We use only the standard mixed-model coordinate and Fourier identities in this argument; their exact form in the source is `MixedModel.exists_section` of `Lib/Weil/MixedModel/RapidSections.lean`.

Lemma 19.7 (Moment estimate in the section construction)**.** *Let $d(a)$ be smooth of compact support in real dimension p, and let $B(a)$ be a smooth matrix on a neighborhood of that support. For the coherent kernel $\omega_{B(a)y}(w)$, define*

$$H_y(\xi,w) = \int_{\mathbb{R}^p} e^{-2\pi i a\cdot\xi} d(a)\omega_{B(a)y}(w)\,da.$$

For any M,N there are C,L depending on d,B,M,N but not y such that

$$\int (1+|w|)^N |H_y(\xi,w)|\,d\mu(w) \le C(1+|y|)^L(1+|\xi|)^{-M}. \tag{19.4}$$

Hence for a rapid vector $v(y)$ the integral of $v(y)H_y(\xi,w)$ against $d\mu(y)$ has integrable moments of arbitrarily large degree in both ξ and w.

Proof. The coherent kernel is exponential of the expression $w\cdot\overline{B(a)y} - |w|^2/2 - |B(a)y|^2/2$, possibly with a smooth unit-modulus quadratic phase in the coordinates. The first two derivatives of the phase in a are polynomial expressions in y,w with uniformly bounded coefficients on the compact support, and the same statement is true for every finite derivative order. By induction and the product rule each derivative of $d(a)\omega_{B(a)y}(w)$ of order at most $2l$ is bounded in absolute value by

$$C_l(1+|w|+|y|)^{4l}e^{-|w-B(a)y|^2/2}.$$

One may increase the exponent without changing the argument. Multiplying by $(1+|w|)^N$, integrating over w, and using $|B(a)y| \le C_B|y|$ gives $C_{l,N}(1+|y|)^{N+4l}$ uniformly in a. Apply $(1-\Delta_a)^l$ under the integral by integration by parts; the multiplier on the Fourier exponential is $(1+4\pi^2|\xi|^2)^l$. The Fourier integral is thus bounded in the weighted L^1 norm in w by $C_{l,N}(1+|y|)^{N+4l}(1+4\pi^2|\xi|^2)^{-l}$. Take $2l \ge M$. The last assertion follows by integrating against $|v(y)|$ and taking M larger than p plus the required degree of the ξ moment. □

In the mixed model the Fourier-transformed vector in this construction is further translated by a linear function of ξ in the enlarged Fock space. Lemma 19.7 still supplies all its moments, because Heisenberg translation loses at most $(1+C|\xi|)^N$. Therefore those integrals are Bochner integrals of rapid Fock data. Applying the slice map, Fourier inversion on the compactly supported smooth integrand and coherent reproduction recover the prescribed $t(a)\omega(U(a))v$. This proves the section existence for smooth choices used here, using the mixed-model transform identities; in the formal development the weaker C^8 regularity is

permitted by a further scaled Fourier bound. No assertion about that minimal regularity is needed for a smooth construction.

For a fixed lower datum v take $U(a) = 1$ and a nonnegative, nonzero smooth compactly supported t. The slice at zero in the opposite polarization is a fixed nonzero scalar times the integral of the forward slices, since the opposite polarization is the Fourier transform, evaluated at zero. The scalar is fixed by the two Haar measures and the mixed-model vacuum convention. Normalize t by that scalar and $\int t$. This gives one rapid enlarged datum whose opposite slice is exactly v. It is the construction `exists_rapid_with_realSlice` in `New/Exchange/RealPair.lean`. The scalar is independent of the varying finite character; this is the feature needed in the equal-rank starting identity.

19.4 A plane removal on determinant vectors

At an infinite place let W, V have sign involutions J_W, J_V (with diagonal entries ± 1 in orthonormal coordinates). The positive coordinates of $U \in \mathrm{Hom}(V, W)$ are those with equal source and target signs. If P is the projection onto these coordinates, it is $P = (1 + J_{\mathrm{Hom}})/2$. For a prescribed coefficient matrix F form the square matrix $M(U)$ by the positive-coordinate projection followed by multiplication by F in the convention of the source. Denote by M_+ and M_- its two principal sign blocks. A determinant vector of the family `detVectors` is a Gaussian times

$$\overline{\det(M_+)^m \det(M_-)^{m+2k}};$$

in the adelic real model multiply these factors over the infinite places with $m = m_w$. The exponent difference $2k$ will cancel the character arising from the hyperbolic plane.

Let $W' = W \oplus H$ and $V' = V \oplus H$ be enlarged by one hyperbolic plane. Its adapted orthonormal positive-negative coordinates can be chosen so that in the isotropic basis the sign matrix exchanges the two line coordinates, and the original block has sign matrix J_W or J_V. The adapted mixed-model slice at a line variable a sends a positive-coordinate polynomial to its evaluation at an old vector with the two new isotropic coordinates equal to $-a$. Indeed if J_W in the new basis has line block $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, then

$$P_+(u_1, u_2, v) = ((u_1 + u_2)/2, (u_1 + u_2)/2, P_+(v)).$$

At $(u_1, u_2) = (-a, -a)$ this is exactly the required projection. Orthogonality of the adapted frame also gives squared Euclidean norm $2|a|^2 + |v|^2$ in these coordinates. In the mixed-model polynomial slice, the Gaussian is evaluated at the same vector. This follows by inserting the coherent exponential in the Fourier transform of the hyperbolic line and differentiating in its polynomial variables. Explicitly the Fourier transform of the Gaussian with linear coefficient is again a Gaussian with the transformed linear coefficient; independent Gaussian derivatives generate every polynomial in the projected coordinates. The normalized transform formula for the chosen splitting is `PositivePart.exists_slice_eq` in `New/Exchange/Determinant Vectors.lean`; the projection and norm calculation above show why the two line coordinates in this formula really are equal, rather than opposite.

Proposition 19.8 (Determinant enlargement and its scalar)**.** *In the preceding adapted coordinates extend the coefficient matrix F by an identity on the new hyperbolic plane, with its normalization chosen by the line scale of the mixed model. After the line slice with local nonzero scalar a the two principal minors become respectively*

$$\det(M'_\pm) = a \det(M_\pm).$$

Suppose the rank identity $\mathrm{rk}(W') = 2k + 2q(W')$ *holds. When the cone integration on the enlarged pair is formed with weight* $k+m_w$ *at the infinite place* w*, the sliced determinant-vector integrand is a positive real scalar times the original determinant vector, with no* U*-dependent multiplier. In particular if the standard cone integrand is integrable and incorporates a nonzero positive line measure, it is possible to normalize the enlarged determinant vector so that the integral is the originally prescribed determinant vector.*

Proof. The chosen coefficient extension is block diagonal between the hyperbolic planes and the old coordinates. Under the slice the enlarged matrix in the two new isotropic coordinates is aI_2; their image in the positive-negative orthonormal coordinates is still aI_2. On the old block the matrix is M. Thus the new equal-sign principal blocks have one additional diagonal entry each, equal to a, and determinant multiplicativity proves the first assertion. The ratio of the polynomial factors is

$$\overline{a}^{m_w}\overline{a}^{m_w+2k} = \overline{a}^{2(k+m_w)}.$$

In the cone integrand the line character of weight $k + m_w$ contributes the compensating complex phase. The product of the line factors in the normalization of `coneIntegrand` is $|a|^{2(k+m_w)}$ times the positive idele-absolute-value factor of the line. To see the cancellation, write $a = re^{i\alpha}$, $r > 0$; multiplying the preceding polynomial ratio by $e^{2i(k+m_w)\alpha}$ gives $r^{2(k+m_w)}$. The maximal compact action on the extended polynomial in this step has no further phase, since the determinant exponents on the two signs differ by $2k$ and the rank identity says the vacuum exponent difference is the same. This is the sign and rank check for the character in the cone integrand.

In addition the Gaussian acquires just the positive line Gaussian, since its restriction has squared norm $2|a|^2 + |v|^2$. It follows that at the product of infinite places the integrand is of the form

$$D(q)v(U), \qquad D(q) = \exp(-\textstyle\sum_w |a_w(q)|^2)|d(q)|^z \prod_w |a_w(q)|^{2(k+m_w)}$$

in the source's line coordinates and positive Haar density. Here q denotes the line integration variables, $d(q)$ the corresponding idele in the line torus, and z the cone exponent. All of these factors are nonnegative and strictly positive on the nonzero line locus. If v is nonzero, integrability at one point with $v(U) \neq 0$ implies finiteness of $\int D$; nonzero positive line measure implies this integral is positive. Multiplying the extended coefficient by $(\int D)^{-1}$ produces the prescribed v. For $v = 0$ take the zero datum. Polynomial Gaussian data are rapid by direct Gaussian integration, so this normalization stays in the determinant-vector family. □

This calculation gives the matrix and positive-scalar part of `surjOn_cone_integral_detVectors`, `New/Exchange/DeterminantVectors.lean`. The hypotheses about the adapted mixed model, the line-scale vacuum multiplier, and convergence are part of the genuine global line identity `integrable_coneIntegrand` in the source; they are not consequences of an arbitrary polynomial extension. The proposition with those hypotheses gives surjectivity, and it provides an explicit inverse on every determinant-vector ray. Multiple plane removals compose these inverses. The zero slice in complementary polarizations likewise composes by Fubini on the rapid sections: separating the two line variables and integrating in either order is valid by their joint absolute moment bounds.

19.5 How the equal-rank normalization is obtained

We now state the intermediate identities used in `New/Exchange/EqualRankValue.lean`, to separate what is needed analytically from what is pure normalization. Let X be of even rank,

and $X^+ = X \oplus H$ an enlargement by a hyperbolic plane. For a real datum v and character of weight α write

$$R_{X^+}(\nu, v) = \frac{|\theta_{X^+,X^+}(\nu; v, \mathrm{standard})(1)|^2}{\Pi_\nu^{\mathrm{int}}(\mathrm{rk}\, X^+/2)}.$$

Here Π_ν^{int} denotes the positive product `lValueProdInt` of the earlier chapter, including its precise conductor normalization. For the top identity X^+ is read with its adapted splitting. The number R_{X^+} is called `equalRankThetaRatio`. The relevant source declarations are as follows.

1. `thetaLarge_eq_apply_one_of_even_rank` in `LowRankIdentity.lean`: for the equal-rank theta integral on standard-column data the value on general standard columns is obtained from the value at the identity by the real slice and Weil determinant covariance of the columns. This uses the rational open orbit in the enlarged hyperbolic splitting.

2. `stdExchange_of_periods` of the same file: when the top periods on the two source groups have the appropriate real slicing relation, it gives a standard-column identity for all target vectors and a factor whose magnitude is expressed using R_{X^+}.

3. `exists_stdExchange_top` in `RealPair.lean`: for X^+ of rank $2k + 2q + 2$ and the corresponding top pair $V' \oplus X$, V', there are real data realizing the preceding relation simultaneously for all tame characters of the prescribed weight. Its construction uses the smooth Witt lemma and the slice sections, with one compactly supported bump around the nondegenerate frame.

4. `exists_constant_term_thetaLarge` in `Lib/SiegelWeil/Tower/LargeSideConstantTerm.lean`: integrate one line radical of the larger-side theta integral with standard-column data. The remaining cone integral is a normalized lower-rank theta datum; the scalar is independent of the varying datum, and its squared magnitude is the appropriate last Hecke factor.

These identities of the genuine global theta integrals are still taken as geometric and analytic inputs. In items 2 and 3, the exact open-orbit statement, measure factors, and real action on the two data have to agree; Lemmas 19.5 and 19.7 give the local construction and rapidity in the latter item, not the global period identity.

For clarity consider the real geometry of item 3. Let the base subspace of the combined source have frame $x_0 = (I, 0)^T$ and ambient form $J = J_0 \oplus J_1$. For a frame x with x^*Jx invertible put

$$z(x) = x(x^*Jx)^{-1}x^*Jx_0.$$

The difference $x_0 - z(x)$ is orthogonal to $\mathrm{im}\, x$. At $x = x_0$, $z(x_0) = x_0$. If $z(x)$ is sufficiently near x_0, its Gram matrix is invertible, and therefore its b columns span exactly $\mathrm{im}\, x$. The smooth Witt map $u(z)$ sends $\mathrm{im}\, x_0$ to $\mathrm{im}\, z = \mathrm{im}\, x$, so $u(z)^{-1}x = x_0C$ for a unique invertible b by b matrix C. Its Gram matrix is $C^*J_0C = x^*Jx$. This is the factorization enabling the two periods to be evaluated in a common set of real coordinates near a nondegenerate frame. All these maps are smooth where defined. Choose a nonnegative smooth bump $t(z)$ supported in the indicated neighborhood, and a second bump on $u(z)^{-1}x$. Their product has compact support as a function of x because $x = u(z)x'$ with z, x' in compact sets and u continuous. Multiplying the product by a fixed polynomial determinant and the requisite compact-unitary translates of the tensor product of the two lower real data gives the smooth slice used in item 3. The precise theta integral equality for these real data uses the normalized period and open-orbit identities of items 1 and 2; the smoothness, support, and existence of one rapid Fock section follow from the earlier lemmas.

We next make the normalization argument explicit. For removal of one plane the usual standard-column constant-term formula has a zero-vector term and nonzero isotropic terms. On the smaller side a primitive standard-column datum kills the nonzero isotropic terms: their finite invariance under a line translation would force a primitive column frame to annihilate a nonzero coordinate. The formal proof uses `constTermAt_smul_k` and `name_eq_zero_of_invariant` in the mixed model to prove this assertion for the actual subspace. On the larger side the nonzero line integral is exactly the cone integrand, so it creates the lower datum in Proposition 19.8. After canceling the common nonzero vacuum factor the ratio of squared scalar magnitudes is the transition quotient of Lemma 18.10. Since removing a plane decreases the rank difference by two, its product over all removed planes telescopes to the requisite $\Pi_\nu^{\mathrm{int}}(e)$. This is the induction in `exists_stdExchange_removing_planes`. The surjectivity asserted there on determinant vectors follows from Proposition 19.8; down the other side the zero slices compose, as just explained. For the arbitrary real datum used in the starting exchange the normalized cone integral, not an inverse of the polynomial map, is used.

The proof of nonvanishing of a universal equal-rank constant in the source has a useful structural form. Put J for the definite group of Hermitian rank $2k$ and Z for the zero-dimensional group. Begin with the alternative that all equal-rank ratios $R_{X^+}(\nu, v)$ are zero. Apply the top period identity to the product with the zero-dimensional partner, and remove all the hyperbolic planes using the determinant surjectivity. This makes the zero-constant identity reach a nonzero determinant (in the case of Z, vacuum) datum at the definite edge, forcing the rank-zero theta value to vanish. On the zero-dimensional space there is a single rational vector, the zero vector, and the normalized vacuum at it is 1. The class integral is over the trivial group and has positive mass. Thus that value is nonzero. The application of this contradiction to the exact top identity is the first part of `exists_equalRankThetaRatio_value`. It uses the period identities above, in particular the lift and descent of the standard-column equality even if the scalar on that equality were zero; it does not assume the result to be shown.

It is not yet enough to have $R_{X^+}(\nu_0, v) \neq 0$ for one character. The datum v must work for all characters simultaneously. The second part of `exists_equalRankThetaRatio_value` exchanges twice between the two extreme ranks (J, Z) and (Z, J), with real data chosen independently of the characters. If B, C are the two universal exchange constants and q the real scalar of the rank-zero datum, the resulting scalar identities give

$$R_{X^+}(\nu, v)^2 = BC|q|^2. \tag{19.5}$$

Here the conductor and local Hecke factors cancel: in one direction the dimensions differing by $2k$ give $\Pi_\nu^{\mathrm{int}}(k)$ and in the other direction $\Pi_\nu^{\mathrm{int}}(-k)$, with product one by the signed-product convention. This is precisely the identity called `F` in the proof of that declaration, obtained by pairing one exchange with the adjoint of the other. Equation (19.5) therefore gives one number R_{X^+} independent of ν, since every ratio is a nonnegative real number. Nonvanishing at ν_0 implies its strict positivity for all ν. In the last exchange, to return from the flipped Hermitian form to its original sign, the Weil flip contributes a unit-modulus determinant character. Its square magnitude does not change the equality. A separate scalar from the change of Haar normalization is independent of the twist. Hence the resulting constant has

$$|c(\nu)|^2 = C_0 \Pi_\nu^{\mathrm{int}}(e), \qquad C_0 > 0,$$

with $e = \operatorname{rk} V - \operatorname{rk} W/2$. This is the normalization of `exchange_of_partners` and `exchange_theorem`. The nonzero constant C_0 cannot be obtained from the local matrix Gauss sum alone. In the reconstruction the double-period identity leading to (19.5), the geometric coefficient vanishing and the normalized line constant-term formula are still the principal global inputs to the exchange theorem. The algebraic factorization, stable-space step, smooth-real-data mechanism and scalar bookkeeping of their deductions have now been unpacked.

19.6 From a convergent exchange to an operator square

Suppose a complementary-rank exchange is available for one universal rapid vector on a direct sum of two real pair spaces. To use it in a chain of theta lifts one must pass from that rapid vector to the prescribed joint harmonics on the two constituent spaces. We give the analytic and linear portions of the argument in `New/Exchange/OneSquare.lean`. Write G_1, G_2 for the exterior groups of the square and A for the larger intermediate Hermitian space. The source estimate for a primitive finite datum is

$$|K_v((g_1, g_2), a)| \le C m_N(v) H_{G_1 \oplus G_2}(g_1, g_2)^B, \tag{19.6}$$

uniformly in the intermediate element a; the more general estimate contains a negative power of $H_A(a)$. This is the specialized kernel estimate used in `pairTheta_decay_of_onto`; its proof uses Poisson summation in the finite named datum, specifically its primitive support. We do not take a growth bound in a in place of the uniform bound in this assertion.

For rapidly decreasing y_i on the exterior quotients, choose C_i with $|y_i(g_i)| H_{G_i}(g_i)^B \le C_i$. The direct-sum height satisfies $H_{G_1 \oplus G_2}(g_1, g_2) \le H_{G_1}(g_1) H_{G_2}(g_2)$, in the source's normalization with heights at least one. Multiplying (19.6) by $|y_1(g_1) y_2(g_2)|$ gives a constant bound

$$|y_1(g_1) y_2(g_2) K_v((g_1, g_2), a)| \le C C_1 C_2 m_N(v). \tag{19.7}$$

Character factors are unitary and have no effect. With fixed quotient measures of finite total mass, the product integrand is absolutely integrable, and the resulting linear functional L of v obeys $|L(v)| \le C' m_N(v)$, where C' includes the three masses. This proves `exists_integrable_kernel_bound` once the stated genuine kernel estimate is available. The estimate is needed also for the smaller side, with the primitive block in its matching datum.

By Lemma 19.6, approximate the exchanged rapid vector by polynomial Gaussian vectors in all moment norms. A polynomial in the real coherent coordinates of a direct sum is a finite sum of products of polynomials on the two factors. Thus every approximant is a finite sum of elementary Gaussian tensors. For each elementary tensor the real convolution with the fixed type picks out a matrix coefficient scalar, independent of the finite twist. This statement is the compact-type equivariant convolution identity `IsEdgeFamily.exists_lift_eq_realConv`. It is an identity about the prescribed joint harmonic and the source and target type, not an assertion that two arbitrary vectors in an abstract irreducible representation are automatically proportional. After applying this identity and the adjoint relation to both exterior groups, an elementary tensor contributes some scalar times the prescribed composition of lift and descent. The scalar may be zero.

For the limit argument let I index all test forms, test functions, twists, and primitive finite data to which the universal real vector must apply. Write P_i for the value of the prescribed composition at test i and $D_i(v)$ for the integral of the approximating tensor datum. Each finite polynomial approximant v_n has $(D_i(v_n))_{i \in I} = \kappa_n (P_i)_{i \in I}$, with the same scalar for every i. By (19.7), $D_i(v_n) \to D_i(v)$ for each fixed i. If all P_i vanish then all the limits vanish and one may take $\kappa = 0$. Otherwise select i_0 with $P_{i_0} \neq 0$; then $\kappa_n = D_{i_0}(v_n)/P_{i_0}$ converges, and each limit equals κP_i. This proves the closed-span step without finiteness or a uniform norm on the index set I. It is essential here to approximate by one sequence independent of all the twists. Dividing by the nonzero complementary-rank exchange scalar $c(\nu)$ gives the square factor $\kappa / c(\nu)$. The determinant covariance of the two edge kernels comes from the local Weil action on named point data. Once this covariance and the two kernel estimates are verified, the convergent exchange implies the entire square law; this is the analytical argument of `one_square`. It does not prove that κ is nonzero, nor is nonvanishing needed for the crude-row transport. In the injective weighted-chain application this issue is settled separately by the first-arrow norm formula.

19.7 Terminal interlacing and the signatures of the two grids

We give a proof of the combinatorial grid identities of `New/Exchange/GridParameters.lean`, rather than using the terminal branching as an unspecified representation-theoretic requirement. Keep m, k even, $r = k - 1 + m$, $t = 5r$ odd, $h = 2r^8$ and $h_2 = h/2$ as in the admissible data. At an infinite place write $w_l = w_0 + lh$, where w_l includes h_2 (that is, $w_l = b_a + lh + h_2$ with $a = 1$ or 2 as appropriate to the place). Use $z = 0, 1$ for the two grids, i, j for their coordinates and o for the triangle offset. The rank at (z, i, j) is

$$n(z, i, j) = r + z + i(t + z) + j(t - z).$$

There are $r + z$ base coordinates, i added blocks of size $t + z$, and j added blocks of size $t - z$. The base signs at coordinate u, $0 \le u < r + z$, are

$$(-1)^{o+a+\min((a-o)_+, j)}, \qquad a = (u + 1 - k)_+,$$

and the signs on all added blocks are $+1$. Here $v_+ = \max(v, 0)$ is the truncated natural subtraction in `entrySign`.

Lemma 19.9 (Grid signatures and parity)**.** *At offset $o = 0$ the number of negative signs is $\lfloor (m - j + z)/2 \rfloor$ for $0 \le j \le m$. Moreover $n(0, i, j) \equiv 1 + i + j \pmod 2$ and $n(1, i, j) \equiv 0 \pmod 2$. The diagonal member of grid 1 has the signature of the diagonal member of grid 0 plus one line of sign $(-1)^j$ (when $i = j$).*

Proof. The base signs with $a \le j$ have sign $(-1)^{2a} = 1$. For $a > j$ they have sign $(-1)^{a+j}$; the positive integers a run exactly once from 1 to $r + z - k = m - 1 + z$. Thus the negative signs correspond to odd positive integers $b = a - j \le m - 1 + z - j$. There are $\lfloor (m + z - j)/2 \rfloor$ such integers. Added blocks are positive. Since r, t are odd, the displayed ranks have the stated parity. At $i = j$ the ranks in the two grids differ by one. Their negative counts differ by $\lfloor (m - j + 1)/2 \rfloor - \lfloor (m - j)/2 \rfloor$, which equals one exactly when j is odd since m is even. This proves the sign of the added line. A basis matching preserving the signs exists by separately matching the positive and negative coordinates; it gives the bijection used in `gridDiagEquiv_spec`. For $j = 0$ the extra base coordinate has positive sign, so the original coordinate inclusion may be used. □

The doubled Harish–Chandra parameter on the base block has entries

$$2\rho_u = \begin{cases} k - 2 - 2u, & 0 \le u < k - 1, \\ 2h_2 - 2w_{u+2-k}, & k - 1 \le u < r + z. \end{cases}$$

The positive added block numbered l, $1 \le l \le i$, is $2w_l + t - 1 - 2b$ for $0 \le b < t + z$, and the negative-centred added block numbered l, $1 \le l \le j$, is $t - 1 - 2(w_{l+o} + z) - 2b$ for $0 \le b < t - z$. This last description concerns the values of the parameter, not the signs of the summands, which are all positive. To all entries add 1 if $z = 1$; if $z = 0$ subtract 1 when $i + j$ is odd, and otherwise add zero. The rank parities in Lemma 19.9 show that every doubled entry is congruent to $n(z, i, j) - 1$ modulo 2. The spacing of the centres $w_l \mapsto w_{l+1}$ is h, and the width of an added block is at most $2t$. The base entries away from the central block lie halfway between the centres of the negative added blocks. From $h > 4t + 4$, $2w_1 > k + t + 2h$, and the parameter bounds (these follow from the large-log inequalities), they are distinct from the added entries and from one another. In particular when $o = 0$ every pair of oppositely signed coordinates is separated by at least $h - 2t - 2$: the negative signs only occur on base outer coordinates. Distinct entries within any same-sign arithmetic block are separated by 2, not by

h. This is why a gap condition between opposite signs must not be confused with a lower bound h on the separation of *all* entries. The source verifies the exact `HCParam.HasGap` convention at offset zero and along the first vertical strip by the same arithmetic-block comparison; for the joint harmonic conditions a further comparison of the centered parameters on adjacent edges is required.

Proposition 19.10 (Terminal branching interlacing)**.** *At offset zero and node (m, m) let V be the multiset of doubled parameters in grid 0 and let Z be the multiset of negatives of the doubled parameters in grid 1. Then for every real x,*

$$0 \le \#\{z' \in Z : z' > x\} - \#\{v \in V : v > x\} \le 1.$$

Proof. We enumerate entries from largest to smallest. On the positive-centred band of centre $2w_l$, $1 \le l \le m$, the V entries are

$$2w_l + t - 1,\ 2w_l + t - 3, \dots, 2w_l - t + 1,$$

and those of Z, obtained by negating the negative-centred block of grid 1, are

$$2w_l + t - 2,\ 2w_l + t - 4, \dots, 2w_l - t + 2.$$

They alternate with a V entry at each end. The extra outer base entries of Z on the positive side are $2w_l - h - 1$ for $1 \le l \le m + 1$. For $l > 1$ such an entry lies strictly between the positive bands of centres $2w_l$ and $2w_{l-1}$, since $w_l - w_{l-1} = h$ and $h > t + 1$. At $l = m + 1$ it lies above the largest band, and at $l = 1$ below the smallest band but above the central block, by the lower bound on w_1. Thus descending across the positive side, the difference starts at zero, rises to one on a base Z entry, drops to zero on the next V, rises to one on the intervening Z, and so on, ending at one just before the central block.

On the central block V has $k - 2, k - 4, \dots, 2 - k$ and Z has $k - 3, k - 5, \dots, 1 - k$. These alternate in the order V, Z, so starting at one the difference remains zero or one and ends at one.

The negative-centred band of V at $-2w_l$ has endpoints $-2w_l + t - 1$ and $-2w_l - t + 1$, whereas the reflected positive block of grid 1, giving the negative band of Z, has endpoints $-2w_l + t$ and $-2w_l - t$. They alternate with a Z at each end. The base negative extra entries of V are $h - 2w_l$, $1 \le l \le m$. In decreasing order each occurs immediately before its respective negative band and strictly after the preceding band or central block. Indeed it is halfway between consecutive negative band centres (with the same endpoint inequality as above). Thus each such base V lowers the difference from one to zero, and the following band keeps the difference in $\{0, 1\}$ and raises it to one at its end. This proves the count inequality for every cut between entries, and at each entry itself since no entries coincide. All the separation requirements used here are only the indicated simple inequalities on h, w_1, k, t. □

This gives `gridParam_interlace`. At the terminal nodes both groups are compact by Lemma 19.9. The Weyl dominance convention for the dual representation reverses and negates its highest weight, so this interlacing of doubled parameters is exactly the condition for restriction through the codimension-one embedding. To see the half-integral shift, a rank d ordered highest weight $\lambda_1 \ge \cdots \ge \lambda_d$ has doubled parameter $2\lambda_i + d + 1 - 2i$, and a rank $d+1$ highest weight μ of the other (dual) member has doubled parameter $2\mu_i + d + 2 - 2i$. The displayed count inequality says $2\mu_i + d + 2 - 2i > 2\lambda_i + d + 1 - 2i > 2\mu_{i+1} + d + 2 - 2(i+1)$. Since the weights are integral, these inequalities are equivalent to $\mu_i \ge \lambda_i \ge \mu_{i+1}$. The ordinary branching rule for unitary groups now gives a nonzero pairing. Its existence over the rational coefficient field follows from algebraic representation theory: the intertwining equations in the

highest-weight models are rational equations, so a nonzero complex solution implies a nonzero rational solution. Clearing its denominator gives an integral matrix for the pairing. The needed *quantitative* bound for such a matrix in the specific Weyl realization, including the powers of the determinant characters, is the additional estimate in `exists_branching_pairing` and `endBranching`. The preceding elementary interlacing is enough for existence, not for that norm bound.

19.8 The finite square move with no accumulated denominator

The standard local exchange kernel of Chapter 18 has an orbit denominator. We now prove the finite transformation of tidy data that removes it. This is the main finite part of `IsTidy.exists_square_move` in `New/Exchange/GridMove.lean`.

Work over the finite valuation-pair ring R and use p, v, u, r, T for the free modules attached to the relevant vertices and the extra Fourier block. In the notation of the cited declaration the old edges are $(p, v), (Q, U)$ with $Q = p \oplus T$, $U = u$, and the new edges are $(v, r), (U, R')$ with $R' = r \oplus T$. We identify the row and column bases by the bijections prescribed in that declaration. The complementary dimensions satisfy $\operatorname{rk} p + \operatorname{rk} r = \operatorname{rk} v + \operatorname{rk} u$. For a tuple S the old data are

$$A = S(p, v) : p \to v, \qquad B = S(Q, U) = [B_0 \ B_1] : p \oplus T \to u.$$

Set $Y = (A, B_0)^T : p \to v \oplus u$. In the new tuple choose matrices

$$D : v \to r, \qquad C = \begin{pmatrix} C_0 \\ C_1 \end{pmatrix} : u \to r \oplus T, \qquad X = [D \ C_0] : v \oplus u \to r.$$

Write $S'(S, C, D)$ for the tuple agreeing with S on the unaffected edges w, with these two new edge matrices and zero elsewhere. The old and new edges are disjoint from w. The phase in the move is $\psi(\operatorname{tr}(C_1 B_1))$, for the additive character ψ of the residual ring of the grid modulus. This is the trace of an endomorphism of T. In the ladder of just one grid, T is the empty block, so the phase is 1.

The two old edge characters are χ_1, χ_2 in the order above. The new characters on $(U, R'), (v, r)$ are χ_1, χ_2 in that order. Thus put $\chi = (\chi_1 \chi_2)^{-1}$. Let $\Delta_\chi(X, Y)$ be the complementary-frame kernel of (18.4); in this ordering its support is $\ker X = \operatorname{im} Y$ with both frames primitive. The kernel is zero elsewhere.

Let Λ be the old tidy functional. For any scalar function d of Y and any test function f on tuples define

$$G_d(f) = \Lambda\left(S \longmapsto \sum_{C,D} d(Y) \Delta_\chi(X, Y) \psi(\operatorname{tr}(C_1 B_1)) f(S'(S, C, D)) \right). \tag{19.8}$$

Initially set $d = 1$; we shall show that the same functional is obtained by a choice of d for which the nonzero coefficients in this expansion are roots of unity.

For a family g of vertex automorphisms as in the admissibility condition, write $g_{R'} = \operatorname{diag}(g_r, 1_T)$ under the prescribed embedding and $b = \operatorname{diag}(g_v, g_u)$ on $v \oplus u$. We need the other compatibility hypothesis of the source: for any $k \in \mathrm{GL}(p)$ we may replace the old vertex at p by k and at Q by $\operatorname{diag}(k, 1_T)$, leaving the rest (including the unaffected edges except for the permitted vertices) the given admissible family. In particular the action on the matrices in w is unchanged. Denote that old family by $\widetilde{g}$.

Lemma 19.11 (Covariance of the replacement functional)**.** *Under these compatibility hypotheses, write $\alpha_{\mathrm{old}}(\widetilde{g})$ and $\alpha_{\mathrm{new}}(g)$ for the products of edge characters of the old and new path respectively. Then*

$$G_{d(b^{-1}Yk)}(f \circ g) = \alpha_{\mathrm{new}}(g) G_d(f).$$

Proof. The old tuple transforms Y to $Y_1 = b^{-1}Yk$ and B_1 to $g_u^{-1}B_1$; the new tuple transforms X to $X_1 = g_r^{-1}Xb$ and C_1 to $C_1 g_u$. The trace phase is unchanged. The row-column kernel obeys

$$\Delta_\chi(X_1, Y_1) = \chi(\det g_r \det k / \det b)\Delta_\chi(X, Y)$$

by (18.5). The change of variables on the finite sum over the matrices (C, D) is bijective. This gives the stated transformed integrand in (19.8) with factor $\chi(\det g_r \det k / \det b)^{-1}$. Using old covariance of Λ multiplies by $\alpha_{\mathrm{old}}(\widetilde{g})$. Finally the ratio of the two edge character products is

$$\begin{aligned}\frac{\alpha_{\mathrm{new}}(g)}{\alpha_{\mathrm{old}}(\widetilde{g})} &= \frac{\chi_1(\det g_r / \det g_u)\chi_2(\det g_r / \det g_v)}{\chi_1(\det g_v / \det k)\chi_2(\det g_u / \det k)} \\ &= (\chi_1\chi_2)(\det g_r \det k / \det b),\end{aligned}$$

as required. The unaffected edge product cancels. This proves the formula with exactly the determinant powers of the two edges; using the same character on all four edges would in general give a different formula. □

Taking $d = 1$ proves covariance for the new functional. Taking $g = 1$ gives $G_{d(Yk)} = G_d$ for every $k \in \mathrm{GL}(p)$. Choose d to be $|\mathrm{GL}(p)|$ at one representative of each primitive column-frame orbit and zero at the other frames. Then its average over the right action is 1 at every primitive Y. Since the kernel is supported on these Y, linearity and averaging show $G_d = G_1$. The coefficient $d(Y)\Delta_\chi(X, Y)$ is either zero or a root of unity. The additive phase is also a root of unity, because the finite additive group has finite exponent.

If the old datum is supported on primitive matrices at its distinguished μ-edges, then Y is primitive as soon as either A or B has a left inverse (a left inverse for B restricts to one for B_0). The support condition propagates across the move. Indeed if $A = \mathrm{pr}_v Y$ has a left inverse L_A, the projection $1 - Y L_A \mathrm{pr}_v$ vanishes on Y and restricts to the identity on the complementary u component. Compose it with a right inverse of X as in Lemma 18.3; it follows that $C_0 = X|_u$ has a left inverse, hence so does C on the new μ-edge (U, R'). If $B_0 = \mathrm{pr}_u Y$ has a left inverse, use the analogous projection to conclude that $D = X|_v$ has a left inverse. This is exactly the propagation needed when the old edge carrying the twist and the new edge carrying the twist are opposite. Each new matrix has at most Q_* possibilities. If the old datum has order q then (19.8), with the chosen d, has at most Q_*^{q+2} nonzero summands in a presentation by root-of-unity coefficients. Thus it is short of order $q + 2$ as well as covariant.

Finally suppose the operator square law has scalar c and tests $L(X)$ and $M(Y)$ on the two routes. For each primitive Y it says $M(Y) = c \sum_X \Delta_\chi(X, Y) L(X)$ by Lemma 18.3. Here tests may depend on the matrices in w and on the extra Fourier variable $K : T \leftarrow u$. Multiply this equality by $\psi(\mathrm{tr}(KB_1))$ and sum over K, then apply Λ. The variable K can be read as the new block C_1, independently of $X = [D\ C_0]$. The right side is precisely cG_1 applied to the new route with $K = C_1$ and old unaffected matrices in w. This proves the equality of functionals in `IsTidy.exists_square_move`. The use of $G_d = G_1$, not a pointwise replacement of Δ_χ by $d\Delta_\chi$, is what proves both the identity and absence of an orbit denominator in the new datum.

19.9 Moving the descent along the path

Let the two paths of stage j descend from $(0, m)$ to $(0, j)$ and then lift to (j, j), one in each grid; the grid type z is suppressed in this description. The stage-j comparison is a bilinear

restriction period at the diagonal pair, evaluated against the tidy datum. To move a single descent past a lifts one repeats the preceding square move a times with $T = 0$. The datum has order increased by $2a$ and the scalar is a product of a square constants. This is seen formally by induction on a: for $a = 0$ the leg descent is literally the first arrow of the longer vertical path, with no change in the tuple or scalar. For the step $a \mapsto a + 1$, apply the induction with the final horizontal edge temporarily in the unaffected set w, then cross the square at the end and use (19.8) to remove the two old matrices. The test depends only on w and on the current forms, so the functional identity from that move is valid with precisely the same test. This reconstructs the combinatorial induction in `Data.descent_past_gridPathMap`, `New/Exchange/RowLadder.lean`, once the square law is available for each genuine edge.

To go from stage j to $j + 1$ do this in both grids for $a = j$. These $2j$ square moves use $2j$ square costs and $4j$ more point evaluations in the order exponent. The final diagonal period is transported by the real seesaw and change of the line sign, using one more move with the Fourier block T corresponding to that line. Its modulus scalar, from the Fourier-normalized flip of the line, is a pure conductor-cardinality power with one unit of cost. Together with the other diagonal scalar in `Data.hasRoute_succ`, the source counts at most $4j + 2$ square costs and $4j + 2$ increases in order. This last step uses the genuine line-kernel identity `exists_modKernel_flip_line` and the real seesaw with the bordering Hermitian embedding; it is not a consequence of the abstract path induction alone. For the finite part, compatibility of g with the bordering embedding is precisely the rule coupling the two diagonal vertex matrices in Lemma 19.9; the unaffected path has no edge incident to the released vertex. This is the situation of Lemma 19.11, while one of the two old edges is a μ-edge. Thus the square move is applicable as a finite transformation at this last step as well.

Once the stage identity with those two scalar factors is known, the total equality follows by multiplying, not adding, the nonzero factors $c_j(\theta)$ and by adding their square cost bounds. The total cost and order increase are $\sum_{j=0}^{m-1}(4j + 2) = 2m^2$. The initial datum of order $4m$ is thereby replaced by a terminal datum of order $4m + 2m^2$. The role of the predicate `IsTidy`, stronger than covariance just at the endpoint levels, is now apparent: without the independent admissible change at the intermediate vertex p in Lemma 19.11, the averaging over column frames would be invalid. A denominator of the order of a large finite general linear group would have been introduced at each of the squares. The proof above removes it before the next square can be applied.

We record also the finite-dimensional surjectivity step in reaching the initial rows. If a weighted chain A from one conjoined real type to the definite endpoint is injective, and its bilinear transpose D is the corresponding descent, then the restriction of D to the endpoint level space is onto the initial level space, provided the transpose identity holds on those level spaces. In fact the image is finite dimensional, and any vector annihilating it has A-image orthogonal to the entire endpoint space, hence zero. For the application the conjugate real type is identified by the sign-flip isometry, so the bilinear transpose is identified with the Hermitian adjoint. The level preservation of both chains follows from their finite point-datum covariance for principal level, and the equality of transpose pairings follows by Fubini for each lift-descent pair and cancellation of the finite weight w_j at that edge. This is `Data.iotaPlus_adjoint_weightedChain` and `Data.gridLevelForms_eq_map_iotaPlus` in `New/Exchange/RowsReached.lean`. Thus it is the injectivity of the long weighted chain that gives surjectivity of a fixed map, not a choice of forms separately for each row.

To pass from the second-grid starting level to the first-grid principal form space, one must still prove the quantitative restriction assertion `Data.eq_zero_of_orthogonal_gridLevel Forms`. Its proof constructs a real translate from a two-factor discrete-series coefficient with a nonzero integral, approximates local convolution by a Poincaré sum at the deep principal

level, and treats points near the boundary by a first-minimum estimate. The estimate at the finite level is uniform over the whole form space. The finite-dimensional argument just given shows what it suffices to prove, but does not justify replacing this assertion by a qualitative restriction-density theorem. We keep this quantitative real restriction input (in addition to the genuine square law and line-kernel identity). With it, bilinear duality gives the surjection in Proposition 17.9, and the proved finite transport of tidy data applies to all of its period tests simultaneously.

19.10 The divided lattice that glues the two comparison groups

We now turn from the first (forward) row volume estimate to the reverse estimate. This part uses the definite target of the single lift of good forms of Chapter 17. Its coefficient vector is a Cartan product of a vector on the large definite part and one on the small definite part. Their ramified finite support must be combined into an integral datum whose descent is their product. Splitting the real vector alone is not enough; it is the lattice constructed here that permits the two finite conditions to lie on the same point form.

At the chosen split places let O be the product of valuation rings (the two chosen places above the primes of g), π a generator of their product ideal, and $R = O/\pi^2 O$. We use N for the rank of the big part, s_0 for that of the small part, and r for the common row rank. Let $F : O^r \hookrightarrow O^N$ and $Z : O^r \hookrightarrow O^{s_0}$ be coordinate inclusions, so each has an integral coordinate left inverse F^T, Z^T. Let S be an a by N matrix over R, with $a \le N$; these are the rows indexed by the member data in the source. Put $T_S = \binom{S}{\pi I_N}$ over R. In the actual splitting, the inclusion inc of the common r coordinates in the full target is $(0, Z)^T$ after the prescribed basis reordering, and the condition on the big block uses the fixed inclusion F. These are the embeddings in `DividedSplitting` of `New/Rows/StackedDeterminants.lean` and `New/Rows/GluedFormTheta.lean`.

The first fundamental observation is that T_S has an integral lift with row lattice of full rank, even if S has very low rank after reduction. This allows the lattice to vary with the row-test matrix S, rather than attempting to use one splitting at every residue point.

Lemma 19.12 (A basis of the divided row module)**.** *Choose an integral lift $\widetilde{S}$ of S. There are integral square matrices D, E of size N such that $ED = DE = \pi I_N$, and the rows of D form a basis of the O-module generated by the rows of $\binom{\widetilde{S}}{\pi I_N}$. In particular, over R, for any column matrices X, X' of the corresponding size,*

$$T_S X = T_S X' \quad \Longleftrightarrow \quad \bar{D}X = \bar{D}X'.$$

Proof. Over each discrete valuation ring factor the row module contains πO^N and is contained in O^N, hence is free of rank N. Choose an ordered basis, assembled as rows of an invertible matrix D over the fraction field. Since every πe_i is in the row module, the coefficients expressing πe_i in that basis are integral; putting them into the rows of E gives $ED = \pi I$. As D is invertible over the fraction field this gives $DE = \pi I$ as well. Combine these choices over the finitely many factors. Since the two row generating sets are integral linear combinations of one another, the equivalence of the two matrix tests holds after any reduction, in particular modulo π^2. □

Let $P_Z = ZZ^T$ be the diagonal projector onto the selected small coordinates and $B = 1 - P_Z$ its complement. Form two square matrices of size $N + s_0$ over O:

$$C = \begin{pmatrix} E & FZ^T \\ 0 & P_Z + \pi B \end{pmatrix}, \qquad G = \begin{pmatrix} D & -DFZ^T \\ 0 & \pi P_Z + B \end{pmatrix}. \tag{19.9}$$

Then

$$CG = GC = \pi I_{N+s_0}. \tag{19.10}$$

Indeed the upper-left and lower-right blocks have product πI; the upper-right block of CG is $-\pi F Z^T + F Z^T(\pi P_Z + B) = 0$. The reverse equality follows by inversion in the fraction field. These are the matrices constructed by the Hermite basis argument in the source's `DividedSplitting.exists_isGlued`.

Proposition 19.13 (Indicator identity for divided gluing)**.** *Write $\mathcal{H} = \{Y \in \mathrm{GL}_r(R) : Y = I \bmod \pi\}$. For every full-size column matrix $X = (X_+, X_-)^T$ of r columns over R,*

$$\begin{aligned} &1_{\{\bar{G}X = \pi(0,Z)^T\}} \\ &= \sum_{Y \in \mathcal{H}} 1_{\{X_- Y = Z\}}\, 1_{\{T_S X_+ Y = T_S F\}}. \end{aligned} \tag{19.11}$$

Furthermore for an integral change of basis $K \in \mathrm{GL}_N(O)$, chosen independently of X, the big column block of $\bar{G}K$ satisfies $\pi(\bar{G}|_+ K)$ has zero columns outside a prescribed set of a positions. Finally $|\mathcal{H}| = |O/\pi O|^{r^2}$.

Proof. The lower block of the equation on the left of (19.11) is $(\pi P_Z + B)X_- = \pi Z$. Multiplying by B shows $BX_- = 0$, so $X_- = ZM$ for $M = Z^T X_-$. The equation then says $\pi M = \pi I$. As the annihilator of π in $O/\pi^2 O$ is precisely πR, we have $M = I + \pi A$. It is invertible, with inverse $I - \pi A$. This is the unique element $Y \in \mathcal{H}$ satisfying $X_- Y = Z$. The upper block of the left-hand equation is $\bar{D}(X_+ - F Z^T X_-) = 0$, or equivalently $\bar{D} X_+ Y = \bar{D} F$, because $Z^T X_- Y = I$ and Y is invertible. By Lemma 19.12, this is the upper test on the right of (19.11). Conversely any summand equal to one gives the unique Y and these two equations, and implies the original full equation by reversing the calculation. Thus either both sides are zero or exactly one term on the right is one.

For the column assertion, reduce the row matrix S modulo π, factor by factor. Its rank over a residue field is at most a. There is therefore a basis over the residue field such that all but a prescribed columns are zero: extend a basis of the kernel to a full basis and label the remaining at most a vectors by those positions. Lift the basis matrix to $K \in \mathrm{GL}_N(O)$; any lift of an invertible residue matrix is invertible over the valuation ring. The same K works for the row module generated by $\widetilde{S}$ and πI modulo π, because every row of D is an integral combination of those rows. Hence all columns of DK outside the prescribed positions vanish modulo π, giving the asserted vanishing of $\pi \bar{D}\bar{K}$. The big block of G consists of D and zero rows, so the same assertion holds for it. Finally $Y = I + \pi A$ with A unique modulo π, and multiplication is addition of A because $\pi^2 = 0$ in R. There are $|O/\pi|^{r^2}$ choices. □

The proof gives all the finite assertions of `DividedSplitting.IsGlued`, as well as `DividedSplitting.exists_congruenceGroup_inv_eq_one_sub` and the bad-column assertion in `DividedSplitting.exists_isGlued`. Notice two separate features of the use of modulus π^2: the factor πI_N ensures the existence of the row basis and bounds its denominator by one π, whereas the quotient of unit groups at the two depths makes the normalizing Y unique. Modulus π alone would not make the latter argument.

Realizing the matrix equation on a genuine finite unitary point. At a split place of E the unitary group of a standard definite Hermitian space identifies with $\mathrm{GL}_{N+s_0}(E_w)$, the matrix at the conjugate place being inverse conjugate-transpose of the given matrix. Define the finite adelic unitary point λ by taking C of (19.9) at the selected places, the inverse conjugate-transpose of C at their conjugates, and the identity away from these places. This defines the source's `DividedLattice`. Indeed $\pi\lambda^{-1} = G$ at the chosen places

by (19.10), and λ^{-1} is integral at their conjugates because it is the conjugate-transpose of C. The generating idele π has value one away from the selected places, so the entire matrix $\pi\lambda^{-1}$ is integral at every place, with reduction $\bar{G}$ at the chosen modulus. This proves `DividedLattice.exists_reducedInverse_eq`.

If a finite unitary point u is integral and equals I mod π^2 at the selected places, then $\lambda^{-1}u\lambda = I$ mod π there. In fact write $u = I + \pi^2 A$, and compute

$$C^{-1}uC = I + \pi GAC,$$

which is integral and congruent to I mod π. At their conjugates integrality follows from the inverse-transpose relation and the fact that the calculated matrix has determinant a unit, as its reduction is the identity. At all other places the conjugate is just u. This is the local-level statement `DividedLattice.reduceAt_conj_eq_one`. Together with Proposition 19.13, it specifies the actual adelic lattice rather than just a mod-π^2 congruence.

The row-fixing subgroup of the large definite factor used in gluing fixes T_S by right action modulo π^2; in particular it fixes its lower block πI_N. Thus each of its matrices is congruent to I mod π. This proves the containment `rowFixingLevel_le_levelOf` in the principal chosen level. All rational stabilizers of this subgroup are trivial by the free-level result of the earlier adelic chapter, since they are contained in rational stabilizers at the chosen level of the appropriate rank. The source uses this freeness when taking a class sum, so no inverse stabilizer appears in the comparison operator.

19.11 The loss of the Cartan product

Consider a rational prime ℓ not dividing the twist modulus, and the valuation ring $\mathcal{V}$ of a chosen coefficient embedding into $\mathbb{C}_\ell$. Let M be a nonzero rational integer belonging to the chosen product ideal (in the application take $M = g$). Multiplication by M clears denominators at each chosen split place. Because π times the whole inverse matrix of λ is integral, all columns of $M\lambda^{-1}$ are integral at every finite place. The particular bad-column assertion of Proposition 19.13 shows more: after the basis change K in the big column block, all but a of the columns of $\lambda^{-1}|_+K$ were already integral, without the factor M. Indeed $\pi G|_+K = 0$ mod π^2 outside the specified positions is equivalent to $G|_+K = 0$ mod π, and $\lambda^{-1} = \pi^{-1}G$ at the chosen places. At conjugate places the original inverse was integral, and the basis change can be taken trivial away from the chosen places. This proves the missing divisor-clearing assertion used in Chapter 18.

For comparison, let ρ_τ be the polynomial Schur action of shape with columns of multiplicities $\kappa_\tau(h)$, at an infinite weight label τ, now construed over $\mathcal{V}$ through its corresponding coefficient embedding. If A is a rectangular matrix with only columns in a set J failing integrality and MA_j is integral for $j \in J$, then the bound of Lemma 18.17 is valid for this rectangular polynomial Schur action as well. Indeed $A = (AD)D^{-1}$ with the diagonal matrix D on the column space; the integral polynomial action of AD is lattice-preserving and the action of D^{-1} is diagonal in tableau weights. Each symbol can occur at most once in a column, so its exponent is at most $\sum_h \kappa_\tau(h)$. Thus the loss is

$$|M|_\ell^{-|J|\sum_h \kappa_\tau(h)}.$$

To include the finite point h_+ of the big factor, factor its inverse column matrix as $\lambda^{-1}i_+h_+^{-1} = (\lambda^{-1}i_+K)(K^{-1}h_+^{-1})$, with i_+ the column inclusion. The compact matrix K^{-1} and its inverse are both integral, so its Schur action is an isometry for the lattice sup norm. The norm of $\rho_\tau(K^{-1}h_+^{-1})y$ is consequently the gauge of the coefficient vector y at h_+, not the possibly

much larger raw norm of y. This is the effect of the integral change of basis on the Weyl gauge in `norm_schur_inv_cols_le_weylGauge`, `New/Rows/CartanIntegrality.lean`. The proof of that declaration includes the multiplicative determinant twist; it cancels in this block because the prescribed Cartan product uses determinant exponent zero on the two factors and the product's exponent is the one dictated by its weight. We retain here the integral polynomial Schur model of the earlier adelic coefficient chapter.

Proposition 19.14 (Cartan gluing estimate)**.** *Let V_+, V_- be the coefficient modules of the big and small factors with multiplicities $\kappa_\tau^+(h), \kappa_\tau^-(h)$ as in the divided splitting. For their Cartan product $C(y,z)$ (the small factor with conjugate weights), at the point* $\mathrm{diag}(h_+,h_-)\lambda$ *the Weyl gauge is bounded by*

$$|M|_\ell^{-L}\,\|y\|_{h_+,\ell}\,\|\bar z\|_{h_-,\ell\circ\mathrm{conj}}, \qquad L=\sum_\tau\left(a\sum_{h<N}\kappa_\tau^+(h)+s_0\sum_{h<s_0}\kappa_\tau^-(h)\right). \tag{19.12}$$

Proof. The big part has at most a bad columns, so the preceding estimate over all weight labels gives the first exponent. For the small part it is enough to use all its s_0 columns with denominator at most M. Since its action is by the conjugate matrix, the coefficient lattice to be used is that for the embedding composed with complex conjugation, applied to $\bar z$. One cannot silently commute the chosen embedding with conjugation. Apply the same diagonal-weight estimate to those s_0 columns, after the conjugate embedding and the inverse-transpose relation for the unitary factor. The source's Cartan product in the integral polynomial model is an integral bilinear map and compatible with block Schur actions: on the weight monomials this is the product of the two families of minors. Hence the gauge of the product is at most the product of these two bounds by the ultrametric inequality. The resulting exponent is L. □

This proves the finite divisor part of `weylGauge_cartanProduct_le`, with the prescribed divided splitting and its sign convention. The estimate is much less than bounding the full polynomial degree of the product by charging each box for the pole. For the large factor a bad column can occur strictly at most once in every Young column, which gives the coefficient a in place of the full rank N. This economy is necessary because the eventual comparison form must have small Archimedean norm after its denominator has been cleared.

19.12 Descent of a glued point form

We now trace how the indicator identity produces the two kernels in `DividedSplitting.theta_productForm`, `New/Rows/GluedFormTheta.lean`. Let O_θ be the auxiliary pair; the chosen split level is M^2Q where Q is the twist conductor and the modulus M corresponds to π. Denote by b, b' the characters of the two definite factors satisfying $bb' = O_\theta.b$, and by $K_+(g,h;y)$ the kernel on the big block with datum $T_SX = T_SF$, including its determinant twist at Q. Denote by $K_-(g,h';z)$ the kernel on the small block with datum $X_- = Z$. Both use the prescribed harmonic polynomial in the conjugate coefficient vector. For a point form with large and small coefficient vectors y, z at $\mathrm{diag}(h,h')\lambda$ the normalized theta descent has finite datum

$$1_{\{\bar G X=\pi(0,Z)^T\}}\,\chi_Q(\det X_e)$$

at the ramified places; e selects the fixed rows for the twist. At M the finite matrix of the point has precisely the displayed support, since the integral argument of the point kernel is

$\lambda^{-1}X = \pi^{-1}GX$. At Q the finite point λ is the identity, so the character factor is unchanged. The multiplier of the Weil action at M is

$$\beta_\lambda = \chi_{O_\theta.a}(\det_M \lambda)|\det_M \lambda|^{-r/2},$$

where $\det_M \lambda$ is the idele restricted to the selected places, and $\chi_{O_\theta.a}$ denotes the space character of the pair with the row rank r. This is `DividedLattice.weilScalar`. The exponent $-r/2$ is the Jacobian exponent for the Weil change of variables on that block. It is important that λ, $\det_M \lambda$ and the value of this multiplier are independent of the varying tame twist on Q.

By Proposition 19.13, the indicator is the sum over $Y \in \mathcal{H}$ of the product of the two block indicators, evaluated after the right column action by Y. Choose a lift $\widetilde{Y}$ at the selected places of the row group; it is congruent to 1 mod M and is trivial at Q. Consequently the determinant character of the twist has value one on it. The genuine finite naming identity for theta kernels changes the row-group point g to $g\widetilde{Y}$. The joint-harmonic polynomial of the Cartan product factors into the two polynomial kernels on the two orthogonal definite blocks. In the tensor-product realization of theta, the rational summation variables of the two blocks are independent, so the theta sum for the product datum is the product of the theta sums. Absolute convergence, for the rapid data at issue here, justifies multiplying them. These give the formula

$$\theta(\mathrm{productForm}_\lambda(h,h';y,z))(g) = \overline{\beta_\lambda \sum_{Y\in\mathcal{H}} K_+(g\widetilde{Y},h;y)K_-(g\widetilde{Y},h';z)}. \tag{19.13}$$

This explains the two kernels and the single Weil scalar in `DividedSplitting.theta_pointForm` and `DividedSplitting.theta_productForm`. The indicator identity and character compatibility have been proved above. In (19.13) we use the exact point-form evaluation and local Weil change-of-polarization law of the underlying normalized theta construction, and the compatibility of the chosen polynomial basis with block restriction. These are genuine model identities; a different normalization of the finite Weil action could change β_λ and would not be sufficient. The formal proof verifies them by the identities for `pairWeil_name_on`, `modDatum_factor`, `modKernelI_block` and `mixedCartan_mulVec_cartanProduct`; the needed statements are precisely these ingredients, not a new existence theorem for the glued lattice.

If $y(h)$ is the big-part coefficient vector depending on its class, let $C_i(g)$ be its coset kernel sum, summing $K_+(g,h;y(h))$ over the row-fixing level; the index i records S and the determinant-row choice. If c is a scalar, taking the class sum of the point forms with coefficient $\overline{cy(h)}$ and applying (19.13) gives

$$\overline{\beta_\lambda c}\sum_{Y\in\mathcal{H}} \overline{C_i(g\widetilde{Y})K_-(g\widetilde{Y},h';z)}.$$

Every sum here is finite, by definiteness of that group and finiteness of its class set. This is `DividedSplitting.theta_stdClassSum_productForm`. If it is integrated against a row form of the principal chosen level, then the right action by $\widetilde{Y}$ can be changed in the integral without changing that form, because $\widetilde{Y}$ belongs to this level. Thus the finite congruence sum contributes exactly $|\mathcal{H}|$ in the final inner product. This factor is a multiplicity, not a denominator, and the independence from the twist will enter the following comparison.

19.13 Rank-one reflections and the small-root descent argument

The comparison forms must not merely be integral. A good form on the large definite group can share a root of the Hecke polynomial of a lift without being in that lift. The denominator

theorem applies to its complement only after the test vector is orthogonal to these false lifts. Two different facts are used to arrange this orthogonality: at a ramified place a reflection-sum annihilator promotes one vanishing descent to vanishing for all local data, while off level a small-root comparison propagates vanishing up the hyperbolic tower. We expand the algebraic parts of both.

Lemma 19.15 (Annihilator or standard frame)**.** *Let O be a discrete valuation ring with fraction field F, residue field k, and standard coordinate inclusion $J : O^r \hookrightarrow O^n$. Suppose $x : F^r \to F^n$ is injective. There is an integral matrix A annihilating x, with the following property. If all the entries of AJ reduce to zero in k, then there is $b \in \mathrm{GL}_r(F)$ such that xb is integral and congruent to J modulo the maximal ideal.*

Proof. By the lattice argument at the beginning of this chapter there is $c \in \mathrm{GL}_r(F)$ such that $P = xc$ is primitive. Choose a left inverse L of P over O, and set $A = PL - I_n$. Then $Ax = 0$. If $AJ \equiv 0 \bmod \varpi$, we have $\bar{P}\bar{L}\bar{J} = \bar{J}$. As $\bar{J}$ has rank r, so does the square $\bar{L}\bar{J}$, hence LJ is invertible over O. Take $b = cLJ$; the asserted congruence holds. If the assumption on AJ fails, a row of A gives an integral row l annihilating x for which the reduction of lJ is nonzero. □

This proves `AuxPair.exists_annihilator_or_congr_inclMatrix`, `New/SiegelWeil/AuxPairVanishing.lean`. It selects an annihilator of the entire local frame, not merely of a residue matrix. The distinction matters for the action on an arbitrarily small coset around x.

For a nonzero row $l : k^r \to k$, define the subgroup

$$H_l = \{\, I_r + ul : u \in k^r, \quad 1 + lu \neq 0 \,\} \subset \mathrm{GL}_r(k).$$

It is indeed a subgroup, since $(I + ul)(I + vl) = I + (u + v + u(lv))l$ and the inverse is $I - (1 + lu)^{-1}ul$. Its determinant is $1 + lu$, by the determinant lemma for a rank-one perturbation. Let χ be the relevant determinant character. In any representation ρ of this finite group,

$$\frac{1}{|H_l|} \sum_{h \in H_l} \chi(\det h)^{-1} \rho(h)$$

is the projection onto the simultaneous $\chi(\det h)$ eigenspace of H_l: translation in the sum gives the required eigenrelation, and on that eigenspace every summand acts as the identity. The numerator is the reflection sum `reflectionSum` of `Lib/LocalGL/Green/FrobeniusOrbits.lean`.

Suppose now l is the integral row provided by Lemma 19.15, with $lJ \not\equiv 0$. For every integral column u of length r and $h = I_n + Jul$ with $1 + lJu \in O^\times$, we have $hx = x$. The matrix h and its inverse are integral. Consequently h fixes every integral-box coset of the form $x + \varpi^k O^{n \times r}$, even for k large. Its reduction has J-columns within the image of J and induces $I_r + \bar{u}\overline{lJ}$ on them; every member of the latter group is obtained in this way. The other entries of the matrix are unconstrained in the level precisely because the indicated columns determine the block in the source's parabolic. If $D_x(y)$ is a period linear in a form y, and equivariance of the local theta action on this fixed coset gives $D_x(\rho(h)y) = \chi(\det h)D_x(y)$, summing the normalized reflection operator in this identity yields $D_x(y) = D_x(P_{H_l}y)$. Thus $D_x(y) = 0$ whenever the specified reflection sum annihilates y. At a twisting place where the action on the frame is by a scalar determinant character alone, the same argument can use a single h for which $\chi(1 + lJu) \neq 1$: the map $u \mapsto lJu$ is onto the residue field, so nontriviality of that character ensures such an h. The period is then equal to a nontrivial scalar times itself and must vanish. For the row convention instead of the column convention use the transpose construction. This is the elementary reflection and nontrivial-character argument at M and Q in `AuxPair.descentsVanish_of_reflectionSum`.

If no such annihilator occurs, the frames are congruent to the standard frame after a change of source basis, by Lemma 19.15. The exact proof that the associated finite coset period reduces to the specified auxiliary descent is a further step in the same declaration; it uses the named-datum translation law, the source-equivariance law and decomposition of arbitrary Schwartz–Bruhat data into such local coset functions. The genuine normalizations in that transport, as well as the approximation in the real Fock datum, are needed to deduce `DescentsVanish` from these local calculations. The rank-one reflection calculation does not assume the conclusion of that theorem.

The small-root propagation through one or more planes. Let V be the definite Hermitian space of rank n, let W have rank r, and let $W_s = W \oplus H^s$ with rank $r + 2s$, for $s \le p$. A descent of a form on V against a rapid Weil datum is a function on W_s. Suppose all the lower descents on W_{s-1} vanish. Integrating the descendent function on the radical along an adapted last line expresses its constant term as the lower descent, by the mixed-model identity for the zero term. Thus in forming the descendent function we may subtract that zero term. This is the identity `DescentsVanish.integral_sourceTerm_eq_zero`, which uses Iwasawa coordinates with respect to the adapted plane and the assumption that the finite datum is the vacuum off the allowed places. The induction in `descentsVanish_add_planes` requires a boundedness estimate for the difference between the original theta summand and its zero term. To see how the estimate is used, denote by $F(g,h)$ the theta kernel on the larger target and by $F_0(g,h)$ its radical average. If

$$\sup_{g,h} |F(g,h) - F_0(g,h)| \le C$$

after the chosen short-vector and finite-coordinate reductions, then the function $D(g) = \int F(g,h)y(h)\,dh$ is bounded by $C\|y\|_{L^1}$, since the lower term integrates to zero. The same estimate with an integrable majorant in h in place of the constant C suffices. The boundedness derived in the source uses the short rational isotropic vector of W_s, the Witt theorem to choose its splitting, then translation of the named finite data and cancellation of the zero term. This is the application of `exists_norm_name_le` with exponent zero after subtracting the average. It relies on a genuine adelic estimate and remains a quantitative input in this reconstruction.

Once boundedness is known, the Hecke root calculation is completely explicit. Let w be a split spherical place of norm q, and suppose the original Hecke polynomial annihilates the form at an argument X of absolute value $q^{r/2+1-n}$, the value `AuxPair.testValue` used by the source. The Hecke transfer identity at the conjugate place (the adjoint in the integral pairing exchanges the two places) gives a product of trivial factors and a smaller Hecke polynomial on the descendent function on W_s at argument X' of modulus

$$|X'| = q^{(n-(r+2s))/2}|X| = q^{1-n/2-s}.$$

Here character values have modulus one after the appropriate unitary normalization. All the trivial factors have the form $1 - q^{j+(r+2s)/2}\zeta X$ with $|\zeta| = 1$ and $0 \le j < n - r - 2s$. Their largest possible absolute norm of the term subtracted from one is

$$q^{n-r-2s-1+(r+2s)/2}|X| = q^{-s} < 1$$

for $s \ge 1$. Hence none of these factors vanishes. For the remaining polynomial,

$$q^{r+2s+1}|X'| = q^{r+s+2-n/2} \le 1$$

whenever $n > 2r + 2p + 3$, since then $n \ge 2r + 2p + 4$ and $s \le p$. The finite Hecke-operator contraction proof of Proposition 18.15 applies to this polynomial on bounded spherical

functions on W_s. Thus the descent is zero. Iteration in s proves the propagation from level zero to level p, under the just-stated boundedness and genuine Hecke transfer inputs. It is important to use $s \geq 1$ at the trivial-factor cancellation; at $s = 0$ the initial vanishing is instead the special auxiliary-descent condition, not a strict inequality $q^0 < 1$.

The exact transfer identity is `hecke_transfer_of_equivariant`, `Lib/Weil/Unramified/HeckeTransfer.lean`, and its adjoint is `heckePolyOp_adjoint`. Combined with the reflection mechanism and the analytic boundedness estimate, the above proves the tower conclusion of `descentsVanish_hypSum` in the actual application to conjugate matrix coefficients of algebraic forms. The further period conversion in `stdClassSum_period_eq_zero`, `New/Rows/MemberOrthogonality.lean`, uses that a class sum on the smaller block paired with a translate of the algebraic matrix coefficient is an auxiliary theta period on the larger descent W_p. To derive it, expand a kernel vector in matrix coordinates, form the orthogonal block product of the corresponding two Weil data, and integrate against the conjugate coefficient of the given form. The twist level and the gluing conjugation property proved after Proposition 19.13 make that integrand fixed by the required compact subgroup. The precise global period identity identifying this integral with the class sum (the statement `period_eq_zero` of `Lib/Shimura/ComplexTheta/PeriodVanishing.lean`) is a further model input. It must preserve both levels and conjugate weights. No eigenform uniqueness principle at one Hecke root can substitute for it.

19.14 The comparison map and the precise scalar it sees

We return to the notation $u_\vartheta : U_j \to V_j(\vartheta)$ for a single good lift. The small definite group is denoted V^- and the large one is V^+. Its coefficient space has a fixed positive Hermitian Gram matrix, the one of the integral Weyl model. Let I be the finite set of member indices, and $\mathcal{C}^-$ the class set of V^- at the squared chosen level M^2. Let T be the Hilbert direct sum of the small coefficient spaces indexed by $I \times \mathcal{C}^-$. This is `Data.Tables` of `New/Rows/ComparisonForms.lean`; T is independent of the twist conductor. Write $\Xi(i, c)$ for a vector of this table at (i, c).

For each i take the glued point λ_i of Proposition 19.13 and its multiplier β_i. Let $y_{\vartheta,i}(h)$ be the large-factor line multiplier vector, and c_* the line scale of the source (one fixed nonzero scalar). The comparison map $g_\vartheta : T \to V_j(\vartheta)$ is the following sum of actual point forms:

$$g_\vartheta(\Xi) = g^L \sum_{i,c} \sum_{h \in \mathcal{C}_i^+} \mathrm{productForm}_{\lambda_i}(h, c; \overline{c_* y_{\vartheta,i}(h)}, \Xi(i, c)). \tag{19.14}$$

The set $\mathcal{C}_i^+$ is the class set of the row-fixing level of the big group associated with S_i. All its stabilizers are trivial, as proved above. The scalar g^L, with L as in (19.12), is the fixed denominator clearing of the Cartan product. This is exactly `Data.comparisonForm`. For all appropriate coefficient places, the norm bound on the large factor of the form is the integral-line-scale condition `lineMultiplier_integral_and_small` of the source. Specifically it asserts that $\overline{c_* y_{\vartheta,i}(h)}$ has gauge at most one at every coefficient embedding at a prime prime to $Q(\vartheta)$; in the untwisted case the squares of the Archimedean conjugate coefficient norms are bounded by $\exp(rB_0)$, with B_0 the small budget of Chapter 17.

By Proposition 19.14, every summand of (19.14), including the factor g^L, is integral if the conjugate table coefficient $\overline{\Xi(i, c)}$ has gauge at most one at the conjugate embedding for its class c. The point-form operation itself preserves such integral vectors: it puts their translated coefficients on one adelic class and vanishes on the others, with only an integral compact action and a unit-modulus finite level character on its support. All the point forms are well-defined because of the prescribed level conditions in the auxiliary pair. Therefore

the finite sum in (19.14) is integral. This proves `comparisonForm_mem_integralForms` from the line-scale assertion and the genuinely constructed divided lattice, not from an abstract integral representation of the comparison map.

Write $C_{\vartheta,i}(g)$ for the coset kernel sum corresponding to the large line vectors in this comparison. The essential geometric identity in `Data.cosetKernelSum_eq_fNatural`, `New/Rows/Separation.lean`, is

$$F_i(g) = \eta_\vartheta C_{\vartheta,i}(g), \qquad |\eta_\vartheta|^2 = C/\Gamma_j(\vartheta), \tag{19.15}$$

for a fixed family of functions F_i (the `fNatural`) independent of ϑ. The number C is a fixed real constant, not initially known nonzero in the proof. The identity is an application of complementary-rank exchange to the large carrier and to the flat line kernels. Its actual construction also uses the determinant polynomials of `New/Rows/Carrier.lean`; the three line-pair products in Γ_j must appear exactly, since they are the special values that repay the idle directions. We retain this exact period equality as an input.

For a table subfamily indexed by $c \in \mathcal{C}^-$, let $T_i(\Xi)(g)$ denote the small-block sum of conjugate kernels (`Data.tableTheta`), with datum equal to the indicated small coordinate inclusion at M^2 and the prescribed polynomial of the conjugate coefficient. It too is independent of ϑ. Integrate the descent identity (19.13) first against a good row $F \in U_j$, then use (19.15). The class sum over h gives $C_{\vartheta,i}$, the sum over c gives $T_i(\Xi)$, and the congruence translates each give the same integral by principal-level invariance of F. Since the inner product is conjugate-linear in the first variable, this yields

$$\bar{\eta}_\vartheta \langle F, u_\vartheta^* g_\vartheta \Xi \rangle = |\mathcal{H}| g^L \bar{c}_* \sum_i \overline{\beta_i} \int \overline{F(g')F_i(g')}, T_i(\Xi)(g'), dg'. \tag{19.16}$$

This is the identity `inner_adjoint_comparisonForm`. Every factor on the right is independent of ϑ. The genuine theta point-form identities isolated in the preceding subsection assure the indicated normalization; the finite grouping and phase cancellation have now been shown.

One additional separation theorem for the fixed functions is needed:

$$\Big[\ int\overline{FF_i}\, T_i(\Xi) = 0 \text{ for all } i, \Xi\ \Big] \quad \Longrightarrow \quad F = 0. \tag{19.17}$$

This is `Data.eq_zero_of_integral_fNatural_tableTheta`. A crucial part of its proof is the statement `exists_fNatural_approx` in `New/Rows/NaturalFormNear.lean`: for $n_* =$ `nStar`, some function K_0 with nonzero $K_0(1)$ can be uniformly approximated by linear combinations of translates $F_i(gy)$ at points where the first adelic minimum of the moved lattice exceeds its specified cut-off. The error is at most $e \max(1, \|y\|)^{6n_*}$, with $4e(5b)^{7n_*} + e < |K_0(1)|$. The function is formed from the nonzero carrier theta lift and holomorphic minors; reproduction and a separate isolation estimate convert orthogonality to a contradiction. The archimedean localization and quantitative approximation, including the first-minimum threshold, are necessary to use this statement and remain inputs. This separates the problem of detection by actual theta products from the linear algebra of the table map.

Proposition 19.16 (Scalar and surjectivity of the comparison)**.** *Assume the genuine identities* (19.15)–(19.16)*, separation* (19.17)*, and* $\dim U_j > 0$*. Put* $A = u_1^* g_1$*. Then* $A : T \to U_j$ *is surjective and*

$$u_\vartheta^* g_\vartheta = \gamma_\vartheta A, qquad \gamma_\vartheta = \overline{\eta_1/\eta_\vartheta}, \qquad |\gamma_\vartheta/\gamma_1|^2 = \Gamma_j(\vartheta)/\Gamma_j(1).$$

Moreover, suppose for every false lift z *at an allowed split place the small-block periods of the table* $(g_\vartheta)^* z$ *all vanish, i.e.* $T_i((g_\vartheta)^* z) = 0$ *for every* i*. Then* $g_\vartheta A^* y$ *is orthogonal to these false lifts for every* $y \in U_j$*.*

Proof. If one $\eta_\vartheta = 0$, (19.15) makes every $F_i = 0$. Then (19.17) makes every $F \in U_j$ zero, a contradiction. Thus every $\eta_\vartheta \neq 0$ and the constant in (19.15) is positive. Comparing (19.16) for ϑ and 1 gives $\overline{\eta_\vartheta} u_\vartheta^* g_\vartheta = \overline{\eta_1} A$ by nondegeneracy of the inner product on U_j. The formula for the absolute value is immediate from (19.15).

If F is orthogonal to $\operatorname{ran} A$, the right side of (19.16) is zero on every table. Choose a table supported at one member index i and allow its small-class coefficients to vary freely. All other $T_{i'}$ then vanish. Since $|\mathcal{H}| g^L \bar{c}_* \bar{\beta}_i \neq 0$, the integral in (19.17) is zero for every such table. By separation $F = 0$; finite dimensionality of U_j gives surjectivity.

For a false lift z, let $v = g_\vartheta^* z$. By hypothesis $T_i(v) = 0$ for all member indices. Equation (19.16) at the trivial twist shows $Av = 0$, hence $\langle z, g_\vartheta A^* y \rangle = \langle A g_\vartheta^* z, y \rangle = 0$ for all y, as desired. □

The last hypothesis in Proposition 19.16 is exactly the consequence of the tower-vanishing and period-conversion assertions in the preceding subsection. Indeed a false lift belongs to the good idempotent, the kernel of the polynomial at the small test root, and the orthogonal complement of $u_\vartheta U_j$. The first property gives the reflection-sum annihilators, the third the vanishing auxiliary descent. The class-sum period conversion then kills tableTheta$((g_\vartheta)^* z)$. This is `tableTheta_adjoint_eq_zero` in the source. The proposition shows why testedness is asserted on the special range $g_\vartheta A^* U_j$, not on the entire form $g_\vartheta(\Xi)$ for arbitrary tables. The proof of `exists_comparison_forms` has now been reduced to explicit finite construction, norm, and the genuine fixed-carrier and orthogonality statements.

19.15 The Archimedean budget of the comparison operator

The remaining size requirement for this construction, at the trivial twist, is $h_- \|\sigma(g_1(\Xi))\|^2 \leq e^{2rB_0} \sum_{i,c} \|\sigma \Xi(i,c)\|^2$, where h_- is the hook factor of the small coefficient space. We detail why the large denominator exponent L is still affordable here. The untwisted line-scale hypothesis says that each big coefficient vector has squared conjugate norm at most e^{rB_0}. Cartan multiplication has norm at most the product of coefficient norms in the integral Fock-Weyl realization, with the same weights as the gauge estimate; this follows either from tensoring the Fischer forms and taking the Cartan component or from the polynomial Gram calculation in the adelic coefficient chapter. The point-form norm at any class is at most M_* times that coefficient norm, with $M_* = \text{chosenGroupOrder}$, because the point-form construction sums over a stabilizer of order dividing M_*. The finite character values have modulus one also after every field automorphism, and the Schur matrices are mapped with the conjugate coefficient coordinates rather than with a new choice of bases.

Let H_V be the number of target classes at the untwisted level, H_I the number of member indices, H_- the number of small definite classes and H_+ an upper bound on the number of classes of each big row-fixing level. Summing first over h, then over (i,c) in (19.14), the triangle inequality and Cauchy–Schwarz give the pointwise bound

$$g^L M_* e^{rB_0/2} H_+ \sqrt{H_I H_-} \left(\sum_{i,c} \|\sigma \Xi(i,c)\|^2 \right)^{1/2}.$$

The squared class-sum norm of the resulting form is bounded by its squared maximum times H_V (class weights are at most one). Thus

$$h_- \|g_1(\Xi)\|_\sigma^2 \leq h_- H_V H_I H_- H_+^2 M_*^2 g^{2L} e^{rB_0} \sum_{i,c} \|\sigma \Xi(i,c)\|^2. \tag{19.18}$$

These bounds hold for the fixed conjugate coefficient-space Hermitian form under the field automorphism, not for the image under the automorphism of a positive real norm inequality. The former is what is required in the compressed determinant.

We check the remaining logarithm along the admissible family of Chapter 26. The class bounds used in `normSqAt_comparisonForm_le` in `ComparisonForms.lean` are $\log H_V, \log H_-, \log H_+ \le b$ eventually; the analogous bound holds for the rank $6m+1$ flat class set occurring in H_I. Indeed each is bounded by the class-count estimate at the principal level M^2 in the relevant rank (all ranks are bounded by a power of x), and the row-fixing and comparison levels contain that principal level. The logarithm of the finite residue ring at M^2 is $2f \log g$, with f the number of infinite places. By `Data.MemberIndex` and `Data.LineChoice` the other choices in H_I consist of a matrix with $3m$ rows and $N = \operatorname{rk} V^+$ columns over this ring, and for every infinite place a list of $2m$ row choices among $6m+1$ repeated $h_2 - \lfloor (t+1)/2 \rfloor$ times. Their logarithmic cardinalities are bounded by

$$6mNf \log g + 2mfh_2 \log(6m+1) = x^{O(1)}$$

uniformly along that family. Since the small coefficient multiplicities are each h, as in `ComparisonForms.lean`, the total size of its Young diagrams and the logarithm of their hook number h_- are bounded by fixed powers of x. The group-order denominator M_* also has logarithm bounded by a fixed power of x, by the free-level bound of the adelic chapter.

For the nonpolynomial part of L, the source's `card_detIndex` expresses $\sum_{h<N} \kappa_\tau^+(h)$ as the determinant-index count minus the reduced half gap; the identity `card_detIndex_add_bigCharExp` bounds this count by $w_j(\tau) \le 2b$ eventually. The negative-column sum is at most $s_0 h$ per weight label. Since $a = 3m$ in the actual splitting, this gives

$$L \le 6fmb + fs_0^2 h.$$

Here $a = 3m$ refers to the number of divided row coordinates, not the prime-sum cutoff (which is still denoted by a elsewhere). The second term is polynomial in x. Also $\log g \le 22 \log b$. Write $Y = \log b = x^{\varepsilon}$, and for this estimate denote the large prime-sum cutoff by a_{pr}. Then

$$rB_0 = \frac{r}{k} \frac{fa_{\mathrm{pr}} Y^3}{m} \quad \ge \quad \frac{fmbY^3}{4},$$

using $r \ge k$, $a_{\mathrm{pr}} \ge xb$ and $m^2 \le 4x$ for large x. Therefore $2L \log g = O(fmbY) = o(rB_0)$. All the other logarithms in (19.18) amount to $O(b) + x^{O(1)} = o(rB_0)$ by Lemma 26.1. Consequently their sum is below rB_0 for all sufficiently large x. We recover precisely

$$h_- \|g_1(\Xi)\|_\sigma^2 \le e^{2rB_0} \sum_{i,c} \|\sigma \Xi(i,c)\|^2. \tag{19.19}$$

This proves the budget calculation in the norm bound of the comparison forms. It is an asymptotic verification on the cofinal admissible family, which is enough for the limiting argument. It uses the fixed untwisted line-scale norm estimate and the indicated class-count, tableau and integral point-form bounds. In particular it does not infer such a norm estimate uniformly at all twist conductors from the bound for the trivial conductor.

19.16 Two further finite calculations for the rank-one values

In the height estimate for $\Gamma_j(1)/\Gamma_j(\vartheta)$ there are two types of rank-one input. One is the arithmetic size of a normalized integral kernel coefficient; the other is the identification of the pairing of a test form with an actual Hecke value. The first is a delicate attainability result

(the depth indicator in Chapter 17). The second involves the exact normalization of class sums, which we detail here.

Fix a pair $p = (\nu, z)$ with target rank $2z$ and the remaining data of `LValueData`. Let Q denote a fixed squarefree twist modulus and consider its tame weight-zero characters ϑ. Write I_Q for the finite indexing set consisting of the chosen torus class set and the local groups of units at Q. The class sum defining $\Xi_\vartheta = \mathrm{trivialLift}(H, \vartheta)$ can be expanded by decomposing each principal-level torus class into these finitely many translates. If g_i, $i \in I_Q$, denote these translates, the finite character occurs only in the factor $\vartheta(\det g_i)$ of each kernel coefficient. Thus, after fixed normalization, write

$$\Xi_\vartheta = \sum_{i \in I_Q} \vartheta(g_i) K_i \tag{19.20}$$

with K_i the prescribed point kernel at modulus Q independent of ϑ. A choice of inverse of the character just interchanges i, j in the formulas below. This is the class expansion `trivialLift_apply`; in the source's proof the nontrivial character factors are those of the Weil determinant multiplier on the split torus. In particular the point-kernel data must be kept at the fixed modulus Q, including for the trivial character viewed at that modulus. One may not replace K_i by the point kernels of a varying truncated level.

Put $d_p = \mathrm{denomBound}$ as in Chapter 17 and $M_* = \mathrm{chosenGroupOrder}$. The normalized value is $\mathcal{A}_p(\vartheta) = M_* d_p^2 \|\Xi_\vartheta\|^2 / N(Q)^{2z-1}$. On inserting (19.20),

$$\mathcal{A}_p(\vartheta) = \sum_{i,j \in I_Q} c_{ij}\, \vartheta(g_j/g_i), \quad c_{ij} = \frac{M_*}{N(Q)^{2z-1}} \langle d_p K_i, d_p K_j \rangle. \tag{19.21}$$

This uses $\overline{\vartheta(g_i)} = \vartheta(g_i)^{-1}$, valid since the characters have finite order. It follows without assuming that an arbitrary automorphism of $\mathbb{C}$ commutes with complex conjugation.

The reason the coefficients in (19.21) are integral at primes prime to Q is as follows. For every such embedding the two point kernels multiplied by d_p have gauge at most one, by `weylGauge_pointKernel_le_one` of `New/Exchange/TriangleLevelExponent.lean`. This applies both to the embedding and its composition with conjugation. The Gram pairing of two such coefficient vectors is integral in the integral Gram model (one factor is evaluated at the conjugate embedding). The class-sum pairing divides each contribution by the order of its finite rational stabilizer. These orders divide M_*, and therefore M_* clears every denominator in the pairing. Finally the conductor power $N(Q)^{2z-1}$ is a unit at the given embedding. Hence every c_{ij} belongs to the corresponding valuation ring. This proves the finite-measure calculation in `isIntegralMeasure_intAlgPart`. It uses the point-kernel integrality and the fixed-modulus expansion in the genuine Weil model, and not an interpolation theorem for special values. Once the exact Rallis norm formula for the normalized characters is known, it is an integral measure for the removed Hecke L-values as well.

For completeness we explicate the three-class-sum test formula. In the following calculation the flat blocks have all coefficient degrees zero. Suppose a factor of even rank m_0 has flat theta lift Ξ_0, normalization d_0, and conductor Q_0 prime to the considered coefficient prime. Let $\mathcal{C}$ be the prescribed finite torus class set and let $\mu_c = 1/|\operatorname{Aut}(c)|$ be its weights. Write $\mathcal{N}_c = \sum_c \mu_c$. A vector-valued test family P_h, indexed by the classes h on the flat target, is combined with its flat lift by the operator

$$\mathcal{X}(P) = \frac{M_* d_0^2}{N(Q_0)^{m_0-1}} \sum_h \mu_h\, \Xi_0(h) P_h.$$

The placement of a complex conjugate in the class sum is determined by which entry of a Gram inner product is linear; the displayed placement agrees with using the flat Gram

functional $\langle 1, \Xi_0(h)\rangle$ in a linear functional of P. This is the operation `Data.Pair.xiSum`. It preserves integrality of P_h at the coefficient prime: $d_0\Xi_0(h)$ is integral by summing the integral normalized point kernels (and accounting for any source stabilizer by its fixed normalization); one remaining factor of d_0 is integral, $M_*\mu_h$ is integral, and the conductor factor is a unit. In the flat case the Gram coefficient factor is just a scalar, so this calculation is componentwise in the valuation ring.

Let L be a complex linear functional of the test vectors and suppose $L(P_h) = C_1 K_{g_0}(h)$ with K_{g_0} the corresponding flat kernel coefficient complementary to Ξ_0. The reproducing class-sum identity in the source is

$$\mathcal{N}_c\kappa_0^{m_0}\sum_h \mu_h\,\Xi_0(h)K_{g_0}(h) = \|\Xi_0\|^2, \tag{19.22}$$

after the included twisted central normalization of K_{g_0}. Here κ_0 is the vacuum scale and the torus level is the *chosen* level, not the level with additionally varying Q_0. To see the mechanism of (19.22), expand Ξ_0 as the sum of translates of a single point kernel on torus classes. The normalized flat lift is fixed by the twisted central action. Pairing each translated kernel with Ξ_0 therefore gives the same number, and their weighted sum is the squared norm of Ξ_0. Canceling the total class mass and the vacuum power gives the displayed identity with the correct supplementary norm factor. The exact twisted center action and normalized point-form adjunction are `inner_trivialLift` and `twistedCentralOp_trivialLift` used in `Pair.smul_apply_xiSum` of `New/Rows/HeckeValueHeight.lean`. They are genuine theta identities; the class-sum argument proves the scalar extraction once they hold. It follows directly that

$$\mathcal{N}_c\kappa_0^{m_0}L(\mathcal{X}(P)) = C_1\mathcal{A}_{p_0}(\vartheta_0).$$

For a rank-one pair of rank $2z$ in the application choose three auxiliary flat blocks of ranks $2z-20, 4, 4$ and a remaining block of size twelve. The sum of these ranks and twelve is $2z$. The finite residue datum fixes the appropriate columns at the chosen places and has determinant character the product of the four block characters, as required by the direct-sum Weil law. Choose the degree distribution $\alpha_{\tau,i}$ of the twelve-jet, so that $\sum_i \alpha_{\tau,i} = d_\tau$. When a power of a linear expression of degree d_τ is expanded in these coordinates, the coefficient of the monomial of multi-index α_τ is

$$\binom{d_\tau}{\alpha_\tau} = \frac{d_\tau!}{\prod_i \alpha_{\tau,i}!}.$$

Equivalently this is the ratio between the Fock Gram pairing of the degree-d_τ vector and that of the single monomial under the chosen polynomial normalization. Real tensor multiplication over the places multiplies these factors. On the flat blocks the monomial is 1 because their weights are zero. Therefore the initial point form of this monomial has, by the block product identity, kernel pairing equal to this multinomial product times the product of the three flat kernels and the twelve-coordinate kernel. Apply the class-sum operation $\mathcal{X}$ once for each flat block, with the other factors held fixed. Each application is linear and changes one flat kernel into its normalized value at the cost of the factor on the left of (19.22). Since all three use the chosen torus level and the sum of their ranks is $2z-12$, the final equality has exactly

$$\mathcal{N}_c^3\kappa_0^{2z-12}\langle u, K_g\rangle = \left(\prod_\tau \frac{d_\tau!}{\prod_i \alpha_{\tau,i}!}\right)\mathcal{A}_{p_1}\mathcal{A}_{p_2}\mathcal{A}_{p_3}\,K_\alpha^{(12)}(g). \tag{19.23}$$

The three normalized values are evaluated at the chosen auxiliary twists matched with the twelve-coordinate jet character, two of them trivial or flat auxiliary twists as in the source. This is the factorization at the end of `exists_test_form_inner_kernelForm`. In particular

the exponent $2z - 20 + 8 = 2z - 12$ of the vacuum factor represents only the three flat blocks; charging the twelve-coordinate kernel again would give an incorrect power. The construction gives an integral test vector if the normalized monomial vector of the twelve-block has gauge at most one and the indicated block-column point data are admissible. The source constructs such a vector by the integral harmonic-polynomial identity `exists_int_vector_harmonicPoly_eq`. The orthogonality of the vector to the false rank-one lifts is the small-rank application of the preceding reflection-sum and larger-rank period vanishing. These statements about the actual representation and period support are still needed to obtain testedness; the class-sum and multinomial computation alone does not imply it.

We recall why the depth acquisition is a separate problem. The chosen scale of `Data.depth_indicator_jet`, say t_0, must be nonzero, small under every Archimedean conjugate, and good at all finite embeddings in the sense that at each one some degree distribution and some integral point produce an attainable jet coefficient. If one has that input and (19.23), the finite character selection and multiplicative selection in Lemma 17.20 produce the required K inequality of the source. The first selection uses Fourier inversion on the finite torus class group, giving one twist at a chosen point whose Fourier coefficient is at least its original value. The second uses a product of the flat values over the full character group. Specifically for functions a, b on that character group, take a root V with $V^{|G|} = \prod_{\eta \in G} a(\eta) b(\eta)$. For a prescribed product character χ, multiply $a(\chi\eta^{-1}) b(\eta)$ over $\eta \in G$; this is exactly $V^{|G|}$. Hence for every coefficient embedding one factor has norm at least $|V|$, while for every Archimedean conjugate a uniform upper bound on all the a, b directly bounds $|V|$. A priori the selected factors may vary with the coefficient embedding, but the height witness V is fixed. This is the reason that the witness in the height proof does not depend on the chosen test form u.

The remaining novel ingredient is precisely the attainability assertion of `depth_indicator_jet`, which uses the CM inverse-series and Frobenius packing argument. An estimate that bounds every coefficient above cannot provide a nonzero $K_\alpha^{(12)}(g)$ of the necessary local size. The uniform exact normalized Rallis pairing for the rank-one values and their conjugate norm formula (`LValueData.map_trivialLift`, `Data.Pair.log_norm_map_norm_sq_trivialLift`) must also be supplied by the actual CM theta model. Equations (19.21) and (19.23) show what has been reduced to elementary finite arithmetic and where these two sources of information enter the height argument. With those genuine inputs, plus the testedness and the rank-one projection bound, the calculation in Chapter 17 gives the height saving of one direction.

19.17 Summary of the strengthened reconstruction in C

The following list distinguishes the completed constructions of this chapter from their remaining geometric inputs.

1. The local algebraic factorization through column partners, including the two-sided lattice coset representatives and compactness proof, is now established (Proposition 19.2); the synthesis of a stable space from unimodular coefficient vanishing is proved in Proposition 19.4. The required automorphic equivariance and actual geometric coefficient vanishing must hold on the normalized theta difference, not just on the finite data.

2. The real smooth local Witt map, simultaneous moment approximation and smooth-section estimate are proved, and the determinant-block calculation gives the positive scalar of the cone integration. The exact top-period identity, normalized mixed-model line constant term and its convergence and global invariant measures still have to be supplied for the

compatible data. The scalar induction and double-exchange argument explain the nonzero common normalization but rely on these genuine global identities.

3. The finite tidy square transformation, including the removal of the frame-orbit denominator by covariance, is now proved; the terminal interlacing, signatures, and square-cost recursion are explicit. The genuine real square law with its kernel height and type-convolution inputs, and the global quantitative restriction map reaching all initial rows, are still used from the specified declarations. The latter is needed before the finite transport can establish the crude-row theorem.

4. The divided lattice gluing, including its realization at the split places, level conjugation, indicator identity, and sharp bad-column denominator bound are proved (Propositions 19.13 and 19.14). The finite assembly and scalar/surjectivity argument for the actual comparison operator are expanded, as is the Archimedean budget calculation. They use the genuine block theta evaluation, line-scale estimate and fixed-carrier separation, and the testedness requires the exact tower-vanishing period conversion.

5. The reflection sum and the small-Hecke-root induction have been made explicit; the latter still needs boundedness after subtraction of the correct constant term, the genuine Hecke transfer, and its period conversion. The finite character measure and flat-block test pairing arithmetic are reconstructed subject to the stated normalized theta and point-form identities. The depth-indicator attainability and normalized rank-one Hecke values and their Galois comparisons remain the specialized arithmetic inputs.

These are concrete advances beyond identifying the modules in which the original proofs occur. In particular the finite gluing and square transformations do not require a rational theta theorem or an assumption that the desired integral bounds already hold. The unresolved parts, on the other hand, are not removable merely by citing a classical theta correspondence in a stable range. They are precisely the analytic and normalized integral comparison statements for the genuine data of the proof.

Chapter 20

Congruences, auxiliary primes, and the fit family

Chapters 21 and 22 reconstruct the auxiliary-prime and triangular relations, the finite duality and matrix-span arguments, the odd-class and twisting estimates, and the Fitting and first-order calculations. In particular the local algebraic implications should be distinguished from the still substantive existence of the compatible automorphic/Galois data.

We now address the integrality issue in the determinant calculation. The notation for ranks and form spaces is as above. In this section a row of a Galois representation and a column of such a representation refer to the distinguished block and the entire representation, respectively, not to vectors in W and C. Let $B_0 = B/(mk)$, where $B = fa(\log b)^3$. We explain the reduction of the existence of a fit family to a global twisting theorem, a first-order congruence theorem, and the local crystalline estimates. The distinction between established linear algebra and the remaining specialized number-theoretic statements is important. Assertions identified as *formal-source inputs* below are precise imports from the accompanying Lean proof, not results independently proved in this manuscript. Their proofs require the auxiliary constructions indicated by the source references.

20.1 The integral projection problem

Consider an auxiliary unitary pair of odd source rank $r' \le r$ and odd rank difference $t' \le t$; its definite target has rank $r' + t'$. Write $\mathfrak{A}$ for its finite dimensional space of algebraic forms, with the class-sum pairing, and $Z \subseteq \mathfrak{A}$ for an image of theta lifting. Let P_Z be the orthogonal projection. The auxiliary pair has a finite level group, a chosen group-algebra idempotent e, and commuting split-place Hecke polynomials $H_w(X)$. As in Chapter 17, their test value is denoted X_w. The condition imposed on a form v is

$$v \perp (Z^{\perp} \cap \operatorname{im} e \cap \ker H_w(X_w)) \quad \text{for every split } w \text{ outside the level and its conjugate.} \tag{20.1}$$

It is essential that the subspace in (20.1) involves a kernel at each such place, rather than merely the space of theta lifts. These kernels can contain other forms with a Galois representation resembling that of a lift. We call a form satisfying (20.1) tested. The source definitions are `AuxPair.IsTested` and `DenominatorLe` in `New/Selmer/FitTwists.lean`.

Fix a prime ℓ and an abstract field isomorphism $\iota : \mathbb{C} \simeq \overline{\mathbb{Q}}_\ell$. The absolute value on the latter is normalized by $|\ell|_\ell = \ell^{-1}$; put $v_\ell(z) = -\log|z|_\ell/\log \ell$. For an integral coefficient ring $\mathcal{O}$ in a finite extension of $\mathbb{Q}_\ell$, let $\lambda_{\mathcal{O}}$ be its ordinary module length and $e_{\mathcal{O}} = \lambda_{\mathcal{O}}(\mathcal{O}/\ell)$. We use $\lambda = \lambda_{\mathcal{O}}/e_{\mathcal{O}}$ for normalized length. Thus $\lambda(\mathcal{O}/(a)) = v_\ell(a)$ for nonzero $a \in \mathcal{O}$. The

conventions extend to the value $+\infty$ when necessary. Integral forms carry the lattice defined by their Weyl-coordinate values; integrality at all ℓ-adic embeddings is distinguished from integrality at one fixed embedding.

If $q = \dim_{\mathbb{C}} Z$, the denominator condition at ι with allowance β is

$$\left|\iota\left(\det\left(\langle u_a, P_Z v_b\rangle\right)_{a,b=1}^{q}\right)\right|_{\ell} \le \exp(qr'\beta) \tag{20.2}$$

for all $u_a \in \mathfrak{A}$ integral at every ℓ-adic embedding and $v_b \in \mathfrak{A}$ integral at ι and tested as in (20.1). The conditions concern the determinant of Hermitian pairings, not the norm of an arbitrary matrix of coefficients. In particular, r' is the rank of the source unitary group, distinct from q; the factor $r'\beta$ represents the loss allowed per form in the determinant. If $q = 0$ the empty determinant is one. For the good lifts one uses $e = e_0$ and Z the image of the auxiliary lift on good forms. For the rank-one comparison pairs one uses the tuple of labels $X - 1$, $r' = 1$, and Z the line generated by the particular trivial lift. The subspace in (20.1) is then the corresponding kernel of descent.

The height condition on the special values concerns ratios of twisted to untwisted Hecke values. Denote by $\rho_p(\vartheta, s)$ the ratio `Data.Pair.ratio` for a ramified pair p and a sign datum s. The letter p in a pair-label is not necessarily a rational prime. We require ℓ-adic bounds on the inverse ratios for all such pairs. Recall that `Data.Pair.IsRamified` specifies that the tame character of a ramified pair is nontrivial exactly at the chosen places above $g_0 g_1$.

Definition 20.1. For an admissible datum $\mathcal{D}$ a fit family consists, for every ι (including the choice of ℓ), of an m-fold twist θ_ι with conductor prime to ℓ. For every $1 \le j \le m$ there exists a family of real allowances $\beta_{j,\ell}$ such that

$$\sum_{\ell \in T} \beta_{j,\ell} \le B_0 \quad \text{for every finite set } T \text{ of rational primes,}$$

and (20.2) holds at the character θ_ι^j both for the good auxiliary lift at step j and for all rank-one pairs of index j. For every j, p, s as above, the family of inverse ratios $\rho_p(\theta_\iota^j, s)^{-1}$ has height at most x^{999} in the sense of the height predicate of the determinant argument.

In `Data.IsFitFamily`, the index j is a lift index everywhere. The rank-one pairs in `Data.Pair.IsRankOne` have integer parameter z with $m < z \le r$, prescribed tame part ω^j, and weights between $w_j + r - t$ and w_j. All conditions are imposed for the same character and its powers. The definition does not require a single twist working for all primes, nor does it allow independent choices for the arrows of one chain.

Theorem 20.2 (Fit family)**.** *Every admissible datum has a fit family.*

The intended proof is `exists_isFitFamily` in `New/Selmer/FitTwistsExist.lean`. We give its quantitative deduction from the other inputs in this section. An integrality estimate of this strength is necessary: a bound on a projection over $\mathbb{C}$ says nothing about its image of an integral form at ℓ, and a global estimate with an uncontrolled contribution from a new auxiliary conductor need not be usable in the final height sum.

20.2 Labels and Galois representations of lifted eigensystems

We summarize the normalization used for the attached Galois representations. Put $G_E = \mathrm{Gal}(\overline{E}/E)$, let c denote complex conjugation in the ambient absolute Galois group, write $g^c = cgc^{-1}$, and put $\epsilon_\ell = \chi_{\mathrm{cyc},\ell}$. Thus ϵ_ℓ^{-e} is the character of $\mathcal{O}(-e)$. A character of $U(1)$ is

transferred to an anticyclotomic character of G_E by base change and class field theory; ψ denotes the avatar of its inverse. For the auxiliary pair the source rescales the representation attached to a form by a particular algebraic character, `chiOneInv`. If A is a place over a split unramified $w \nmid \ell$ and g a geometric Frobenius normalized as in the source, the scalar of this character is $\iota(a_O(w))^{-1}$, where $a_O(w)$ is the Frobenius character value of the pair (distinct from the Hecke test point X_w). This normalization is designed so that the string in the theta-lift Hecke relation becomes the fixed Tate string $1, \epsilon_\ell^{-1}, \dots, \epsilon_\ell^{1-t'}$.

The key input is not merely an assignment of semisimple Galois representations to eigensystems. It is the simultaneous assignment to all the lifted systems of a row representation T of rank r', independent of the twist, with prescribed inertia and jumps, and an identity of traces for the whole normalized representation U:

$$\operatorname{tr} U(g) = \psi(g) \operatorname{tr} T(g) + \sum_{e=0}^{t'-1} \epsilon_\ell(g)^{-e} \qquad (g \in G_E). \tag{20.3}$$

For the first block the field of coefficients and lattice can depend on ℓ and ι; the row T is the one associated to the untwisted system. It must be possible to work with integral lattices that satisfy the balance and matrix-span hypotheses below.

The labels at a chosen place v are monic squarefree polynomials c_v of degree r' over its residue field. The roots label residue-field characters, and twisting by the inverse tame character scales the roots by the corresponding residue-field element. The label thus modified is denoted $\widetilde{c}_v$. The aptness conditions include $1 \notin \operatorname{Roots}(\widetilde{c}_v)$; for $r' = r$, at least one primitive label occurs above each of g_0, g_1, while for $r' = 1$ the unmodified label used is $X - 1$. The ranks r' and t' are odd. The other conditions specify the chosen places, the extra tame level, and the tame part $\omega^j \vartheta$ of the auxiliary character. This is `IsApt` in `New/Selmer/AptTupleDefs.lean`; aptness is checked for the good and rank-one pairs by `isApt_auxPair` and `isApt_rank_one` in `New/Selmer/HeckeStableSubspaces.lean`.

Here are the row conditions that enter the congruence argument. Let M_g denote the integer `Data.chosenGroupOrder`, the order of the finite product of general linear groups used to bound automorphism and level denominators; in particular

$$g_i^a - 1 \mid M_g \quad (i = 0, 1,\ 1 \le a \le rt). \tag{20.4}$$

Define

$$c_\ell = v_\ell(M_g) + 2^{-\ell}.$$

A row is continuous, trace-nondegenerate, and conjugate self-dual with the normalizing Tate shift: in the rank-difference convention of the source the trace identity has the factor $\epsilon_\ell(g)^{1-t'}$. It is unramified off ℓg and crystalline at primes above ℓ prime to g, potentially crystalline at the other stipulated places, with the indicated multiset of jumps. The row type requires, for each complex embedding, jumps either in a low interval bounded by h_2^2 in absolute value (and bounded below by $h_2/2$), or within t of the appropriately signed w_j. The number of jumps with absolute value at most $2h_2^2$ is bounded by $r'm/k$, and the number of negative jumps at conjugate embeddings together is at most r'. Further:

1. At each place above g not above ℓ, some tame inertia generator γ (the specified one at the chosen places) has determinant bound

$$v_\ell(\det(T(\gamma) - 1)),\ v_\ell(\det(T((\gamma^c)^{-1}) - 1)) \le r' c_\ell.$$

   The inertia action is killed by an integer prime to the residue characteristic at that place.

2. In all sufficiently large coefficient rings, each lattice representation $T_{\mathcal{O}}$ spans the full algebra $M_{r'}(\mathcal{O})$ as an $\mathcal{O}$-module by its group elements. There is a Galois element having scalar action and cyclotomic character $g_i^{r'}$ for one of the two chosen primes.

3. There is also a scalar element in inertia above g off ℓ, with scalar s such that all powers $(s^{\ell^a})^j$ for $1 \le j \le 4r'+2$ differ from one by a unit, for all sufficiently large a as in the source. In particular one can arrange $s-1$ and $s+1$ to be units.

See `IsColumn`, `IsRow` and `HasRow` in `New/Selmer/AptTupleDefs.lean`, and the row-type definition in `New/Selmer/DeterminantBound.lean`. The precise phrase "for all sufficiently large coefficient rings" is the eventual assertion over the directed collection of ℓ-adic coefficient rings; it permits further finite extension after choosing the twist or an auxiliary level. We do not identify the algebraic closure with one finite coefficient ring.

Proposition 20.3 (Formal-source input: rows for all twists)**.** *For every nonzero Hecke eigensystem occurring in the good forms at the final single-pair source W_m, and every auxiliary lift index j, there is a row T as above satisfying (20.3) for each twist of the auxiliary pair and every attached normalized representation of a lifted eigensystem. The corresponding statement holds for every rank-one pair, with the untwisted rank-one system. All relevant row representations can consequently be put in one finite list for a fixed $\ell, \iota, \mathcal{D}$.*

This is `exists_hasRow_auxPair` and `exists_row_for_all_twists` in `New/Selmer/RowsExist.lean`, applied using the apt labels and rank-one statements. The passage from an automorphic eigensystem to a labelled column uses the local-global compatibility packages in `Lib/GaloisRep` and `New/Selmer/LabelledColumns.lean`. The latter obtains the conjugate self-dual inertia polynomials and proves `IsLabelledColumn.isRow`. At a chosen place, for inertial s, there are $r'+t'$ roots y annihilated by $y^{g_i^d-1}=1$ for some $1 \le d \le r$, and the product of $X-y$ kills s and the conjugate inverse on every labelled column of the specified type. For an untwisted row no root is 1 at the distinguished generator. This is `exists_inertia_eigenvalues`. The proof uses the primitive label above a prime different from ℓ to enforce nondegeneracy and the matrix-span condition. The row-type verification uses the explicit highest weights; the bound on the number of low jumps is what controls the crystalline defect in the local part of the argument.

The finiteness assertion deserves emphasis. The space of good forms is finite dimensional, and the appropriate Hecke operators are simultaneously diagonalizable. For a finite-dimensional space the commuting normal Hecke operators admit a common eigenbasis, and therefore only finitely many characters occur. The equality of relevant multiplicities across signatures identifies the list used in constructing the auxiliary pairs. There are only finitely many pairs with the prescribed weights and tame levels: they are characters in finite sets at fixed conductor and prescribed infinity type, and the construction ranges over finitely many integer weights. Thus the row list may be formed before choosing the finite-order twist. The same is not true of a putative bound obtained directly from the infinitely variable twists of all eigensystems.

We record two algebraic details used in this reduction. First, if the target of a one-dimensional polynomial label is a value z of finite order and neither of z^a, z^b becomes equal after raising to any ℓ-power, then $|z^a - z^b|_\ell = 1$. Indeed write $z = z_{\ell'} z_\ell$ as commuting roots of unity of prime-to-ℓ and ℓ-power order. Reduction modulo the maximal ideal kills z_ℓ and is injective on roots of unity of prime-to-ℓ order (by the separability of X^n-1 and the uniqueness of a lift of the root 1 in the principal unit group). The premise says $z_{\ell'}^a \ne z_{\ell'}^b$, so their difference is nonzero on reduction. Second, if traces of two representations agree on every group element in characteristic zero and each has nondegenerate trace pairing on its image

algebra, their representations are conjugate after matching multiplicities; this is the semisimple trace argument used to turn (20.3) into block diagonal lattices for the distinguished systems. To see the role of trace nondegeneracy, consider finite sums in the group algebra. If such a sum acts as zero in one representation, pairing its product with every group element gives trace zero in the other; nondegeneracy implies that it acts as zero in the other as well. Thus the image algebras identify with their trace functionals. Over an algebraically closed field the multiplicity of each simple module is determined by this functional, giving equivalence. For the source proof the relevant finiteness and trace-nondegeneracy properties of the attached representations are part of the input, not automatic from comparing unramified Euler factors alone.

20.3 Selmer modules over finite coefficient rings

For a stable lattice $T_{\mathcal{O}}$ with row action $R : G_E \to \mathrm{GL}_{r'}(\mathcal{O})$ and a finite-order character ψ with values in $\mathcal{O}^\times$, let V_M be the twist of $(\mathcal{O}/\ell^M)^{r'}$ by $\psi \otimes R$. Its continuous cocycles obey

$$z(gh) = z(g) + (\psi R)(g)z(h).$$

For a finite set of places S stable under conjugation, and a prescribed family $\mathcal{L}$ of finitely supported linear relations on functions on G_E, write

$$\mathfrak{S}_M(\psi R; S, \mathcal{L}) = \left\{ z \in Z^1(G_E, V_M) : \ z(I_v) = 0 \ (v \notin S), \ z \text{ obeys } \mathcal{L} \right\}.$$

At primes over ℓ the family $\mathcal{L}$ is the collection of linear relations obeyed by the characteristic-zero crystalline cocycles Z^1_f of the chosen integral lattice, pulled from the decomposition groups to G_E; it is not the assertion that each torsion cocycle itself is the reduction of a crystalline cocycle. This weaker formulation is important in the local comparison; it is `crysRels` in `New/Selmer/DeterminantBound.lean`. The ambient space is a module of actual cocycles, not cohomology classes. The source `SelmerLengthLe` normalizes by one evaluation $z(\gamma) = 0$, where $\psi R(\gamma) = sI$ and $s - 1$ is a unit, $\epsilon_\ell(\gamma) = 1$.

The normalization has a simple conceptual meaning. If any module of local conditions is stable under adding coboundaries, the evaluation at γ maps its coboundary subgroup isomorphically onto V_M: the coboundary defined by $v \in V_M$ is $\delta v(g) = ((\psi R)(g) - 1)v$, and $\delta v(\gamma) = (s - 1)v$. Thus for every cocycle z there is a unique $v = (s-1)^{-1}z(\gamma)$ such that $z - \delta v$ vanishes at γ. The normalized cocycles then represent the corresponding Selmer classes uniquely. This argument explains why an actual cocycle-module bound can stand in for a more familiar cohomological bound. The source includes the scalar and the normalization in its definitions because the use of both at the integral level is essential later.

Several times we shall use the following elementary length facts. Over a DVR, if $A \in M_N(\mathcal{O})$ has nonzero determinant, the kernel of A on $(\mathcal{O}/(a))^N$ has length at most $\lambda_{\mathcal{O}}(\mathcal{O}/(\det A))$. This follows by writing A in Smith normal form with diagonal entries π^{e_i} up to units; the kernel on each factor has length $\min(e_i, v_\pi(a))$. If $Y, Z \subseteq V$ are submodules of a finite-length module, the length lost by intersection with Z equals $\lambda(Y/(Y \cap Z))$, and such relative lengths are subadditive under successive intersections. Thus if each new linear condition takes values in a torsion module of length at most c_j, imposing all of them decreases length by at most $\sum c_j$.

Definition 20.4 (Balanced lattices)**.** A representation R over a coefficient ring $\mathcal{O}$ is balanced up to a constant β_0 at depth M for a ramification set S, relations $\mathcal{L}$, and normalizer γ if it has the following properties. Its reduction modulo ℓ^M is locally constant; it is unramified off S; its group matrices span $M_{r'}(\mathcal{O})$; the residue field has more than $4r' + 2$ elements. The

distinguished inertial element γ lies above S off ℓ, with $R(\gamma) = R((\gamma^c)^{-1}) = sI$ and $s \pm 1$ units. The coboundaries obey $\mathcal{L}$ modulo every divisor of ℓ^M. Finally there are $g \in G_E$ and scalars α, α' with $R(g) = \alpha I$, $R((g^c)^{-1}) = \alpha' I$, such that for each $a \mid \ell^M$ the local cocycle conditions satisfy

$$\begin{aligned}\sum_{v \in S} \inf_{A|v} \lambda_{\mathcal{O}}(Z^1_v(\mathcal{L}, \mathcal{O}/(a))) + r' v_\pi(\alpha - \epsilon_\ell(g^c)\alpha') \\ \leq (|S| + f)\lambda_{\mathcal{O}}((\mathcal{O}/(a))^{r'}) + \beta_0.\end{aligned} \tag{20.5}$$

Here v_π is the valuation equal to one on a uniformizer of $\mathcal{O}$ and Z^1_v denotes the cocycles on the local decomposition group obeying the induced relations. The local fields and closures are chosen as in the global representation. This is the predicate `IsBalanced` of `New/Selmer/Corepresentable.lean`.

Proposition 20.5 (Formal-source input: balanced row lattices)**.** *Fix a row T of rank r' and a chosen place w above g not above ℓ, with tame generator z. For every sufficiently large coefficient ring and depth M there is a lattice representation R of T, a normalizer γ and β_0 making it balanced for ramification at the specified level and the local crystalline relations. They satisfy*

$$\frac{\beta_0}{e_{\mathcal{O}}} + v_\ell \det(R(z) - 1) + v_\ell \det(R((z^c)^{-1}) - 1) \leq 10 f r' c_\ell. \tag{20.6}$$

This is `IsRow.eventually_isBalanced` in `New/Selmer/RowBalanced.lean`; the finite ramification set may be taken as the places over ℓg embedded in the common exceptional set of primes. The proof relies on a local-global comparison of dimensions of crystalline cocycles and the scalar difference. The arithmetic point is the factor $10 f r' c_\ell$ rather than a bound in terms of the depth M: the depth will be much larger than this and is not an available contribution to the determinant allowance. We retain (20.5) to identify exactly where local conditions enter the subsequent use of global duality.

20.4 New level characters and tied cocycle conditions

Let S_0 be a fixed finite exceptional set containing the original ramification places, the places over ℓ, and the places of the prescribed levels that must not be used as auxiliary places. A finite set Q is auxiliary if each place $P \in Q$ has prime norm, both P and P^c lie off S_0, and $P^c \notin Q$. Thus the rational place below P is split in the relevant quadratic extension, and the auxiliary conditions at the two conjugate sides can be specified compatibly. The formal definition is `IsAuxiliary`. We reserve Q for the modifying conductor of an actual twist, and J initially for a set of auxiliary places carrying further linear conditions. The subsequent Taylor–Wiles places used in the projection argument are different sets.

For a depth N write $q_N = \ell^N$. Auxiliary prime data consist of a place P with $\mathrm{N}P \equiv 1 \pmod{q_N}$, a chosen geometric place above it, a tame inertia generator τ_P, an arithmetic Frobenius φ_P, and an anticyclotomic homomorphism

$$h_P : G_E \longrightarrow \mathbb{Z}/q_N\mathbb{Z}, \qquad h_P(\tau_P) = 1,$$

trivial on the prohibited local subgroups and of conductor supported on P, P^c. The condition of coming from $U(1)$ includes the local triviality and anticyclotomic condition needed for the subsequent passage to tame characters. Frobenius and inertia nomenclature follows the conventions of `AuxPrimeData` and `IsAnticyclotomicAt` in `New/Selmer/Corepresentable.lean`.

Proposition 20.6 (One anticyclotomic prime (see Proposition 21.9))**.** *Let z lie in the normal closure of the inertia groups above a finite set R outside the original prescribed set (and its conjugate) and the places over ℓ. Given a neighborhood U of z, and any N, there is a new prime-norm place P outside $R \cup R^c$ with auxiliary prime data at depth N, with $\varphi_P \in U$, $h_P(\varphi_P) = 0$, and h_P trivial at the specified old local places and above ℓ.*

This is the precise auxiliary-prime mechanism of `exists_isAnticyclotomicAt` in `New/Selmer/Corepresentable.lean`. Its proof uses the finite generation of the group of S-units. A Kummer extension is formed by adjoining q_N-th roots of a finite generating system and the q_N-th roots of unity. Inertia at allowed places fixes these roots as required. Chebotarev then produces a prime-norm place with Frobenius close to z and whose residue field satisfies the splitting conditions. The local residue-field unit logarithm, sending a suitable tame unit to 1 modulo q_N, is trivial on the image of the prescribed global units, so it descends to the idelic quotient after extending along the chosen ideal classes. Taking the quotient by complex conjugation realizes an anticyclotomic class-field character. One finally modifies the representative of Frobenius by an inertia element to obtain $h_P(\varphi_P) = 0$ (in additive notation). The compatibility of the global unit condition and the Frobenius prescription is the point of the Kummer construction; it cannot be obtained by selecting an arbitrary prime congruent to 1 modulo ℓ^N.

A pencil of auxiliary places consists of an auxiliary finite set J, a helper $P_0 \in J$, the data just described for every $P \in J$, and elements $t_{P,i} \in \mathcal{O}$ for the different row representations R_i. Besides the auxiliary-place conditions, the following hold:

$$h_{P'}(\varphi_P) = \begin{cases} 1, & P' = P_0, \\ 0, & P' \neq P_0, \end{cases} \quad (P \neq P_0), \tag{20.7}$$

φ_{P_0} lies in a prescribed neighborhood of a chosen element z, and for every i with $t_{P,i} \neq 0$ both φ_P and $(\varphi_P^c)^{-1}$ act trivially on R_i modulo ℓ^M. The tied conditions for row i are

$$\begin{aligned} c(\tau_P) &= t_{P,i}\, c(\varphi_P), \\ c((\tau_P^c)^{-1}) &= t_{P,i}\, c((\varphi_P^c)^{-1}). \end{aligned} \tag{20.8}$$

They are linear relations on cocycles unramified outside $S \cup J \cup J^c$. Notice the special role of P_0: although each row uses different slopes, the chosen auxiliary places and their characters will be common to all rows. This permits one actual twisting character at the end. The source definitions are `IsPencil` and `tieRelations`.

A basic property of these pencils explains how imposing and then removing ties can be controlled. For a tuple $y = (y_P)$ of q_N-th roots of unity, viewed as units modulo some torsion modulus, define

$$\psi_y(g) = \prod_{P \in J} y_P^{h_P(g)}. \tag{20.9}$$

It follows from the unramified condition at the other primes that $h_{P'}(\tau_P) = \delta_{PP'}$, and anticyclotomicity gives the same values on $(\tau_P^c)^{-1}$. Thus ψ_y has value y_P at the two tame generators at P, P^c, and by (20.7) has the value y_{P_0} at both of the appropriate Frobenius elements when $P \neq P_0$. This is `IsPencil.pointChar_idOrConjInv` in `New/Selmer/TiltingLemma.lean`. The point of (20.7) is that one can choose different points on a line $y_P = 1 + t_{P,i}(u - 1)$, with helper coordinate u, for each row, and still use polynomial interpolation to find one common character.

Lemma 20.7 (Removal of a tie; elementary calculation)**.** *Let V be a finite-length module of cocycles over a DVR, with values in $(\mathcal{O}/(a))^{r'}$. Suppose that at one pair of generators (τ, φ)*

the induced affine actions of τ and φ commute. Suppose also that for the twisted action R', one has $R'(\varphi) \equiv uS$ modulo a and

$$R'(\tau) - 1 = t(R'(\varphi) - 1)$$

on the coefficient module. Then the length lost by imposing the relation $c(\tau) = tc(\varphi)$ is at most $v_\pi(\det(uS - 1))$ (with the infinite case interpreted as vacuous).

Proof. Write the affine action as $v \mapsto R'(g)v + c(g)$. Evaluating commutativity at zero gives

$$(R'(\varphi) - 1)c(\tau) = (R'(\tau) - 1)c(\varphi).$$

Hence the defect $c(\tau) - tc(\varphi)$ belongs to the kernel of $uS - 1$. The quotient of V by the tied submodule identifies with the image of the defect map, and the kernel estimate from Smith normal form bounds its length by the indicated valuation. □

In the application, tame local relations say that Frobenius conjugates a tame inertia element by the power $\mathrm{N}P$. With $\mathrm{N}P \equiv 1 \pmod{\ell^N}$ and N larger than the depth of the coefficients and of the twist, the corresponding affine actions commute modulo ℓ^M. The relation between $R'(\tau)$ and $R'(\varphi)$ on the line follows from (20.9): if the slope is nonzero the original row matrix at Frobenius is the identity modulo ℓ^M, whereas if the slope is zero the twisted inertia action is trivial. At the helper, removing ties uses $u = 1$ instead. Summing Lemma 20.7 over both conjugate generators and the subset of places whose ties are removed gives the exact determinant-loss term in `IsPencil.length_le_of_drop_ties`, in `New/Selmer/TiltingLemma.lean`. The full source statement permits specified replacements for the Frobenius elements as long as their matrices agree modulo ℓ^M. The ramification and affine-action assertions just invoked are hypotheses verified there; in particular one does not divide by ℓ or by residue cardinalities inside the torsion module.

20.5 Shrinking several tied Selmer modules

There are two logically separate ingredients in the existence of a pencil that shrinks every module. One is finite-dimensional linear algebra over a DVR; the other is a reciprocity lemma ensuring a new place on which a long module produces a nonzero detection. We record both, because the arithmetic conclusion does not follow from mere genericity of slopes.

Lemma 20.8 (Tilting a pair of linear maps)**.** *Let $\mathcal{O}$ be a DVR with residue field κ, let $A = \mathcal{O}/I$ have finite length, and let $\Phi_0, \Phi_1 : V \to A^u$ be $\mathcal{O}$-linear. For a finite-length submodule $Y \subset V$, all except at most u residue classes $s \pmod{\mathfrak{m}_\mathcal{O}}$ satisfy*

$$\lambda_\mathcal{O}\big(Y \cap \ker(s\Phi_0 - \Phi_1)\big) \le \lambda_\mathcal{O}\big(Y \cap \ker \Phi_0 \cap \Phi_1^{-1}(\Phi_0 Y)\big). \tag{20.10}$$

Proof. Let Y_1 be the sum of the kernels on the left as s varies. Then $Y_1 \subseteq Y$, and $\Phi_1 Y_1 \subseteq \Phi_0 Y_1$ because this holds on each summand. Let $L \subseteq \mathcal{O}^u$ be the inverse image of $\Phi_0 Y_1$ under reduction modulo I. Since L is free of rank at most u, lift the quotient map $\pi : L \to \Phi_0 Y_1$ to a map $\sigma : L \to Y_1$ with $\Phi_0 \sigma = \pi$. Lift $\Phi_1 \sigma$ to an endomorphism T of L using the same projectivity. For every s whose reduction is not a root of the characteristic polynomial of the reduction of T, the endomorphism $sI - T$ of L is invertible. Therefore

$$(s\Phi_0 - \Phi_1)\sigma(L) = \pi((sI - T)L) = \Phi_0 Y_1.$$

The image of $s\Phi_0 - \Phi_1$ on Y_1 is contained in $\Phi_0 Y_1$, hence equal to it. Its kernel on Y_1 has length $\lambda_\mathcal{O}(Y_1) - \lambda_\mathcal{O}(\Phi_0 Y_1) = \lambda_\mathcal{O}(Y_1 \cap \ker \Phi_0)$. The kernel on Y is included in Y_1 by definition. Finally $Y_1 \cap \ker \Phi_0 \subseteq Y \cap \ker \Phi_0 \cap \Phi_1^{-1}(\Phi_0 Y)$, giving (20.10). □

This is the lemma `tilting`, in `New/Selmer/TiltingLemma.lean`. In the arithmetic application $u = 2r'$, because both conjugate sides are included in each evaluation. Avoiding the exceptional slopes at two different moduli, and also the slope zero, requires at most $4r' + 1$ exceptional residue classes. This explains the residue-field requirement in `IsBalanced`.

The reciprocity statement used in conjunction with this lemma is the following special form of a global-duality obstruction.

Proposition 20.9 (Formal-source input: the mirror criterion)**.** *Let R and the pencil be as above, with R balanced up to β_0, and $a \mid \ell^M$, $a \neq 0$. Let $X(J, a)$ be the corresponding normalized module of cocycles obeying the ties modulo a. Suppose some element of $X(J, a)$ has a unit coordinate. If for every admissible enlargement by a new prime P, compatible with the pencil for every choice of its new slope, the image of $X(J, a)$ under evaluation at the two Frobenius elements lies in the image under tame-inertia evaluation of the relaxed module (in which the new ramification is allowed but its ties not imposed), then*

$$\lambda_{\mathcal{O}}(X(J, a)) \leq \beta_0.$$

This is `IsPencil.length_le_of_unseen_class` in `New/Selmer/TiltingLemma.lean`. A unit-valued coordinate is used to detect a nontrivial row in the full matrix algebra. The detailed proof constructs dual cocycles via global reciprocity, with isotropic pairings against the class in question. The dual classes, evaluated on a suitable open subgroup on which the row and cyclotomic character act trivially modulo the torsion modulus, are bounded by the size of the torsion subgroup for the scalar difference $\alpha - \epsilon_\ell(g^c)\alpha'$ from (20.5). The "mirror" estimates use the full span of the matrix algebra and the units $s \pm 1$ to constrain that image. The algebraic assertion is `card_le_torsion_of_isotropic` in `New/Selmer/LinePointPowers.lean`; the reciprocity construction and its local restrictions are developed in the proof of the mirror criterion. The determinant or torsion size term is absorbed by the balance bound, not by a global bound on all cohomology classes.

Proposition 20.10 (Boundedly many ties; formal-source input)**.** *Fix finitely many balanced representations R_i at depth M, a neighborhood of an allowed inertial element z, and a divisor a of ℓ^M. There is an integer J_0 such that for every depth N of the auxiliary characters there is a pencil of at most J_0 places with normalized cocycle modules satisfying*

$$\lambda_{\mathcal{O}}(X_i(J, a)) \leq \beta_{0,i} \qquad \text{for every } i.$$

This is `exists_isPencil` in `New/Selmer/TwistTheorem.lean`. The important quantifier is that J_0 is chosen before N. To explain its proof, one inducts on the length of $\mathcal{O}/(a)$. Lifting a good pencil for a smaller divisor a to πa gives a finite a priori bound on the length of the larger-depth modules, by evaluating on the finitely many tame generators. For the larger divisor minimize the total length of all the tied modules subject to preserving the smaller-depth inequalities and an upper bound on the sum of the number of places and this total length. If a row has normalized module longer than the allowed bound, and there is no unit-valued element, it reduces to the smaller depth: multiplication by π identifies the modules killed by the shorter divisor, preserving all the linear relations. If there is a unit-valued element the mirror criterion applies. Unless its conclusion already gives the desired bound, it supplies a new place where the image under Frobenius evaluation is not contained in that of the relaxed module under tame-inertia evaluation.

At such a place, write Y for the relaxed module, Φ_0 for tame-inertia evaluation, and Φ_1 for Frobenius evaluation. Imposing $\Phi_0 = 0$ recovers the original unramified module. The failure of the mirror inclusion gives a strict decrease when imposing the additional condition $\Phi_1(c) \in \Phi_0 Y$. Lemma 20.8 yields a slope giving at least this strict decrease at the larger

depth, and by avoiding the exceptional classes at the smaller depth as well it preserves the smaller-depth inequality. The new slopes for all other rows can be set to zero, so their modules do not increase. Consequently the minimized total length strictly decreases, while the sum of length and number of primes does not increase. This is the contradiction with minimality. The induction bounds the required number of new places independently of the auxiliary-character depth. We have described the induction's length mechanism in order to identify where the key reciprocity input is used; the actual selection of primes satisfying all neighborhood constraints together is part of the cited formal-source assertion.

20.6 Finite presentations and interpolation of lengths

To convert the linear pencils into common characters, one must avoid choosing a different character for every row. The device is to bound the length of each Selmer module by a polynomial in the tame-character parameters. The polynomial need not have a degree bounded in terms of the original rank. A special feature of the root-of-unity argument below is that it handles arbitrary polynomial degree uniformly, after a common finite bound on the number of variables is fixed.

Here is the underlying finite-presentation lemma. Fix a modulus $m_0 \neq 0$, a representation given by matrices over a coefficient ring Λ, a finite ramification set, a subfamily of linear relations L, and an open neighborhood of the identity on which all the relevant actions on the m_0-torsion coefficient modules are trivial. Continuous cocycles of these types can be described by finitely many coordinates (the values at chosen group elements) obeying a submodule of relations in a finite free Λ-module. One standard justification is to choose an open normal subgroup on which the action is trivial, restrict the cocycle there, and observe that its restriction is a homomorphism to a module killed by m_0. The maximal abelian exponent-m_0 quotient of the subgroup unramified off the prescribed finite set is finite, by global class field theory (or the corresponding finite-character finiteness theorem). Thus there is one common finite quotient through which all such cocycles factor, independently of their coefficient module. More explicitly, the restriction on the open subgroup is a homomorphism with the prescribed ramification; it factors through the common quotient because any nonzero value in a torsion module is detected by a homomorphism to $\mathbb{Q}/\mathbb{Z}$ (and hence to the subgroup killed by m_0). Taking normal cores yields the desired finite quotient of the original group. On its finite set of elements, the cocycle relations, inertia-zero relations, and the imposed linear relations are simply linear equations. This is the principle of `exists_corepresenting_coords` in `New/Selmer/Corepresentable.lean`. Its source proof also handles the universal coefficient ring by constructing the relation module directly.

Let the character parameters be $Y = (Y_1, \dots, Y_q)$, with $h : G_E \to (\mathbb{Z}/\ell^N)^q$ a finite family of homomorphisms. Use the polynomial ring $\Lambda[Y]$ for the pointwise operators with coefficients $\prod_j Y_j^{h_j(g)}$ (least nonnegative lifts of the exponents). At an actual tuple of roots of unity, or modulo a depth at which the tuple is a root of unity, these operators represent the twist. The cocycle condition and all the other local conditions give a submodule of relations $R(Y) \subseteq \mathcal{O}[Y]^u$ which contains $\ell^M \mathcal{O}[Y]^u$. This also makes sense at arbitrary tuples: even if the formal operators fail to be a representation there, the algebraic equations define a submodule, and we only need their specialization at true or approximate characters.

Lemma 20.11 (Sharp Fitting polynomial)**.** *Let $R \subseteq \mathcal{O}[Y]^u$ contain $\ell^M \mathcal{O}[Y]^u$. For $Y = y$, write R_y for the specialization and define*

$$X_y = \{\xi \in (\mathcal{O}/\ell^M)^u : \langle r, \xi \rangle = 0 \text{ for every } r \in R_y\}.$$

Fix y_0 over a DVR $\mathcal{O}$. There is $F \in \mathcal{O}[Y]$ with

$$v_\pi(F(y_0)) = \lambda_{\mathcal{O}}(X_{y_0}), \qquad \lambda_{\mathcal{O}'}(X_y) \le v_{\pi'}(F(y))$$

for every specialization over a DVR $\mathcal{O}'$ above $\mathcal{O}$. Here F is mapped to the appropriate coefficient ring in the second formula.

Proof. Over $\mathcal{O}$, choose a basis of the full-rank sublattice $R_{y_0} \subseteq \mathcal{O}^u$. Each basis element is a finite $\mathcal{O}$-linear combination of evaluations of elements of R, so we can lift the basis to u polynomial rows of R. Take F to be their determinant. At y_0, Smith normal form gives elementary divisors π^{s_i}, all dividing ℓ^M because R_{y_0} contains $\ell^M\mathcal{O}^u$. The annihilator modulo ℓ^M has length $\sum_i s_i = v_\pi(F(y_0))$. At any other specialization X_y lies in the kernel modulo ℓ^M of the same matrix of specialized rows. If its determinant is nonzero, the kernel bound by Smith normal form gives the inequality; if zero, the inequality is vacuous. This proof also works if R was presented by an infinite set of relations, since only the finite chosen rows enter the determinant. □

This is the algebraic mechanism of `exists_sharp_fitting_poly` in `New/Selmer/LinePointPowers.lean`, where R is obtained by the preceding corepresentability lemma. The equality rather than a mere upper bound at y_0 is essential: the tied Selmer theorem controls the actual module length, not an arbitrary chosen determinant. The lemma applies also after imposing a selected subset of ties; its sharp specialization can be different for every row.

Given a root of unity $u^{\ell^k} = 1$ and $t \in \mathcal{O}$, the point $y = 1 + t(u-1)$ is an approximate root of unity:

$$\ell \mid (y-1)^{\ell^k}, \qquad \ell^M \mid y^{\ell^N} - 1 \quad (N \ge k + M). \tag{20.11}$$

Indeed in characteristic ℓ, $(u-1)^{\ell^k} = u^{\ell^k} - 1 = 0$, so $y^{\ell^k} \equiv 1 \pmod{\ell}$. If $v \equiv 1 \pmod{\ell^j}$ with $j \ge 1$, the binomial expansion gives $v^\ell \equiv 1 \pmod{\ell^{j+1}}$; iterate this argument. This is `dvd_one_add_mul_pow_sub_one`. In particular, the rows can each be specialized on their own line of approximate characters without giving up the interpretation of the matrices modulo ℓ^M.

20.7 Degree-free selection of roots of unity

We give the polynomial selection argument, since it supplies the main transition from the tied linear conditions to a single actual anticyclotomic character. It is also an example of a quantitative lemma whose conclusion is uniform in the degrees of arbitrarily many high-degree determinants, an essential feature here.

For a polynomial F over the integers of a finite ℓ-adic extension, and $0 < R \le 1$, write $|F|_R = \max_i |a_i|_\ell R^i$ for its Gauss norm, where $F(X) = \sum a_i X^i$. It is multiplicative and nondecreasing in R. We use this multiplicativity over the finite extensions containing the roots of unity in the following proof.

Lemma 20.12 (Simultaneous one-variable avoidance)**.** *Fix $0 < R < 1$ with $|\ell|_\ell < R^{\ell-1}$, and finitely many positive numbers u_i. For all sufficiently large N, in every sufficiently large coefficient ring containing the ℓ^N-th roots of unity, the following holds: if F_i are integral polynomials, of arbitrary degrees, and $|F_i|_R \ge u_i$, then there is a single ℓ^N-th root of unity ζ with*

$$|F_i(\zeta - 1)|_\ell \ge R u_i \qquad \text{for every } i.$$

The required depth is independent of the degrees.

Proof. Put $F=\prod_i F_i$ and $Q_i=\prod_{h\neq i}F_h$. Suppose, for contradiction, that at every ℓ^N-th root of unity at least one inequality fails. Take an integer $a\geq 1$ so large that $R^a<|\ell|_\ell$. For $0\leq j<N$ set $R_j=R^{1/\ell^j}$ and $n=\ell^{j+1}$. The roots α of $P_j(X)=(1+X)^n-1$ all have absolute value less than R_j, since a primitive root of order ℓ^{j+1} satisfies

$$|\alpha|_\ell=|\ell|_\ell^{1/(\ell^j(\ell-1))}<R_j.$$

For each such root choose $i(\alpha)$ with $|F_{i(\alpha)}(\alpha)|_\ell<Ru_{i(\alpha)}$. Let $b=a(j+1)$ and form

$$S_0(X)=\sum_{\alpha:P_j(\alpha)=0}\frac{P_j(X)}{X-\alpha}F_{i(\alpha)}(\alpha)^bQ_{i(\alpha)}(X)^b.$$

The roots of P_j are simple. At each root, S_0 agrees with $P_j'F^b$, so their difference is divisible by P_j. Write $P_j'F^b-S_0=P_jS$. For $T=R_j$ and $T=R_{j+1}$ we have $|P_j|_T=T^n$ and $|P_j'|_T=|n|_\ell$ (the latter because $P_j'=n(1+X)^{n-1}$). Moreover $u_i\leq|F_i|_T$, so

$$|S_0|_T\leq T^{n-1}R^b|F|_T^b<|n|_\ell|F|_T^b,$$

using $T\leq 1$ and $R^{a(j+1)}<|\ell|_\ell^{j+1}$. The ultrametric inequality and multiplicativity give

$$|S|_T=\frac{|n|_\ell|F|_T^b}{T^n}.$$

Since $|S|_{R_j}\leq|S|_{R_{j+1}}$, it follows that

$$\log|F|_{R_{j+1}}-\log|F|_{R_j}\geq\frac{n}{a(j+1)}\log(R_{j+1}/R_j)=\frac{-(\ell-1)\log R}{a(j+1)}.$$

Summing these inequalities over $j<N$ produces a contradiction for large N: the left hand side is at most $-\log|F|_R\leq-\sum_i\log u_i$, whereas the harmonic sum on the right is unbounded. If a lower bound $u_i>1$ makes the hypotheses impossible there is of course nothing to prove. This establishes the uniform assertion. □

This is the interpolation proof of `exists_root_of_unity_eval_large` in `New/Selmer/RootsOfUnityPoints.lean`.

Lemma 20.13 (Simultaneous selection in several variables)**.** *Fix a positive real δ, positive integers m,a,q, a nonnegative integer L_0, and finitely many nonnegative numbers B_i. For every sufficiently large N and every sufficiently large integral coefficient ring, suppose that $F_{i,j}$ $(1\leq j\leq m)$ are integral polynomials in at most q variables. Suppose there exist tuples $u_{i,j}$ with $\ell\mid(u_{i,j,k}-1)^a$ in each coordinate and*

$$v_\ell(F_{i,j}(u_{i,j}))\leq B_i.$$

Let also D_h be monic one-variable polynomials with total degree at most L_0. Then there is one tuple ζ of ℓ^N-th roots of unity such that simultaneously

$$v_\ell(F_{i,j}(\zeta^j))\leq B_i+\delta,\qquad v_\ell(D_h(\zeta_k))\leq\delta\quad\text{for every } k,h. \tag{20.12}$$

Here the depth can be chosen before the polynomials or tuples are specified.

Proof. Choose $R<1$ so close to one that $R^q,R^{L_0+1}>\ell^{-\delta}$ and $R^{\max(ma,\ell-1)}>|\ell|_\ell$. Assume that some coordinates have already been replaced by ζ_k^j in the points for each pair (i,j), and that the norm bound has deteriorated by at most a factor R per replacement. For the

next coordinate, all other coordinates stay fixed (possibly different for different pairs). The polynomial in X given by substituting $(1+X)^j$ in this coordinate has Gauss norm of radius R at least the currently established lower bound. To justify that assertion, the equation $(1+X)^j = u_{i,j,k}$ has a root with $|X|_\ell \le R$: the monic polynomial $((1+X)^j - u_{i,j,k})^a$ has constant term of absolute value at most $|\ell|_\ell < R^{ja}$; if all its roots had absolute value larger than R, their product would violate this bound. The value at such a root is exactly the previously known value of $F_{i,j}$, bounded above by the Gauss norm. We may use the algebraic closure to compare the norm and do not need the root itself in the coefficient ring.

Include in the one-variable family the polynomial $\prod_h D_h(1+X)$, whose monic leading term gives Gauss norm at least R^{L_0}. Apply Lemma 20.12 simultaneously. This selects one ζ_k for the next coordinate, with loss at most a further factor R, and also gives $|\prod_h D_h(\zeta_k)|_\ell \ge R^{L_0+1}$. Because each factor has norm at most one, this lower bound implies the same lower bound for every factor. After at most q steps, the accumulated factor is at least $R^q > \ell^{-\delta}$. The one-variable lemma is applied with the finite list of lower bounds $\ell^{-B_i}R^u$, $0 \le u \le q$, and R^{L_0}, so its threshold is uniform in the polynomial coefficients and their degrees. This yields (20.12). □

The source declaration is `exists_roots_of_unity_ellVal_le`, in the same file. An important application of the monic polynomials D_h is to $Y^j - 1$: at a selected tuple all the first m powers have nontrivial tame values, with total valuation as small as desired after accounting for the bounded number of variables.

20.8 The global twisting theorem and value families

The twisting theorem controls more than the Selmer modules. It also preserves the valuation of the special Hecke values needed in the norm comparison. The relevant *value family* (`IsFamily` in `New/Selmer/Corepresentable.lean`) is given by continuous unit-valued characters μ_i and a function $V(Q,\psi)$, defined for auxiliary conductors Q and characters. Its untwisted value $V(\varnothing, 1)$ is nonzero; at each fixed Q it is an integral finite linear combination of values of ψ, say

$$V(Q,\psi) = \sum_{a \in A_Q} c_a \psi(g_a), \qquad |c_a|_\ell \le 1. \tag{20.13}$$

In particular, on a character given by (20.9), this specializes an integral polynomial in Y_P. The family satisfies a norm relation in the form of an estimate: if $P \notin Q$ and ψ is anticyclotomic of conductor Q, the inequality

$$v_\ell(1 - \mu_i(\varphi_P)\psi(\varphi_P)^{\pm 1}) \le c \quad \text{for all } i$$

implies

$$v_\ell V(Q \cup \{P\}, \psi) \le v_\ell V(Q,\psi) + 8c. \tag{20.14}$$

For a ramified pair p the source constructs these families by an integral special-value measure, including the factors at the local level places. More precisely, the integral measure F satisfies a distribution relation upon adding v to the modulus: its value is multiplied by a product of at most four factors

$$\prod_{u \in K_0} \prod_{\epsilon \in \{1,-1\}} (1 - (\nu\vartheta)_{\mathrm{bc}}(\varpi_v)^\epsilon q_v^{u-z}), \quad K_0 \subseteq \{0,1\},$$

where z is the integer of the pair and ν its underlying character. The value on a twisted character is $F(1)\,\mathrm{N}(Q)^h \rho_p(\vartheta)^{h+1}$ for some nonnegative integer h. This is the formal-source

input `Data.Pair.exists_isIntegralMeasure` in `New/Selmer/ValueFamilies.lean`; its derivation uses the algebraic CM Hecke special values and integral interpolation. The passage to the fixed-sign ratios and to the Galois-valued functions in (20.13) is `exists_isFamily` in that file. Because the character of a ramified pair is nontrivial at the chosen place w off ℓ, the listed auxiliary characters satisfy

$$v_\ell(\mu_i(z)-1) \le v_\ell(q_w-1) \le c_\ell \tag{20.15}$$

for its tame generator z. The last inequality follows from (20.4). Inverting the special-value ratio changes the valuation inequality to an upper bound on its ℓ-adic norm. No assumption of nondivisibility of every original Hecke value is needed.

Proposition 20.14 (Formal-source input: simultaneous anticyclotomic twisting)**.** *Let R_i be a finite collection of balanced representations with local bounds $\beta_{0,i}$ at depth M, and fix an inertial element z off ℓ as in Proposition 20.6. Suppose $v_\ell(\mu_i(z)-1) \le c$ for the characters of finitely many value families. If*

$$\frac{\beta_{0,i}}{e_{\mathcal{O}}} + v_\ell \det(R_i(z)-1) + v_\ell \det(R_i((z^c)^{-1})-1) \le B_i < \infty,$$

then for every $\eta > 0$ and every m there are an auxiliary set Q and an anticyclotomic character ψ such that, for each $1 \le j \le m$:

1. *the character ψ^j comes from $U(1)$ with conductor Q and is locally trivial at the prescribed old places and above ℓ, and*
$$\sum_{P\in Q} v_\ell(\psi^j(\tau_P)-1) \le \eta;$$
2. *the normalized lengths of the cocycle modules of $\psi^j R_i$ at depth M, with the original crystalline relations and with value zero at the corresponding γ_i, are at most $B_i+\eta$, over all sufficiently large coefficient rings;*
3. *for each value family, $v_\ell V(Q,\psi^j) \le v_\ell V(\varnothing,1) + 8c + \eta$.*

This is `exists_anticyclotomic_twist` in `New/Selmer/TwistTheorem.lean`, with the normalization of lengths as above. To see the use of the preceding lemmas, choose a pencil as in Proposition 20.10, using at most J_0 places. Choose a small allowance $\delta > 0$ with $(8J_0+2)\delta < \eta$. The construction is in two specialization stages. The first uses an ℓ-power root of unity u and the line $y_P = 1 + t_{P,i}(u-1)$ for the ith row, with $y_{P_0} = u$. By (20.11) all such points can be used modulo ℓ^M after the auxiliary depth has been chosen sufficiently large.

Use the sharp Fitting polynomial of Lemma 20.11 for the tied module, with the helper tie omitted. At $u = 1$ the polynomial has valuation bounded by that of the tied Selmer module plus the determinant terms for removing the helper tie. The one-variable avoidance lemma, applied also to the finitely many characteristic polynomials of the matrices at the nonhelper Frobenius places and the Euler factors for the value families, chooses a common u for which all these valuations increase by at most δ. This is why the pencil was required not merely to have suitable slopes but also to have the Frobenii of all nonzero slopes trivial on the corresponding row modulo ℓ^M. Along the chosen line the extra cost of dropping the other ties can then be bounded by Lemma 20.7: these are determinants of $uR_i(\varphi_P) - 1$ and their conjugate counterparts. These are precisely among the polynomials selected for the small-valuation conditions.

For each value family one may instead take the line with only the helper coordinate equal to u and the others equal to 1. Its initial valuation costs at most $8c$ by (20.14), using φ_{P_0}

close to z and the bound (20.15). On inserting the other primes successively, their Frobenius values are the helper value by (20.7). By the monic Euler-factor avoidance in the first stage the remaining cost is at most 8δ per insertion. This bounds the value family on this new polynomial point. There is no need to find one tuple lying on the lines constructed for all different rows and value families.

At the second stage use the sharp Fitting polynomials for the unrestricted modules, now evaluated at the good row-specific points just constructed; also include the polynomials for all the value families. The number of variables is at most J_0, and by (20.11) all these specializing points satisfy the near-1 power condition with one common exponent. Lemma 20.13 therefore gives one tuple of ℓ^N-th roots of unity working for all polynomials, and simultaneously for their evaluations at the first m powers of that tuple. The monic factors $Y^j - 1$ ensure nontriviality and bound the valuations at the tame generators. The total error from these selection stages and from inserting the Euler factors is bounded by $(8J_0 + 2)\delta < \eta$. Formula (20.9) gives the character ψ with common conductor, and the finite coefficient rings can be enlarged without changing the normalized length bounds. This describes the proof's organization and supplies the polynomial and length arguments, while the exact global compatibility of the pencil and integral special-value measures is retained in the stated formal-source input.

Corollary 20.15 (Formal-source consequence: twisting all rows)**.** *For a finite collection of rows T_i of ranks r_i of admissible data, fixed ℓ, ι and any depth M, there is an m-fold twist θ such that for all $1 \le j \le m$, writing $\vartheta = \theta^j$ and ψ for its inverse avatar, one has*

$$\sum_{P \in Q} v_\ell(\psi(\tau_P) - 1) \le c_\ell, \tag{20.16}$$

$$\lambda\big(\mathfrak{S}_M(\psi T_i; S \cup Q \cup Q^c, \mathcal{L}) \cap \ker \mathrm{ev}_{\gamma_i}\big) \le (10 f r_i + 1) c_\ell, \tag{20.17}$$

$$\left|\iota(\rho_p(\vartheta, s)^{-1})\right|_\ell \le \exp(9 c_\ell \log \ell) \quad \textit{for every ramified } p, s. \tag{20.18}$$

The normalizer satisfies the scalar and unit conditions of SelmerLengthLe*.*

Deduction from the stated inputs. Take a chosen place w above one of g_0, g_1 different from ℓ. Proposition 20.5 applies for all rows after one common finite enlargement of the coefficient ring. Choose $B_i = 10 f r_i c_\ell$ and $\eta = c = c_\ell > 0$ in Proposition 20.14, using (20.15) for all the finite value families. The resulting character is the inverse avatar of a tame character, by the class-field identification `exists_tameChar_invAvatar_eq` of `New/Selmer/ValueFamilies.lean`. The nontriviality at the tame generators for all powers makes it an m-fold twist, while being trivial at the original places leaves the normalizers and crystalline conditions unchanged. The bounds on the normalized Selmer length and local ramification give the first two assertions. Using (20.14) and the ramified-pair norm relations in that file gives $v_\ell(\iota\rho_p(\vartheta, s)) \le 9c_\ell$, equivalent to (20.18). This is exactly `exists_mFoldTwist_selmerLengthLe` in `New/Selmer/TwistForRows.lean`. □

20.9 The Tate characters and triangular Taylor–Wiles places

The bound on the distinguished row block obtained above is not a bound on all congruences of the normalized column. The extra characters $\xi_e = \epsilon_\ell^{-e}$ also give torsion cocycles. It is not possible simply to replace their modules by a conjectural bound on the even cyclotomic class groups. The argument uses instead a finite collection of auxiliary primes that depends on the torsion depth but whose cardinality is bounded independently of that depth. These primes are distinct from the conductor of θ in the twisting theorem.

For a Tate character ξ_e, define $\mathfrak{E}_{\nu,e}(S)$ to consist of the locally constant cocycles z in $(\mathcal{O}/\ell^\nu)(-e)$ unramified outside the places above ℓ and $S \cup S^c$, with $z(g^c) = z(g)$ and vanishing on the prescribed subgroup of the local decomposition groups in the kernel of the cyclotomic character. For $e > 0$ the relevant odd module $\mathfrak{H}^-_{\nu,e}$ is given by the same unramified and local conditions without the added primes and with the odd symmetry condition. At $e = 0$ the odd-class package of the source uses a different convention, namely the unramified cocycles with its local condition without the odd restriction. These are the definitions `evenClasses` in `New/Selmer/Corepresentable.lean` and `shaOdd` in the odd-class bound. We do not implicitly decompose into two direct eigenspaces at $\ell = 2$; the use of sums and differences, and factors of 2, is important at that prime.

Let N_E be the cyclotomic conductor specified for E in the datum, so $E \subset \mathbb{Q}(\zeta_{N_E})$. The following background result is one of the few global class-number terms that survive in the final bound.

Proposition 20.16 (One integer for all odd classes (Theorem 22.14))**.** *There exists an integer $X > 0$ with*

$$X \le 2^{4(N_E+t+3)^{27}} h_E, \tag{20.19}$$

where h_E is the class number of E, such that for every ℓ, every $0 \le e < t$ and every ν, the normalized length of $\mathfrak{H}^-_{\nu,e}$ is at most $v_\ell(X)$. The coefficient rings are assumed, in addition to the integral DVR hypotheses, to contain enough roots of unity of order the exponent of $(\mathbb{Z}/N_E)^\times$.

This is `exists_length_shaOdd_le` in `Lib/Cyclotomic/OddSha/OddShaBound.lean`. It is a finite-level statement uniform in ν, not a generic assertion of Iwasawa nonvanishing. The proof uses the cyclotomic embedding of E, annihilating elements of Stickelberger type, the class-number bound in the trivial twist and explicit estimates of their sizes; it constructs the integer X by multiplying the corresponding annihilator factors. We retain this input as stated, since the bound (20.19), not just finiteness for every fixed (ℓ, e, ν), is needed for the global sum.

In contrast, the size of $\mathfrak{E}_{\nu,e}$ will not be included. To see the substitute, let $a \in G_E$ separate the distinct Tate characters up to ℓ-adic loss. A triangular system of depth ν consists of a finite ordered set J', a nondecreasing type function $\mathrm{ty} : J' \to \{1, \dots, t' - 1\}$, and distinct Taylor–Wiles places w_j with tame generators τ_j and Frobenius elements ϕ_j fixing the ℓ^ν-th roots of unity. The places have prime norm and are conjugate-disjoint. Given bounds H_ℓ and C_ℓ, the `IsPrimeTriangle` predicate includes the following conditions:

1. There is a finite detecting set s of group elements such that evaluation at s detects all unramified-outside-ℓ Tate cocycles modulo ℓ^ν in the indicated range, and $(|J'| + |s|)H_\ell \le C_\ell$.

2. For each $0 \le e < t'$, the product $\prod_{k \ne e}(\xi_e(a) - \xi_k(a))$ divides ℓ^{H_ℓ} in the integral ring.

3. For every $z \in \mathfrak{E}_{\nu,\mathrm{ty}(j)}(\{w_i\})$, if $z(\tau_i) = 0$ for $i < j$ and $z(\phi_i) = 0$ for every earlier i of the same type, then $\ell^{H_\ell} z(\tau_j) = 0$.

4. Write $d_e(g) = \xi_e(g) - 1$. For $e > 0$, $4d_e(a)^2 \mid \ell^{H_\ell}$, and if an even cocycle z with no new ramification is zero at the ϕ_j of type e, then

$$d_e(a)(2z(g)) = d_e(g)(2z(a)) \qquad (g \in G_E).$$

   Thus a controlled multiple of such a cocycle is a coboundary.

The precise predicate is in `New/Selmer/Corepresentable.lean`; writing a as in that predicate accounts for the inverse in the expression of the Tate characters.

Proposition 20.17 (Formal-source input: existence of triangular places)**.** *Suppose $4 \le f$, $N_E + t' < y$, $\epsilon_\ell(\gamma)^a - 1 \mid \ell^h$ for $0 < a < t'$, and $\ell \le y^2$ or $h = 0$. Assume the desired Frobenius conditions U hold on the generic Frobenius data near some element of the form $\delta^{\ell^\nu} g$ for each ν, g, where $\epsilon_\ell(\delta) = 1$. Then there exist an ordered index set J' and a type function, both independent of ν, such that at every depth $\nu > 0$ one can choose Taylor–Wiles places satisfying the triangle conditions with $a = \gamma^{-1}$, $H_\ell = yh$ and $C_\ell = y^{22}h$, and with all their Frobenius data satisfying U.*

This is `exists_isPrimeTriangle` in `New/Selmer/TrianglePrimes.lean`. The source constructs the detecting set uniformly and indexes the triangle by the product of its detecting set and the nontrivial types. The unusual order on this product, increasing by type, is necessary in the subsequent polynomial-support argument. One ingredient is a reciprocity functional z_0 on G_E, locally constant modulo ℓ^ν, which on the kernel of the cyclotomic character modulo ℓ^ν is an odd additive map transforming by cyclotomic weight $e+1$. At split places outside a finite exceptional set, if the functional is zero at all other Frobenius elements in a finite collection, its value on the remaining Frobenius, multiplied by 2, annihilates the corresponding even inertia value, by reciprocity. The source proves existence of such a functional, vanishing at a prescribed set of no more points than were previously used, and taking a unit value on some element of the cyclotomic kernel: `exists_isLaw` in `New/Selmer/ReciprocityLaw.lean`. The even cocycles are tested through all cyclic additive quotients, which detect zero by duality for finite abelian groups. By taking a linear combination of elements of the cyclotomic kernel and applying the products of differences of the cyclotomic weights, the construction makes the new functional vanish on all previously chosen Frobenius elements of other types. It then applies Chebotarev with the generic Frobenius conditions simultaneously. The annihilating error in this selection is the controlled product of differences above, not the length of the even module. This is the mechanism of the induction in Proposition 20.17; the integral construction of the reciprocity functionals and uniform detecting points is retained as part of the formal-source input.

For our datum the elementary separation of the characters is available with $h \le v_\ell(M_g)$ and $y = h_2 = r^8$. Indeed for ℓ larger than h_2^2 the field conductor is prime to ℓ and an inertial cyclotomic value whose reduction generates $\mathbb{F}_\ell^\times$ separates all exponents less than t by units; take $h = 0$. For the smaller primes, choose a degree-one place above the chosen prime $g_i \ne \ell$. Its Frobenius has cyclotomic value g_i, and each $g_i^d - 1$ with $0 < d < t$ divides M_g by (20.4); take $h = v_\ell(M_g)$. This is `exists_cycChar_pow_sub_one_dvd` in `New/Selmer/DenominatorBound.lean`. The admissible size conditions give $N_E + t < h_2$, and the field contains sufficiently many infinite places. The generic Frobenius conditions needed in Proposition 20.17 arise from the Hecke string lemma discussed below. In particular the global odd term in (20.19) and the bounded valuations of the separating products are sufficient even if the even-class modules are very large.

20.10 The first-order Ribet construction

The algebraic argument works with a finite set Y of normalized column representations U_y over $\mathcal{O}$. A subset $\mathfrak{X} \subseteq Y$ consists of the distinguished lifts; for $y \in \mathfrak{X}$ the representation has block form $R_y \oplus \bigoplus_e \xi_e$ by (20.3), where R_y denotes the already twisted row block. For the projection we distinguish one index $y_0 \in \mathfrak{X}$. Use the group algebra $\mathcal{O}[G_E]$, and let a separate

its Tate characters. Choose a normalizer γ with $R_{y_0}(\gamma) = sI$, $s-1$ a unit, and $\xi_e(\gamma) = 1$. Put

$$\mathcal{N} = [\gamma a] - s[a] + p_0 D \sum_g n_g[g], \qquad \beta_y = \det\left(\sum_g \mathcal{N}_g(U_y(g) - 1)\right), \tag{20.20}$$

where p_0 is the row-cocycle modulus, D a separating factor, and the n_g are finitely supported nonnegative integers. This is `ribetNormaliser` in `New/Selmer/RibetCramer.lean`. For distinguished systems $\beta_y = 0$ since the block $\xi_0 = 1$ is annihilated by every $U_y(g) - 1$. On other systems one can arrange $\beta_y \neq 0$ whenever some $U_y(g) - 1$ is invertible by the following finite avoidance observation: each determinant is a polynomial in the finitely many n_g, nonzero as a polynomial if at least one leading matrix $U_y(g) - 1$ is invertible. Finitely many nonzero polynomials over a characteristic-zero field do not vanish on all tuples of nonnegative integers (induct on the number of variables). This is the reason the normalizer is allowed the final summand. One must still treat the systems on which every $U_y(g)$ has eigenvalue 1, rather than assuming those systems are precisely the lifts.

Let $A = \mathcal{O}[X_j : j \in J']$. We use a subalgebra $\mathfrak{S}$ of functions from Y to polynomials over the coefficient field, containing all scalar variables and characteristic polynomial coefficients of the $U_y(g)$. Its trace ideal $\mathfrak{J}$ contains, for all g, $p \operatorname{tr} U_y(g)$, with $p = \ell^\nu$, and the trace ties (for the two copies of j)

$$\operatorname{tr}((U_y(\tau_j) - 1)U_y(g)) - X_j \operatorname{tr}((U_y(\phi_j) - 1)e_j(U_y(\phi_j))U_y(g)), \tag{20.21}$$

where e_j is a polynomial with integral coefficients and $e_j(1) = 1$; conjugate inverse generators are included with the same variable. The precise abstract predicate `IsTraceIdeal` is in `New/Selmer/RibetFirstOrder.lean`. For that file the two conjugate copies may be prescribed independently but share X_j.

The Ribet element assertion is stronger than a mere membership of a determinant in an annihilator. It asserts the existence of $F \in \mathfrak{S}$ with

$$F\beta \in \mathfrak{J}, \qquad \beta_y = 0 \ (y \in \mathfrak{X}), \qquad [X^d]F_y = 0 \ (y \in \mathfrak{X},\ d < \mathbf{1}), \qquad [X^{\mathbf{1}}]F_{y_0} = c_0, \tag{20.22}$$

where $X^{\mathbf{1}} = \prod_{j \in J'} X_j$ and $d < \mathbf{1}$ means a proper coordinatewise divisor of this monomial. In addition $c_0 \neq 0$ has a quantitative upper valuation bound, given below. This is `IsRibetElement`. Crucially the vanishing condition is not that all terms of smaller total degree vanish; the source works with the lower order ideal of exponents coordinatewise at most $\mathbf{1}$. It is precisely this vanishing condition that permits the later first-order coefficient extraction at several primes.

Here are some details of the algebra of producing such an element. Because the distinguished block has full matrix span, one can construct group-algebra expressions isolating the Tate blocks. For each $y \in \mathfrak{X}$, write $R_y(\gamma_y) = s_y I$, $s_y - 1$ a unit, and $\xi_e(\gamma_y) = 1$. The element

$$k = \prod_{y \in \mathfrak{X}} (s_y - [\gamma_y])/(s_y - 1)$$

is integral, acts by zero on every distinguished row block, and by one on all Tate blocks. The factors need not commute: each row representation sends its own factor to zero, which annihilates the entire product in that representation. For each index e, Lagrange interpolation supplies an element q_e with $\xi_i(q_e) = D\delta_{ie}$, where D is a common multiple of the products of differences. Explicitly take a multiple of $\prod_{i \neq e}([a] - \xi_i(a))$ by integral coefficients. Both constructions remain valid in characteristic ℓ after retaining the multiplier D, whereas dividing by the differences would not. These are the two block-separation operations at the start of `ribet_first_order`.

To explain the Fitting module in this proof, let $\mathfrak{M}$ be the finite $\mathcal{O}$-module spanned, inside the family of matrices, by the elements $(U_y(g)-1)_{y\in Y}$. It is finitely generated since Y is finite. A linear map from $\mathfrak{M}$ into a torsion row module produces a cocycle when it respects left multiplication by the $U_y(g)$, because

$$U_y(gh)-1=(U_y(g)-1)+U_y(g)(U_y(h)-1).$$

Conversely a cocycle obeying all linear relations between the family $(U_y(g)-1)_y$ can be evaluated by a linear map on $\mathfrak{M}$ with the same equivariance. This motivates `obeyingCocycles` in `New/Selmer/BlocksAndCocycles.lean`. To compare with local crystalline classes it is useful to extend this graph map slightly: if the target is killed by a, Artin–Rees gives q with $\mathfrak{M}\cap a^q\mathfrak{V}\subseteq a\mathfrak{M}$, for the ambient matrix lattice $\mathfrak{V}$. Its evaluation therefore extends by zero on $a^q\mathfrak{V}$ to $\mathfrak{M}+a^q\mathfrak{V}$, and its graph is a subrepresentation with injective first projection. The map is well defined because differences in the intersection are killed. This is the lemma `obeyingCocycles.exists_graph_subrepresentation`; it also shows why the extracted cocycles are continuous and have the necessary unramifiedness, even though the definition by linear relations alone did not initially impose continuity. Where all U_y act trivially on inertia, the corresponding matrix differences are zero, so the cocycle is zero there. The graph description and continuity of the integral matrices show that a finite reduction of them determines the cocycle.

After imposing the trace ties and the relations of (20.20), one works in the polynomial extension of this finite presentation. At each distinguished block, a bound on the torsion length supplies a determinant (a maximal minor of a relation matrix) with small valuation. For the Tate block one must use a first-order variant: the relations between evaluation at a pair of conjugate inertia elements and evaluation at the Frobenius elements, multiplied by X_j, make the corresponding determinant divisible by the products of X_j of that type. To state its precise arithmetic input, for a module Z of Tate cocycles put $I_e=\{j:\mathrm{ty}(j)=e\}$. The triangle cost with multiplier D and allowance b_e consists of a set L of evaluations such that, for cocycles zero on L, agreement at earlier pairs of tame generators and at earlier Frobenius pairs of the indicated type forces

$$D\big(z(\tau_j)-z((\tau_j^c)^{-1})\big)=0\qquad(j\in I_e).$$

For submodules of these cocycles vanishing at all tame generators and agreeing at the Frobenius pairs of type e, their length modulo D-torsion plus the length cost of the initial evaluations on L is at most b_e. This is `HasTriangleCost`.

Proposition 20.18 (Formal-source input: first-order Ribet bound)**.** *Use the preceding block setup and trace ideal. Suppose the normalized row cocycle module, modulo p_0, obeying the linear matrix-difference relations and vanishing at the new tame generators and its normalizer, has ordinary length at most $b_0<v_\pi(p_0)$. Suppose the corresponding Tate modules, modulo p, have triangle costs b_e and detecting set s as above. Include auxiliary inertial relations of the form $(U_y(\delta_v)-1)o_v(U_y(\delta_v))=0$, where $\xi_e(\delta_v)=1$. Assume also $e_j(\xi_e(\phi_j))\equiv 1 \pmod p$ and $\prod_{k\neq e}(\xi_e(a)-\xi_k(a))\mid D$. Then the Ribet element (20.22) exists with*

$$v_\pi(c_0)\le b_0+\sum_e b_e+t'\left(\sum_v v_\pi(o_v(1))+2|s|v_\pi(D)\right)+(|J'|+1)v_\pi(D).\tag{20.23}$$

Here the Tate-module assertions of triangle cost and detection are required for the modules whose members vanish on all δ_v and whose coboundaries are killed by $\xi_e(a)-1$ (as in the source).

This is `ribet_first_order` in `New/Selmer/RibetFirstOrder.lean`. The block Fitting lemma `exists_mem_fittingIdeal_block` in the same file gives, for one Tate block, a relation in the Fitting ideal of the form $G_e \prod_{j\in I_e} X_j$ with constant-term bound $v_\pi(G_e(0)) \le b_e + (|I_e| + |s|)v_\pi(D)$. The proof successively enlarges presentations, uses the stated cost to clear one evaluation at a time, and then multiplies the resulting minors. When combining blocks, terms not supported in the interval of divisors of $\prod X_j$ can be discarded. More precisely, work modulo the ideal spanned by monomials not dividing $\prod X_j$, and at each step additionally keep track of previous types; the type ordering puts the unwanted earlier variables in the annihilator after extracting the earlier block factors. The direct-sum decomposition of the distinguished representations turns the determinant of a block-flattened presentation matrix into the product of its row determinant and its Tate determinants. The factors D and $o_v(1)$ are inserted to make the block operations integral, accounting for the additional losses in (20.23). Finally a group-algebra Cramer determinant identity turns these minors into an element of the trace ideal after multiplying by β. That identity is `exists_ribet_element_cramer` in `New/Selmer/RibetCramer.lean`; it uses integrality of characteristic polynomials of block matrices over the group algebra, supplied by the indicated Amitsur formula. These Fitting, block Cramer and multivariable support steps are specialized assertions of the formal-source proof, not consequences of the elementary determinant inequality alone.

In applying the triangle cost to a Tate component, choose L to be the prescribed local set above ℓ. The length lost by imposing vanishing there is bounded by the crystalline defect, since crystalline Tate cocycles vanish on that set. After imposing vanishing there, taking $z + z^c$ produces an even cocycle, and the triangle condition gives the required agreement of the conjugate pairs up to the factor D. Once all new inertia values vanish the cocycle is unramified outside ℓ. Its odd part $z^c - z$ now provides an element of the odd module, and the even part, by the triangle condition, is a coboundary up to a factor of controlled valuation. The assumption on coboundaries makes the kernel of this odd projection killed by D; in particular its length modulo D-torsion is bounded by the length of the odd module. The assertion is `hasTriangleCost_tateTwist` in `New/Selmer/TateTwistSteps.lean`; it treats $e = 0$ directly. Thus b_e is the local defect plus the length in Proposition 20.16, without any even-class length term.

20.11 Hecke operators realizing the first-order relations

The formal-source proof now passes from the group algebra to forms. For a fixed auxiliary pair, label tuple and twist, choose the base space at the corresponding level and consider its enlargements at the Taylor–Wiles places. On each of these finite-dimensional spaces there is a finite set of eigensystems, and a closed integral Hecke algebra of scalar functions of the eigensystem preserving the lattice of Weyl-integral forms. The characteristic polynomial coefficients of the normalized Galois representations, as functions of the eigensystem, belong to this algebra: the attached representation comparison provides them as limits of integral Hecke coefficients on generic split Frobenius data. In particular $\mathrm{tr}(P(U_y(h))U_y(g))$ belongs to the algebra for every integral polynomial P, by expanding P and using the trace at each $h^i g$. This is `BaseLevel.charpoly_mem_heckeAlgebra` and `BaseLevel.traceElem` in `New/Patching/OnePrimeTie.lean`, with the closedness and local–global compatibility of the Hecke algebra taken as formal-source inputs.

A system is called weakly lifted for this step if every $U_y(g)$ has eigenvalue one; call this set $\mathfrak{W}_1$. It is potentially larger than $\mathfrak{X}$. A normalizer can be chosen as in (20.20) so that $\beta_y \neq 0$ off $\mathfrak{W}_1$, but the argument does not try to separate $\mathfrak{W}_1 \setminus \mathfrak{X}$ by this determinant. A further

polynomial $P(\beta)$ over the ℓ-adic integers is chosen to vanish at all the nonzero determinant values on the old eigensystems and to have nonzero constant term (take the product of their linear factors, with integral scalar factors as necessary). Thus $P(\beta) \neq 0$ on $\mathfrak{W}_1$ and vanishes on all other old systems. It is not asserted that the coefficients of P or the divisions needed in clearing this cutoff are units.

Proposition 20.19 (Formal-source input: Hecke string)**.** *For a weakly lifted eigensystem of an auxiliary pair, its rescaled unramified Hecke polynomial at a generic split place of prime norm q_w has the full string of roots $1, q_w, \dots, q_w^{t'-1}$. More precisely, if its normalized Galois characteristic polynomial at a generic Frobenius is congruent coefficientwise, modulo the maximal ideal, to*

$$S(X) \prod_{i<t'} (X - \epsilon_\ell(\phi)^{-i}) \quad \text{with } |S(1)|_\ell \geq 1,$$

then its rescaled Hecke polynomial factors as the corresponding full string times a residual factor dividing an integral polynomial $e(X)$ with $e(1) = 1$.

This is `isRoot_hecke_string` and `eventually_heckePoly_eq_string_mul` in `New/Selmer/HeckeString.lean`. The underlying representation-theoretic input is the string lemma for spherical unitary representations and the prescribed local type. The integral polynomial assertion follows from comparing monic string factors modulo the maximal ideal and from the bound $|S(1)|_\ell \geq 1$: for a monic B dividing the string, the Gauss norm of a product with B equals the product of the norms; hence comparison of AB and $A'B$ forces comparison of A, A'. Once the residual factor A is integral and $A(1)$ a unit, use $e = A/A(1)$. This is the elementary part of `exists_integral_multiple_eval_one`. For finitely many relevant eigensystems one takes the product of the resulting polynomials, still with $e(1) = 1$.

The required Frobenius conditions in Proposition 20.17 ensure that this proposition can be applied at every chosen auxiliary prime simultaneously. A simple reason such conditions are available is that at the distinguished row block a scalar inertial element with long prime-to-ℓ order allows one to move a Frobenius condition without changing its cyclotomic value. If an integral matrix A has rank r', the polynomial in Y given by $\det(1 - YA)$ has degree at most r' and constant term one. Among $r' + 1$ distinct residual scalars Y, at least one avoids its roots. The scalar-inertia condition of the row provides such scalars even after taking large ℓ-powers of the inertial element. By continuity the nonvanishing is a neighborhood condition. The formal statement is `exists_eventually_no_eigenvalue_one`. The attached representations and the generic Frobenius filter then yield the simultaneous string conditions for the eigensystems congenial to the target residual system. This is the condition imposed in the choice of the triangular primes in `denominatorLe_rowBudget`.

To realize the relations (20.21), at a Taylor–Wiles place let $\Gamma = \mathrm{GL}_{r'+t'}(k_w)$ and $H \subseteq \Gamma$ the stabilizer of a line. The character of H to $\mathbb{Z}/\ell^\nu$ given by the logarithm of its scalar action on the line is a diamond-operator character. Its transfer to Γ factors through the determinant. The relevant normalization, after multiplying by $(r' + t')!$, sends a chosen primitive determinant scalar to one modulo the logarithmic modulus. This follows from the transfer identity on scalar matrices and the count of lines; it is `exists_line_stabilizer_log` in `New/Patching/OnePrime.lean`. The proof then constructs a first-order rather than a semisimple deformation of the trace. At one place, if ρ denotes the residue-group action, define for a function f_0 on the logarithmic group

$$W(f_0) = \sum_{h \in H} f_0(\log h)\rho(h).$$

Changing h to $g^{-1}h$ or shifting the logarithmic parameter shows that multiplying by a diamond element shifts the argument of f_0. Taking a representative of the logarithm gives

the infinitesimal version of this identity modulo ℓ^ν. The source performs the induced-representation transfer, using coset representatives for Γ/H, to construct a local operator F_j and the norm-sum operator $N_j = \sum_{\gamma\in\Gamma} \rho(\gamma)$, with the trace relations acting as follows: if $t_j(g) = \mathrm{tr}((U(\tau_j) - 1)U(g))$ and $r_j(g) = \mathrm{tr}((U(\phi_j) - 1)e_j(U(\phi_j))U(g))$, then, modulo ℓ^ν and up to the specified integer factors, $F_j t_j(g) = N_j r_j(g)$ and $N_j t_j(g) = 0$. There are additional identities controlling how F_j moves the old-space invariants. This is the formal-source input `BaseLevel.exists_tie` in `New/Patching/OnePrimeTie.lean`, proved using `one_prime_first_order` in `New/Patching/OnePrime.lean` and local line-stabilizer operators. The nonzero factorial factor is retained rather than divided out at the torsion level.

The multiple-prime coefficient extraction is motivated by the following simple polynomial observation. If operators F_j, N_j at distinct places commute, multiplying the first-order relation by products of operators at the other places transports the contribution from a monomial X^d to the monomial X^{d+e_j}. Terms with exponents larger than one are discarded. The product of the sums N_j is the full residue-group norm mapping to the old-level subspace, while the complementary product of the F_j probes its infinitesimal deformations. Thus evaluating the trace-ideal relations on this first-order coefficient system annihilates the trace ideal modulo ℓ^ν. On the distinguished systems, all coefficients of F in (20.22) properly dividing $\prod X_j$ vanish; only the top coefficient remains. The source proves that this coefficient gives an integral diagonal operator at old level, localized to the cutoff determined by $P(\beta)$, after controlling the movement from the old space. The auxiliary depth ν must exceed the valuation of an explicit factor involving the constant term of P, the level automorphism factors and the factorial. It can be made arbitrarily large since the index set of the triangle was bounded before ν was chosen.

Precisely the result needed is `BaseLevel.exists_diagOp_mem_ellIntegralForms` in `New/Patching/OnePrimeTie.lean`, a formal-source input. Given the string condition near the distinguished residual eigensystem, enough divisibility by ℓ^ν, and a Ribet element of top-coefficient valuation with a factor M_0 at most Γ_0, it supplies a scalar function s of the eigensystem with $v_\ell(M_0 s(y_0)) \le \Gamma_0$. The diagonal operator with scalars s sends a vector supported on $\mathfrak{X}$ to an integral vector up to a specified level factor and any original denominator, provided that vector agrees with an integral old form on the eigensystems where $P(\beta) \neq 0$ in the sense of the diagonal-operator test in the source statement. Thus the integral output does not require that all weakly lifted systems themselves belong to $\mathfrak{X}$. A tested form minus its projection to Z is orthogonal to the unwanted weakly lifted components by the Hecke string assertion and (20.1), after inserting the chosen label projector. It follows that the diagonal operators satisfy the required integrality hypothesis on the projected tested forms.

20.12 The denominator bound from the local length budgets

We now combine the congruence terms numerically. Write b_1 for the rounded value of the thick weight and let

$$F_\ell^{\mathrm{num}} = (2b_1)! \left(\prod_{\max(b_1-30h_2^3,0)<i\le b_1+30h_2^3} i \right)^{h_2^3}.$$

This integer is independent of ℓ; the notation emphasizes its use in the local estimate. With $g = g_0 g_1$, define

$$\mathrm{loss}_\ell = 10^6 \frac{m}{k} \left(b_1 v_\ell(g) + v_\ell(F_\ell^{\mathrm{num}})\right), \qquad D_\ell = h_2^{27}\left(c_\ell + v_\ell(X)\right) + 4f\,\mathrm{loss}_\ell .$$

Here X is from Proposition 20.16. These are `loss`, `lambdaLoss` in `New/Selmer/DeterminantBound.lean`, and `rowBudget` in `New/Selmer/RibetElement.lean`. The detailed local comparison of torsion cocycles and the crystalline relations is developed in the next part of the proof. Its crucial formal-source bound, `sum_blockLoss_le` in `New/Selmer/RibetElement.lean`, is

$$\sum_{i \text{ blocks}} \lambda(\text{loss of crystalline relations on block } i) \le 4r' f \operatorname{loss}_\ell, \tag{20.24}$$

for arbitrary torsion depths of the individual blocks. This controls the failure of extracted cocycles to satisfy the crystalline local conditions in the Selmer theorem. One cannot apply that theorem directly to all cocycles extracted from congruent matrix differences.

Proposition 20.20 (Formal-source input: arithmetic Ribet element)**.** *Use the apt columns and rows of an auxiliary pair and a twist satisfying* (20.16) *and* (20.17) *at a depth $M > \lfloor rD_\ell \rfloor$. Suppose their normalized matrices satisfy* (20.3) *on the distinguished systems, and for new twist inertia the normalized columns satisfy $(U(g) - \psi(g))(U(g) - 1) = 0$. At a triangular auxiliary level, with the trace ideal as above, the determinant from* (20.20) *has a Ribet element with top coefficient c_0 such that*

$$v_\ell(M_g^3 c_0) \le r' D_\ell.$$

This is `ribet_columns` in `New/Selmer/RibetElement.lean`. The required quadratic inertia relation is `normalisedRep_quadratic_on_inertia` in `New/Selmer/LabelledColumns.lean`; it follows from the shape of the parahoric type at the twist conductor. The cocycle length for the row block is at most its normalized Selmer bound $(10fr' + 1)c_\ell$ plus its share in (20.24). The Tate costs in (20.23) contribute at most $v_\ell(X)$ per block plus their shares in (20.24). The scalar-separation and detecting terms contribute multiples of $h_2 h$ and $h_2^{22} h$; the old and new tame polynomial factors have valuations bounded by the number of places times c_ℓ and by the sum in (20.16). The aggregate of these noncrystalline terms, including the factor M_g^3, fits inside $r' h_2^{27}(c_\ell + v_\ell(X))$ by the admissible size inequalities. The source proves the explicit inequality in the last step of `ribet_columns`. Notice why the row modulus M is chosen larger than rD_ℓ: it is larger than the resulting bound on the row-cocycle length, as required in the first-order Ribet lemma.

We give the final linear algebra of combining individual integral operators, since it avoids multiplying the denominator cost by the number of distinct eigensystems. Let $\{p_i\}_{i=1}^q$ be a basis of Z consisting of nonzero vectors from the spanning family of lifted forms. Given integral test vectors u_a, v_b in the denominator condition, write $P_Z v_b = \sum_i P_{ib} p_i$. Suppose for each j a diagonal comparison operator acts on p_i by scalars a_{ji}, after a common injective comparison map, with $|M_0 \iota(a_{jj})|_\ell \ge e^{-\Gamma}$, and that the matrices

$$(\iota\langle u_a, \textstyle\sum_i a_{ji} P_{ib} p_i\rangle M_0)_{a,b}$$

are integral for every j. Choose coefficients t_j in the integral valuation ring of $\overline{\mathbb{Q}}_\ell$ so that for every i,

$$|M_0 \sum_j t_j \iota(a_{ji})|_\ell \ge |M_0 \iota(a_{ii})|_\ell \ge e^{-\Gamma}.$$

This is possible since for each i the finite linear polynomial has its Gauss norm attained at some tuple of units, and a finite set of nonzero linear polynomials over the infinite residue field can be avoided simultaneously after dividing each one by a coefficient of maximal absolute value. The linear combination of the integral pairing matrices is integral. Its determinant equals the determinant of the matrix $(\iota\langle u_a, P_Z v_b\rangle)$ multiplied by $\prod_i (M_0 \sum_j t_j \iota(a_{ji}))$, by writing all pairing matrices as a product of the basis-change matrices and the relevant diagonal matrices. Consequently the absolute value of the desired determinant is at most $e^{q\Gamma}$. This is the linear algebra in `norm_gram_det_le_of_operators` in `New/Selmer/DeterminantBound.lean`.

Apply the arithmetic Ribet element proposition and the first-order trace deformation over the triangle to a single basis eigensystem in the space of lifts. In the setting of the tested forms, the local operator integrality described above, together with the class-sum pairing and automorphism bounds absorbed into M_g^3, provides exactly the operators for this linear algebra with $\Gamma = r' D_\ell \log \ell$. This yields

$$\mathrm{DenominatorLe}(\iota, Z, D_\ell \log \ell), \tag{20.25}$$

that is (20.2) with $\beta = D_\ell \log \ell$. The precise formal-source theorem, including the spanning hypothesis that each generator is apt, lifts a fixed system having a row for all twists, and has the needed Selmer bound, is `denominatorLe_rowBudget` in `New/Selmer/DenominatorBoun d.lean`. The use of tested forms, not arbitrary forms, is why the weakly lifted eigensystems suffice. Equation (20.25) marks the end of the congruence argument.

20.13 Global numerical summation and proof of the fit theorem

All the terms in the local allowances are nonnegative, and the factorization of each positive integer A gives $\sum_{\ell \in T} v_\ell(A) \log \ell \le \log A$ for every finite set of primes T. The other part of c_ℓ has bounded sum, since $\sum_{n \ge 2} 2^{-n} \log n < \infty$. Therefore

$$\begin{aligned}
&\sum_{\ell \in T} 9 c_\ell \log \ell \le 9 (\log M_g + \sum_{n \ge 2} 2^{-n} \log n), \\
&\sum_{\ell \in T} D_\ell \log \ell \le h_2^{27} \left(\log M_g + \log X + \sum_{n \ge 2} 2^{-n} \log n \right) \\
&\qquad\qquad + 4 \cdot 10^6 \frac{fm}{k} \left(b_1 \log g + \log(2b_1)! + h_2^3 \sum_{\max(b_1 - 30h_2^3, 0) < i \le b_1 + 30h_2^3} \log i \right).
\end{aligned}$$

These estimates reduce the infinite collection of local lengths to a handful of ordinary integer sizes. For example $\log(2b_1)! \le 2b_1 \log(2b_1)$, and the short-interval contribution is at most $60 h_2^6 \log(b_1 + 30h_2^3)$. Also $\log M_g$ is bounded by a constant times $f(rt)^2 \log g$ by the order formula for finite general linear groups, and $\log X$ is bounded by $4(N_E + t + 3)^{27} \log 2 + \log h_E$. The explicit admissibility inequalities in the source bound these contributions by

$$\sum_{\ell \in T} 9 c_\ell \log \ell \le x^{999}, \qquad \sum_{\ell \in T} D_\ell \log \ell \le B_0. \tag{20.26}$$

See the numerical calculation in `exists_isFitFamily` in `New/Selmer/FitTwistsExist.lean`, using the class-number estimate and `Data.le_of_hasLogBound`. These are formal-source bounds with the full numerical hypotheses on the datum. To see their scale, along the admissible family with fixed $\varepsilon > 0$, b exceeds every fixed power of x; h_2, N_E, r and t are polynomial in x. Thus the first term of the second bound is smaller than B_0 by a power of b. In the remaining factorial term b_1 has order b, $\log g$ and $\log b_1$ have order $\log b$, and the leading contribution is $O(fmb \log b/k)$. Its ratio to $B_0 = fa(\log b)^3/(mk)$ is $O(m^2 b/(a(\log b)^2)) = O((\log b)^{-2})$ by $m^2 = O(x)$ and $a \ge xb$. This highlights the logarithmic slack without replacing the exact inequalities by asymptotic claims.

Proof of Theorem 20.2, using the stated formal-source inputs. For each ℓ and ι, make the finite list of rows in Proposition 20.3, including all rank-one pairs, before choosing any twists. Apply Corollary 20.15 to this list at depth $M = \lfloor r D_\ell \rfloor + 1$. The resulting m-fold twist θ_ι has conductor prime to ℓ and satisfies the local conditions for its first m powers. At each arrow j,

the good lift is spanned by the apt systems with the chosen labels and this common Selmer bound. For the rank-one lines the corresponding assertion follows from their apt label $X - 1$ and the rank-one row. Consequently (20.25) applies simultaneously to the desired spaces, with the same allowance $\beta_{j,\ell} = D_\ell \log \ell$ independent even of j. The second inequality of (20.26) is its required finite-sum bound. The special-value inequality (20.18) and the first inequality of (20.26) give the required height bound on the inverse ratios: in the height definition take the constant archimedean factor 1 and the local log allowance $9c_\ell \log \ell$. All these choices are independent between the pairs (ℓ, ι) but uniform over the finite lists for a given pair, as required. Thus the family satisfies Definition 20.1. This reproduces the final assembly of `exists_isFitFamily`. □

The upper bound on projection denominator and the lower bound on theta-lift covolume now meet at exactly the same error budget. The global Selmer argument has been used to produce the local denominator estimate, not to establish a vanishing theorem for a class group, a modularity lifting theorem, or a nonvanishing assertion about automorphic L-functions at critical points. In particular the omission of any even-class-group length term is a consequence of the first-order ties at the auxiliary triangle, and not of assuming those classes vanish.

Chapter 21

The foundations of the congruence argument

The preceding congruence chapter explains how integral projection determinants reduce to a family of bounds on cocycles. In the present chapter we develop the algebraic and arithmetic foundations of that reduction. We isolate the precise places at which the method uses global Galois representations, reciprocity, the construction of auxiliary levels and the extraction of a first-order Hecke relation. The point is not to introduce more abstract hypotheses than the construction needs. Rather, it is to understand why the particular auxiliary data of the proof are effective in the congruence problem: why the primitive labels at two distinct primes are useful, why anticyclotomic twists have been distinguished from triangular auxiliary primes, and how one can account for congruences of Tate characters without a bound for their even classes. We refer to the preceding chapter for the statements of the fit-family and determinant estimates. In particular r' denotes the rank of the distinguished row block, t' the number of Tate characters, and neither is the dimension of the automorphic form space. We use v_ℓ normalized by $v_\ell(\ell) = 1$; in intermediate algebraic arguments v_π denotes the integer-valued valuation of a DVR with uniformizer π. The normalized lengths of finite modules are related to their ordinary lengths by the ramification index, exactly as in Section 20.1. All claims imported from the formal proof without a reconstructed argument are marked as such.

21.1 From weak Hecke compatibility to a trace identity

The version of Galois attachment in the compact development is weaker in one important respect than a usual local-global compatibility theorem. Rather than specifying the complete local Weil–Deligne representation at every place, it specifies a convergence condition at the set of Frobenius data and polynomial annihilation conditions at individual places. We explain the distinction because an argument using the full Weil–Deligne parameter from the outset may inadvertently assume an assertion that is substantially stronger than the one provided in the code.

Let $G_E = \mathrm{Gal}(\overline{E}/E)$. A Frobenius datum comprises a degree-one place w, a choice of a place above it in the algebraic closure, and a Frobenius element of its decomposition group; all finite exceptional sets of places can be omitted. The source has a filter on these data, denoted `mostFrobenii`. It also has its localization near $g \in G_E$, denoted `mostFrobeniiAt`; the latter includes arbitrarily small neighborhoods of g in the Frobenius variable. This formalizes a useful way of applying Chebotarev. If $F(g)$ is continuous, then as Frobenius data tend to g the values of F at the corresponding Frobenius elements tend to $F(g)$. One should not replace this filter with an unspecified density-one subset: a density statement needs to be

strong enough that after excluding finitely many places there are still Frobenius data in *every* permissible neighborhood. The code derives the relevant properties from its Chebotarev package in `Lib/ClassField/Chebotarev/FrobeniusFilter.lean`.

More precisely, if $P_w(X)$ is a family of monic polynomials of degree n over $\mathbb{C}$, a representation $R : G_E \to \mathrm{GL}_n(\overline{\mathbb{Q}}_\ell)$ is attached to P in the sense used here if it is continuous, trace-nondegenerate and, for every $\eta > 0$, eventually along the Frobenius data one has

$$\left|[X^i]\left(\det(X - R(g_w)) - \iota P_w(X)\right)\right|_\ell \le \eta \qquad (0 \le i \le n). \tag{21.1}$$

The more general definition permits an open group map in place of the identity on G_E, and requires the estimate for every element mapping to the chosen Frobenius. See `Attached` in `Lib/GaloisRep/Main/AttachedDefs.lean`. In these notes trace-nondegenerate means that the trace pairing on the finite-dimensional algebra spanned by the matrices of the representation is nondegenerate. We use the topology of the ℓ-adic field, not the Zariski topology, in (21.1).

Lemma 21.1 (Trace extraction from attachment)**.** *Let R be attached to P as in (21.1). For every $g \in G_E$, the functions $-\iota[X^{n-1}]P_w$ tend, along the Frobenius data localized at g, to $\operatorname{tr} R(g)$. Consequently, if two attached systems have polynomials related eventually by multiplication by a product of known linear factors, and the coefficients of those factors tend to prescribed characters at g, their traces differ by the corresponding sum of characters. The same assertion holds with a rescaling of the roots by the values of one continuous character.*

Proof. The identity $[X^{n-1}]\det(X - A) = -\operatorname{tr} A$ holds for every square matrix. For each positive η the non-Archimedean triangle inequality bounds the discrepancy from $\operatorname{tr} R(g)$ by the maximum of the error in (21.1) and $|\operatorname{tr}(R(g_w) - R(g))|_\ell$. The former is eventually at most η by hypothesis, the latter tends to zero by continuity. The rest follows because the negative next-to-leading coefficient of a product of monic polynomials is the sum of their negative next-to-leading coefficients. Scaling all roots of a polynomial by u multiplies this coefficient by u. Limit uniqueness proves the asserted relations at every g. □

This elementary argument is the passage from Hecke polynomials to the identity of traces used in `exists_row_for_all_twists`. It does not imply that at a fixed exceptional place every particular Hecke polynomial is necessarily the characteristic polynomial of Frobenius. The source instead uses the following separate assertion.

Proposition 21.2 (Formal-source input: spherical compatibility)**.** *For the standard definite unitary group over the chosen field with its Weil data, every nonzero Hecke eigenform of rank n has an attached representation. If the form is spherical at a place w, any attached representation is unramified off ℓ, respectively crystalline above ℓ, at the corresponding decomposition group; it has the jumps prescribed by the weight, and every specified Hecke polynomial at w, transported by ι, annihilates the specified geometric Frobenius action (in the local polynomial sense of the source).*

The exact statement is `sphericalCompat_std` in `Lib/GaloisRep/Main/SphericalAllRanks.lean`, with `SphericalCompat` defined in `Lib/GaloisRep/Main/AttachedDefs.lean`. It includes the local jumps `jumps`, the predicate `IsGoodAt`, and the polynomial predicate `PolyKillsAt`. Its proof uses the construction for odd rank, the theory of theta lifts, and local comparison statements. We retain these geometric and local comparison assertions as formal-source inputs. A considerable algebraic part of the reduction from odd rank to general rank, however, can be reconstructed directly as follows.

Lemma 21.3 (Trace-nondegenerate representations and cancellation)**.** *Over an algebraically closed field L of characteristic zero, a finite-dimensional representation whose trace pairing on its image algebra is nondegenerate is semisimple. Its isomorphism type is determined by its trace function on the group. Suppose R_0, R_1 are such representations and C_0, C_1 are finite direct sums of characters, with no character occurring in both C_0 and C_1. If*

$$\operatorname{tr} R_0 - \operatorname{tr} C_0 = \operatorname{tr} R_1 - \operatorname{tr} C_1,$$

then there is a semisimple representation R such that $R_i \simeq R \oplus C_i$ for $i = 0, 1$. If the initial representations are continuous, R can be chosen continuous.

Proof. Write $A \subseteq \operatorname{End}_L(V)$ for the image of the group algebra. The radical J of the finite-dimensional algebra A is nilpotent. For any $a \in A$ and $j \in J$, the element aj is nilpotent and its trace in V is zero. Thus J lies in the radical of the trace form. Nondegeneracy forces $J = 0$. The algebra is therefore a product of full matrix algebras and V is a direct sum of simple A-modules. In the converse direction the trace form on an image algebra of this type is the direct sum of the matrix trace forms multiplied by the corresponding positive multiplicities; it is nondegenerate in characteristic zero.

For two semisimple representations take the image of the group algebra on their direct sum. It too is semisimple: the direct sum is semisimple and faithful for that algebra, so its radical acts both trivially and faithfully, whence is zero. Let the simple factors be $M_{d_i}(L)$. The central idempotent of the ith factor acts as the identity on every copy of the corresponding simple representation and as zero on the others. Its trace divided by d_i recovers the multiplicity. A trace identity on the group extends linearly to the group algebra and hence to all these idempotents. This proves determination by trace.

Now compare the multiplicities in $R_0 \oplus C_1$ and $R_1 \oplus C_0$. For a character occurring in C_0, its multiplicity in R_0 is at least that in C_0, since its multiplicity in C_1 is zero. The corresponding conclusion holds for C_1. All other multiplicities in R_0 and R_1 agree. Subtract the corresponding sums of characters and take the common semisimple remainder. As a direct summand of a continuous representation on a finite-dimensional valued vector space it is continuous (the coordinates of the action in a fixed basis are obtained by continuous matrix operations). □

The cancellation in the source is `IsTraceNondeg.exists_cancel_chars`. It is applied to two auxiliary theta lifts of odd target rank. To orient this construction, if the original rank is n, put $u = n + 1$, so the auxiliary rank is $n + u = 2n + 1$. Choose two distinct unramified characters of $U(1)$ whose infinity weights differ by a nonzero multiple of the torsion order of E at one place. Associated to either choice is a string of u algebraic characters whose Frobenius values at split places supply the extra factors in the Hecke polynomial of the theta lift. Distinct weight choices are arranged so that the two strings share no character. The construction is the first part of the proof of `sphericalCompat_std`. The character existence and the nonzero theta lifts of suitable joint-harmonic weight are formal-source inputs there.

The source establishes that at every split spherical place the extra factors $Q_{i,w}$ in the two lifts are pairwise disjoint between strings. The resulting algebraic check is useful independently of the global construction.

Lemma 21.4 (Two cancellations of auxiliary factors)**.** *Suppose R_i are as in Lemma 21.3 with common remainder R, and let P_w be the sought Hecke polynomial. Assume at a given place, after the normalization and removal of the known characters, the polynomial operators $P_w(F)Q_{i,w}(F)$ vanish on the relevant local space of R for both i. If $Q_{0,w}$ and $Q_{1,w}$ are coprime, then $P_w(F)$ vanishes there. Furthermore, if $P_w Q_{i,w}$ are monic polynomials with coefficients in a valuation ring and $Q_{i,w}$ has coefficients in this ring, then so does P_w.*

Proof. Choose A, B with $AQ_{0,w} + BQ_{1,w} = 1$ and substitute F; all polynomials in F commute. Multiply the identity by $P_w(F)$. For the integral assertion, successively compare the highest coefficients of $P_w Q_{i,w}$. The coefficient of each degree in P_w is expressed as one integral coefficient of the product minus integral combinations of already determined higher coefficients. No division is needed because the divisor is monic. Equivalently use the multiplicativity of the non-Archimedean Gauss norm and the fact that a monic integral divisor has Gauss norm one. □

This is the final polynomial step in `sphericalCompat_std`: both lift constructions have the same trace remainder and each yields the required local comparison after introducing its own extra polynomial. The local comparison itself, including passage of the crystalline conditions and jumps through the trace limit, is the formal-source input `local_clauses_of_trace_limit` used there. This explains why merely quoting an odd-rank construction and deleting extra characters from a global trace identity would not be sufficient.

21.2 Tame labels and full matrix span

We now turn to the local property of the Galois block on which the Selmer argument relies. The proof requires the strong integral statement $\mathrm{span}_{\mathcal{O}} R(G_E) = M_{r'}(\mathcal{O})$ for every sufficiently large coefficient ring. Irreducibility only over its fraction field is insufficient. The two primes g_0, g_1 and their primitive labels are specifically chosen for this purpose.

Lemma 21.5 (Metacyclic tame-pair lemma)**.** *Let $\mathcal{O}$ be a DVR with fraction field L and let $R : G \to \mathrm{GL}_n(\mathcal{O})$ be a representation. Suppose $s, f \in G$ and an integer $q \geq 2$ satisfy*

$$R(f)^{-1} R(s) R(f) = R(s)^q, \qquad R(s)^{q^n} = R(s).$$

Suppose an eigenvalue $\zeta \in \mathcal{O}$ of $R(s)$ is a unit and the n numbers ζ^{q^i} $(0 \leq i < n)$ are pairwise distinct modulo the maximal ideal. Then n^2 matrices in the image span $M_n(\mathcal{O})$ over $\mathcal{O}$, and $R(f^n)$ is a scalar unit.

Proof. Reduce the matrices modulo the maximal ideal. Let $A = \overline{R(s)}$, $B = \overline{R(f)}$, and choose a nonzero vector v with $Av = \bar{\zeta} v$. The relation implies that $B^i v$ has eigenvalue $\bar{\zeta}^{q^i}$ for A. Consequently the vectors $v, Bv, \ldots, B^{n-1}v$ form a basis of the residue vector space. Over it the matrix A is diagonal with pairwise distinct entries. Its spectral projectors E_i are polynomials of degree $< n$ in A, by Lagrange interpolation. The matrix B cyclically moves the eigenspaces, with nonzero scalar at the wraparound, since it is invertible. Hence $B^j E_i$, $0 \leq j < n$, $0 \leq i < n$, are nonzero multiples of all matrix units in the given basis. Each is a linear combination of $B^j A^k$ for $0 \leq k < n$. There are n^2 such matrices, so their lifts $R(f^j s^k)$ reduce to a basis. Nakayama's lemma, or the determinant criterion over the local ring, implies these lifts form an $\mathcal{O}$-basis of the entire matrix algebra.

The second tame-pair relation gives $R(f)^n$ commuting with $R(s)$; it plainly commutes with $R(f)$. It therefore commutes with the $\mathcal{O}$-basis just exhibited. The centralizer of the entire matrix algebra in itself is its scalar center, as one sees by commuting successively with diagonal matrix units and off-diagonal matrix units. Thus $R(f)^n$ is scalar, and invertibility makes the scalar a unit. □

This proves the key statement `IsLatticeIn.span_eq_top_of_tame_pair` in `Lib/General/Groups/ResidualIrreducibility.lean`; the source takes f^{-n} rather than f^n for its scalar, which is equivalent. Notice that if the n eigenvalues were distinct merely in characteristic zero, the Lagrange denominators might have positive valuations and the resulting basis would

not be integral. In that case there would be no reason for the matrix span to be all matrices over $\mathcal{O}$.

We give the elementary number theory behind the special primitive-label property. Let $q \equiv 1 \pmod 9$, let $n = 3^h$ for $h \geq 1$ and put $n_0 = n/3$. The quotient

$$S = \frac{q^n - 1}{q^{n_0} - 1} = q^{2n_0} + q^{n_0} + 1$$

is divisible by 3 exactly once: reducing q^{n_0} modulo 9 gives $1 + 1 + 1 = 3 \pmod 9$. Moreover $S > 3$, so it has some prime factor other than 3. It follows that S is not a power of any single prime. Suppose ζ has order $q^n - 1$ and fix any prime ℓ. The order of ζ^{ℓ^a} is $(q^n - 1)/\gcd(q^n - 1, \ell^a)$. If its q-orbit had period $< n$, that period, a proper divisor of $n = 3^h$, would divide n_0. Then $q^n - 1 \mid \ell^a(q^{n_0} - 1)$, forcing $S \mid \ell^a$, which is impossible. Thus all n powers $(\zeta^{\ell^a})^{q^i}$ are distinct for every $a \geq 0$. By the characterization of reduction of roots of unity from the preceding chapter, all the ζ^{q^i} are residually distinct. This is `frobenius_conj_pow_injOn` in `Lib/LocalGL/Green/LabelIdempotents.lean`. The case $n = 1$ requires no orbit argument.

The same argument applies when the primitive element is multiplied by an element z of the base residue field of cardinality q. Indeed in the ratio of the ith and jth q-conjugates the base-field factor cancels. This is precisely the scaling of the label by the inverse of the tame part of the automorphic character in `twistedLabel`. In particular the proof uses the *primitive original label*, not a claim that its scaled version is still primitive. If the order of a primitive element is $q^n - 1$ and u has order dividing M_ω in the base field, the product of the n conjugates of the scaled element is the norm of the primitive element times u^n. Its norm before scaling is primitive of order $q - 1$. The order of the scaled product is at least $(q-1)/M_\omega$ in the divisibility sense: if its jth power equals one, then $q - 1 \mid M_\omega j$. This follows by raising the norm identity to M_ω. The source uses the fact that $q - 1$ is divisible by M_ω^2, and $M_\omega = (N_0 d)^{s+6}$ has more than one prime divisor. Choose any prime divisor $p \mid N_0 d$ with $p \neq \ell$; the p-part of the norm of the scaled label then still has order at least p^{s+6}. This grows faster than the permitted small powers $4r' + 2$. A corresponding power of the inertia generator acts by a scalar on the row, while all those small powers of that scalar remain distinct from one after reduction. This is the mechanism of `IsLabelledColumn.isRow`; below we specify the polynomial input needed to justify passing from labels to the tame generator.

Lemma 21.6 (Valuation of a nontrivial root of unity)**.** *If $\lambda \neq 1$ and $\lambda^D = 1$ in $\overline{\mathbb{Q}}_\ell$, then $v_\ell(1 - \lambda) \leq v_\ell(D)$. If a matrix A of size n is annihilated by a polynomial all of whose roots differ from one and whose roots satisfy this condition with a common D, then*

$$v_\ell(\det(A - 1)) \leq n\, v_\ell(D).$$

Proof. Every root of unity is a unit. Since $\sum_{i=0}^{D-1} \lambda^i = 0$, we have

$$D = \sum_{i=0}^{D-1} (1 - \lambda^i) = (1 - \lambda) \sum_{i=1}^{D-1} (1 + \lambda + \cdots + \lambda^{i-1}).$$

The last factor is integral. This proves the first assertion. Every eigenvalue of a matrix is a root of every polynomial that annihilates it, by evaluating that polynomial on the eigenspace. The determinant of $A - 1$ is the product of its eigenvalues minus one, with multiplicities, and the sum of their valuations gives the claimed bound. □

The common integer needed in the determinant bound can in fact be chosen less crudely than $q^{n!} - 1$. The source allows each root of an inertia annihilator to have order dividing $g_i^e - 1$

for some $1 \le e \le r$. Each such integer divides M_g (the group order defined in (20.4)). Thus $v_\ell(1-\lambda) \le v_\ell(M_g)$ separately for each root, and the same proof gives the bound $r'c_\ell$ of the preceding chapter. If $q = g_i \neq \ell$, a tame generator at a place above q generates the required finite cyclic action away from its wild inertia. At a prime above g_i the row eigenvalues are the roots of the scaled label transported by the residue-field character. The other chosen prime can be used if $g_i = \ell$. Their independent presence is the reason a residual irreducibility conclusion holds for every ℓ, even when ℓ is one of the two large level primes.

Proposition 21.7 (Formal-source input: local labels)**.** *For an apt pair at the chosen primes above g_i, each element of inertia has a common annihilating polynomial on the labelled columns and their conjugate inverses. The polynomial is the product of $r'+t'$ linear factors $X-\lambda$; each root obeys $\lambda^{g_i^{(r')!}-1} = 1$ and $\lambda^{g_i^e-1} = 1$ for some $1 \le e \le r$. At the distinguished generator the roots for $t' = 0$ are nontrivial. At the twist primes outside the original level, the normalized representation is annihilated on inertia by $(X-\psi(g))(X-1)$, where ψ is the inverse ℓ-adic avatar of the new twisting character. Twisting by the character needed to normalize a row multiplies the corresponding inertia labels and shifts the Hodge jumps by its local weights.*

The assertions are respectively `exists_inertia_eigenvalues`, `normalisedRep_quadratic_on_inertia`, and `IsLabelledColumn.twist_avatar`, in `New/Selmer/LabelledColumns.lean`. These use the parahoric local-global comparison `Lib/GaloisRep/LocalGlobal/LabelCompatibility.lean`. The latter supplies the polynomial killing conditions, not an assertion that a characteristic polynomial computed by a global trace uniquely determines a local monodromy operator. Subject to this input, Lemma 21.5 and the primitive-root and determinant calculations give the full matrix-span and small inertia determinants of `IsRow`. The remaining row conditions are the jump bounds and conjugate self-duality; these enter via Proposition 21.2 and the particular shifts of the algebraic weight in `exists_hasRow_auxPair` of `New/Selmer/RowsExist.lean`.

21.3 A duality calculation at the mirror

The role of the scalar inertia element is more subtle than simply making a representation irreducible. It is used a second time in an estimate for a family of additive functionals. We now prove the algebraic part of this estimate. It is important that these functionals are additive characters, not necessarily linear functionals over the coefficient ring.

Write c for the conjugation involution of the group, so cg in the following lemma means the conjugate of the element g, not multiplication by a fixed element. Let $A = \mathcal{O}/I$ be a finite quotient of the coefficient DVR, let $M = A^n$, and let R and χ be the row representation and a scalar character on this module. The twist by χ occurs on the dual of the row, not on the row itself. This distinction accounts for the different factors in the mirror functional.

Lemma 21.8 (Mirror calculation)**.** *Suppose $H \subset U \subset G_E$ are normal conjugation-stable subgroups, with U acting trivially through R and χ. Let γ satisfy*

$$R(\gamma) = s\,1, \qquad R(c\gamma) = s^{-1}\,1, \qquad \chi(\gamma) = \chi(c\gamma) = 1,$$

where $s-1$ and $s+1$ are units. Suppose H contains $[g,\gamma]$ for all $g \in G_E$. Let $y : G_E \to M$ be an R-cocycle, with $y(\gamma) = 0$ and with $y(g_0)$ having a unit coordinate. Assume $R(G_E)$ spans $M_n(A)$. Let W be a finite additive family of maps $x : U \to M^\vee = \mathrm{Hom}_{\mathbb{Z}}(M, \mathbb{Q}/\mathbb{Z})$ such that the maps are additive group homomorphisms and satisfy

$$x(kbk^{-1})(u) = x(b)(\chi(k)R(k)^{-1}u) \quad (b \in U).$$

Assume for a fixed $z \in U$ and all $h \in H$, $a \in A$, that

$$x(zh)(ay(zh)) = x(c(zh))(ay(c(zh))). \tag{21.2}$$

If there is an element g with $R(g) = \alpha\, 1$ and $R(cg)^{-1} = \alpha'\, 1$, put $D = \alpha - \chi(cg)\alpha'$. Then

$$|W| \le |M[D]| = |M/DM|.$$

Proof. On U the map y is additive. Write $\tau = [g_0, \gamma] = g_0\gamma g_0^{-1}\gamma^{-1}$. Expanding the cocycle equation gives $y(\tau) = (1 - s)y(g_0)$, so it still has a unit coordinate. For $b, b' \in H$, take the alternating difference of (21.2) at bb', b, b' and the identity. The terms involving z disappear, and what remains is

$$\begin{aligned} A(a) + B(a) &= C(a) + D_1(a), \\ A(a) = x(b)(ay(b')), &\quad B(a) = x(b')(ay(b)), \\ C(a) = x(cb)(ay(cb')), &\quad D_1(a) = x(cb')(ay(cb)). \end{aligned}$$

The cocycle transform on H is $y(kbk^{-1}) = R(k)y(b)$, because $R(b) = 1$. Replace b by $\gamma b\gamma^{-1}$ and a by sa. The four terms become $A(a)$, $B(s^2a)$, $C(s^2a)$ and $D_1(a)$: here both s^{-1} factors on the dual side are essential. Comparing with the untransformed equality at s^2a yields

$$(A - D_1)((s^2 - 1)a) = 0.$$

This step uses additivity in a, not A-linearity of x. Since $s^2 - 1$ is a unit, it follows that $A(a) = D_1(a)$ for every a.

Take $b' = k\tau k^{-1}$. We have proved for each $b \in H$ the stronger identity

$$x(b)(aR(k)y(\tau)) = x(c\tau)(a\chi(ck)R(ck)^{-1}y(cb)). \tag{21.3}$$

Because one coordinate of $y(\tau)$ is a unit and the row matrices span the full matrix algebra, the elements $aR(k)y(\tau)$ generate M additively. Thus the single functional $x(c\tau)$ determines x on H. Taking $b = c\tau$ in (21.3) and comparing its values at k and kg shows

$$x(c\tau)((\alpha - \chi(cg)\alpha')aR(k)y(\tau)) = 0.$$

Indeed the left argument is multiplied by α, while the right one is multiplied by $\chi(cg)\alpha'$, and the identity at k permits us to move that scalar to the left. Full span now shows that $x(c\tau)$ vanishes on DM.

Finally, for $b \in U$, the map x on $[b, \gamma]$ equals $x(b) \circ (1 - s^{-1})$ by additivity and the transformation rule. Since the commutator is in H and $1 - s^{-1}$ is a unit, determination on H determines x on all of U. We have an injection of W into the additive dual of M/DM. The latter has the same size as M/DM by finite abelian duality. The kernel and cokernel of multiplication by D on M have equal cardinalities, proving the assertion. □

This proof is the mechanism of `card_le_torsion_of_isotropic` in `New/Selmer/LinePointPowers.lean`. The actual source formulates the equivariance and subgroup hypotheses with a family of points and line powers. Two features should be emphasized. First, the scalar of γ is required to differ from both 1 and -1 modulo the maximal ideal; the reason is the cross-difference, which has coefficient $s^2 - 1$. Second, a uniform estimate for the additive dual follows from a single scalar defect D, independent of the number of cocycles initially used to formulate W. If D is nonzero, its normalized length contribution is at most $nv_\ell(D)$, since the kernel of multiplication by D on $\mathcal{O}/\pi^u$ has ordinary length $\min(u, v_\pi(D))$. If $D = 0$ this observation yields no saving and cannot be used in a finite quantitative balance argument.

The global application has additional steps not implicit in this calculation. One must extract the additive dual family from a system of local restrictions, realize the mirror relations at split places, and choose a Frobenius point seeing a class that has not yet been detected. The formal-source input for these steps is `exists_split_primes_relations` in `Lib/ClassField/GreenbergWiles/SelmerRelations.lean` and the application `IsPencil.length_le_of_unseen_class` in `New/Selmer/TiltingLemma.lean`. The calculation just given explains why there is a defect term of size $nv_\ell(D)$ in that application. It is not a replacement for the global reciprocity theorem.

21.4 Anticyclotomic primes from Kummer extensions

An anticyclotomic auxiliary prime adds a character variable to the twists already considered. The requirement that the character be trivial at all the old places would be quite restrictive if one tried to choose a ray class character after fixing an arbitrary split prime. Instead the prime is selected to solve that restriction. This is an instance where the order of construction matters.

Here is a convenient class field formulation. For a finite conjugation-stable set T of finite places, the T-units are the elements of $E^\times$ which are units off T. For a new prime $P \notin T$, an additive character initially specified on the units at P extends to the idele class group, trivially on the components at T, if it kills the ambiguity given by the T-units and the remaining ideal classes. Killing the latter ambiguity is accomplished by enlarging T before constructing P. The class group of the ring of T-integers can be made trivial by adding finitely many outside primes. If another finite set $R \setminus T$ is excluded, one can still choose these primes outside R: every ideal class is represented by a prime outside any finite set, by the Hilbert class field and Chebotarev. The idele form of this observation is `exists_sup_idelesUnitOff_eq_top`, used in the following construction. We take the usual reciprocity and Kummer correspondence as inputs here.

Proposition 21.9 (Controlled auxiliary character)**.** *Fix a rational prime ℓ, a positive power Q of ℓ, a finite set S of places at which the character must be trivial, a finite set R, and z in the closed normal closure of inertia of places in R away from $S \cup cS$ and the primes above ℓ. For any neighborhood W of z in G_E there is a degree-one prime $P \notin R \cup cR \cup S \cup cS$ and an anticyclotomic character $h : G_E \to \mathbb{Z}/Q\mathbb{Z}$ such that $NP \equiv 1 \pmod{Q}$. There are a Frobenius lift f at P in W and a tame inertia generator τ at that place with*

$$h(f) = 0, \qquad h(\tau) = 1.$$

The character is trivial at S and at primes above ℓ, and unramified outside P, cP. One may also specify a finite set of previously chosen primes to be avoided by P.

Explanation of the construction. Enlarge $S \cup cS \cup \{v : v \mid \ell\}$ to a finite set T killing the ideal classes as above, but avoid $R \cup cR$ outside this prescribed union. Take all conjugates of the additional primes if necessary. The group of T-units is finitely generated, including its finite torsion group. Write $u_1, \dots, u_d$ for generators. Consider the finite extension obtained by adjoining the Qth roots of unity and the Qth roots of these generators (and take its normal closure if needed). Inertia at each allowed place of R fixes this extension: in residue characteristic prime to Q, adjoining a Qth root of a unit gives an unramified extension. This last fact follows locally by Hensel's lemma after passing to the finite unramified extension whose residue field contains all the needed roots. The normal closure of these inertia groups also fixes the extension. In particular z acts trivially on it.

Choose an open normal subgroup small enough that a Frobenius element in a specified coset of it lies in W and fixes the Kummer extension. Chebotarev supplies such a Frobenius element at P, with P of degree one and split over the real subfield; finitely many primes can be excluded. Splitting in the roots-of-unity extension gives $Q \mid NP - 1$. Every u_i becomes a Qth power modulo P. The residue multiplicative group is cyclic, so choose a logarithm into $\mathbb{Z}/Q\mathbb{Z}$ taking a fixed residue generator to 1. This logarithm vanishes on the residues of all T-units. Extend it to $E_P^\times$ by killing the chosen uniformizer.

On the ideles supported away from T on their unit factors, take the logarithm of the unit component at P, and zero on the other components. Since the class group of T-integers is trivial, every idele is represented, modulo $E^\times$ and the local factors at T, by such an idele. Any ambiguity in this representative comes from a T-unit, on which the logarithm vanishes. Thus the logarithm extends to the idele class group, and is continuous with finite image. Compose with global reciprocity to obtain h_0. The character $h = h_0 - h_0 \circ c$ is anticyclotomic. At P, the term from cP does not modify its inertial value; choose the tame generator with the logarithm 1, as specified by the reciprocity convention. It is trivial at T and unramified outside P, cP.

Finally change the Frobenius lift at P by a suitable power of τ, so that $h(f) = 0$. The finite extension used in defining the Frobenius neighborhood is unramified at P, because its ramified places can have been excluded. This change therefore preserves membership of f in W. The congruence for NP is independent of the Frobenius lift. This obtains all the stated properties. □

The precise theorem for the given data is `exists_isAnticyclotomicAt` in `New/Selmer/Corepresentable.lean`. In Lean the extension from local characters is expressed by `IsFromU1`, whose hypotheses include the restriction to ideles at the set T; the data of the tame generator and Frobenius are stored in `IsAnticyclotomicAt`. Our proof above describes the arithmetic construction with the standard Chebotarev and class field inputs, rather than developing these general theorems anew.

Observe why the hypothesis on z is important. Arbitrary prescribed z need not fix the roots of T-units; one cannot in general ensure a Frobenius near that element and also impose the character condition. Normal closure of allowed inertia solves the problem because that entire closure fixes the Kummer extension. The fact that P depends on the neighborhood is also indispensable, as is freedom to replace the Frobenius lift. The construction permits previously obtained restrictions at finite places to remain valid when new characters are added. This is the anticyclotomic operation used in the pencil argument, not the triangular operation that follows. The latter is directed instead at the scalar summands of the trace identity.

21.5 Odd laws and the triangular removal of even classes

We next give a proof of the triangular construction, conditional on two substantial arithmetic inputs. This construction is logically distinct from the ordinary patching problem in which one first bounds the dimension of a residual Selmer group and chooses enough primes to kill it. Here the classes in a nonzero Tate twist may be even. Their required bound is obtained by arranging the order of the primes, not by supplying an independent bound for the size of the even group in every Tate twist. The technique produces the same number of primes at every sufficiently deep finite coefficient quotient. Their locations and their Frobenius lifts may change with the quotient.

Let $A_\nu = \mathbb{Z}/\ell^\nu\mathbb{Z}$, write ε for the cyclotomic character, ε_ν for its reduction to $A_\nu^\times$, and set $K = \ker \varepsilon_\nu$. Let c act on G_E by conjugation. A scalar cocycle d of twist $-e$ in a finite

quotient of a coefficient DVR has the convention

$$d(gh) = d(g) + \varepsilon_\nu(g)^{-e} d(h).$$

It is even if $d(cg) = d(g)$ and is required to satisfy the stated local Selmer conditions. On K it is additive. We need odd functions $Z : G_E \to A_\nu$ which restrict to additive maps on K, satisfy $Z(cu) = -Z(u)$ on K, and transform under conjugation by

$$Z(kuk^{-1}) = \varepsilon_\nu(k)^{e+1} Z(u) \qquad (u \in K). \tag{21.4}$$

They are locally constant. In addition to these elementary requirements the function must satisfy a reciprocity condition, stated in terms of a finite set F of old Frobenius lifts and a finite set T of exceptional primes of the real subfield. If a collection of new split places has distinct underlying real primes off T, Frobenius in K, and Z vanishes at all of these Frobenius lifts except perhaps the one at P_j, the following inference holds. For every even class d of twist $-e$, unramified away from ℓ and the new places and their conjugates, vanishing on F, and every integer b,

$$2Z(f_j) \mid b \text{ in } A_\nu \quad \Longrightarrow \quad bd(\tau_j) = 0. \tag{21.5}$$

This is the structure `IsLaw` in `New/Selmer/ReciprocityLaw.lean`; the local Selmer conditions in the source are carried by `evenClasses`. It is an assertion about all finite quotients of coefficient rings and not just $\mathbb{F}_\ell$.

Proposition 21.10 (Formal-source inputs: existence and detection)**.** *Under the field hypotheses of the proof, in particular the stable unramifiedness and the lower bound of four for the number of infinite places, given k old split places as above and at most k prescribed elements of K, there exists a law Z of type e for the old Frobenius lifts which vanishes at the prescribed elements and takes a unit value at some element of K. Independently there is a finite set D_0 of detecting points, of size at most y^{20} when the field is contained in the cyclotomic field of order $N_E < y$ and $\ell \le y^2$, on which vanishing detects all cocycles of a Tate twist with the specified Selmer conditions and without extra ramification. The detecting set can be used at every finite coefficient level and for all twists in question.*

The inputs are `exists_isLaw` in `New/Selmer/ReciprocityLaw.lean` and `exists_detecting_points_cocycles` in `Lib/Cyclotomic/OddSha/OddShaBound.lean`. The proof of the law uses global duality over the real subfield. To indicate why the quantitative lower bound on the number of choices matters, we record an algebraic step in the proof of primitive relations.

Lemma 21.11 (A primitive finite relation)**.** *Let $A = \mathbb{Z}/\ell^b\mathbb{Z}$ with $b \ge 1$, let Y be a finite abelian group and let $p : Y \to A^I$ be additive. Fix k coordinate positions in I. If $|Y[\ell]|\ell^{k+1} \le \ell^{|I|}$, there is $w \in A^I$ vanishing in those k positions, annihilating $p(Y)$ under the dot product, and primitive (at least one coordinate is a unit).*

Proof. Take the homomorphism

$$A^I \longrightarrow \mathrm{Hom}_{\mathbb{Z}}(Y, A) \oplus A^k, \quad w \longmapsto (y \mapsto w \cdot p(y),\ (w_i)_{i \text{ prescribed}}).$$

Denote its image by B and its kernel by L. If no relation in L were primitive then $L \subset \ell A^I$, so $B/\ell B$ would have dimension $|I|$ over $\mathbb{F}_\ell$. The structure theorem gives $\dim B/\ell B = \dim B[\ell]$. Since $B[\ell]$ embeds in the ℓ-torsion of the displayed target, its dimension is at most $\log_\ell |Y[\ell]| + k$. For the first factor this follows by decomposing Y into cyclic primary groups; homomorphisms to $A[\ell]$ have one parameter for each ℓ-primary cyclic summand. The asserted bound is strictly less than $|I|$, a contradiction. □

This reproduces the counting step in `exists_primitive_relation` of `Lib/ClassField/GaloisCohomology/DiscreteUnits.lean`. Producing Y with the required bound is part of the global argument, not a consequence of the lemma alone. In particular primitive relations should not be assumed merely because the number of places is large.

We now explain how Proposition 21.10 supplies the exact form of the triangle. Fix an element a with separated cyclotomic values, use $\gamma = a^{-1}$, and suppose each $\varepsilon(\gamma)^u - 1$ with $0 < u < t'$ divides ℓ^h in the coefficient ring. The set J consists of pairs (e, i) for $1 \le e < t'$ and $i \in D_0$, ordered lexicographically, so it has types e and each type occurs $|D_0|$ times. Write

$$b_i = (ai)^{-1}(ia) \in K,$$

using the point i itself as a group element. We seek distinct primes with Frobenius $f_j \in K$ of types e, with a law Z_j at each step. The law Z_j must satisfy reciprocity with respect to the earlier Frobenius lifts of the same type; it should vanish at every later chosen Frobenius lift; and $2Z_j(f_j)$ should divide ℓ^H. For control of the cocycles without extra ramification we will also arrange

$$d(f_{(e,i)}) = 2d(b_i) \tag{21.6}$$

for every even unramified class of type e (unramified here and below means outside ℓ).

Consider the step $j = (e, i)$. The evaluations at all possible laws of all types, restricted to K, form an additive subgroup of the product of the groups A_ν. Indeed the laws are additive on K, so sums and inverses of simultaneous evaluation vectors are realized at products and inverses of elements of K. Project this group to the coordinates of the old laws of type e. A subgroup of A_ν^k needs at most k generators: its minimum number of generators equals the dimension of its ℓ-torsion, bounded by k. Choose elements $y_1, \dots, y_v \in K$ whose evaluation vectors generate that projection. By the law existence assertion choose Z_j vanishing on these elements and primitive at an element $w \in K$. Subtract a suitable integral combination of the evaluation vectors of the y_i from that of w. The result is realized at an element $u_0 \in K$ where every earlier law of type e vanishes and Z_j still takes its unit value.

Earlier laws of the other types are removed by a polynomial in a single conjugation operation. On simultaneous evaluation vectors write C for conjugation by γ, which is defined on this subgroup. Its action on a coordinate of type d is multiplication by m^{d+1} with $m = \varepsilon_\nu(\gamma)$. Replace the vector of u_0 by

$$\prod_{\substack{0 \le k < t' \\ k \ne e}} (C - m^{k+1}) \operatorname{ev}(u_0).$$

The resulting vector is still realized at a point $u \in K$ since all its operations are conjugation, sums and integer multiplications (choose integral representatives of the coefficients in A_ν). The earlier laws all vanish there; the new law takes its original unit value multiplied by

$$m^{t'-1} \prod_{k \ne e} (m^e - m^k).$$

The index 0 in this product causes no harm although the set of positive types omits it. It will allow all Tate-character differences to occur in one bound.

The controlling divisibility is calculated before reducing modulo ℓ^ν. For $e > k$ one has

$$\varepsilon(\gamma)^e - \varepsilon(\gamma)^k = \varepsilon(\gamma)^k(\varepsilon(\gamma)^{e-k} - 1),$$

and for $e < k$ one uses the negative of the reverse difference. All prefactors are units. Hence every difference divides ℓ^h, and their product divides $\ell^{(t'-1)h}$. The source reserves three more powers of ℓ^h to handle a factor 2 and a squared character difference. If $\ell \ne 2$ the factor 2 is a

unit; for $\ell = 2$ the field contains the fourth roots of unity and $4 \mid \varepsilon(\gamma) - 1$, so $4 \mid \ell^h$ provided there is at least one positive type. With $t' + N_E < y$ we may use the uniform bound

$$H = yh, \qquad 2 \prod_{k \neq e} (\varepsilon(\gamma)^e - \varepsilon(\gamma)^k) \mid \ell^H.$$

It is legitimate to use this divisibility even when some factors reduce to zero in A_ν: in that case the right-hand side also reduces to zero. The assertion concerns divisibility, not the division of an element of a quotient ring by a nonunit. At a deeper level with $\nu > H$ it implies a genuine nonvanishing statement.

To prescribe both the older laws and the new detector condition, form the element

$$b_i \, cb_i \, u \, (cu)^{-1} \in K.$$

Every law takes there the value $2Z(u)$ since it is odd, and every even class takes the value $2d(b_i)$. Thus the prospective Frobenius can be chosen simultaneously for both tasks. The genericity condition on Frobenii in `exists_isPrimeTriangle` permits the neighborhood to be translated by a factor δ^{ℓ^ν} with $\varepsilon(\delta) = 1$. This factor has zero value under any of the additive functions being considered. It does not alter the prescriptions. The locally constant laws and the finitely many scalar cocycles in A_ν give finitely many open conditions, and one omits the finitely many exceptional real primes for the law and its predecessors. Chebotarev therefore supplies a degree-one prime P_j not over a previously used real prime. This yields norm congruent to 1 modulo ℓ^ν and the required Frobenius in the chosen generic set. To pass from scalar cocycles in A_ν to arbitrary coefficient quotients use their additive characters into A_ν: those characters are jointly faithful by finite abelian duality and commute with the action of $\mathbb{Z}/\ell^\nu\mathbb{Z}$. They preserve the scalar twist and the even condition. This is the content of `eq_zero_of_mem_evenClasses` in the same file as the laws. This descent is essential for applying Chebotarev with finitely many conditions independently of the eventual coefficient ring.

Let us verify the two consequences of this construction.

The triangular ramification step. Suppose d is an even class of type $e = \mathrm{ty}(j)$, unramified outside ℓ and the places in the list, satisfying $d(\tau_i) = 0$ for every $i < j$ and $d(f_i) = 0$ for earlier i of the same type. It is then unramified at the earlier auxiliary places. For the conjugate places this follows from evenness. At a place prime to ℓ the wild inertia is pro-prime-to-ℓ and contributes no homomorphism to the finite ℓ-group; hence vanishing at the specified tame generator suffices. Apply the law Z_j to the remaining places $i \geq j$ (these avoid its exceptional real primes), with the earlier same-type Frobenius lifts in its prescribed set F. It vanishes on all f_i for $i > j$ by construction, while $2Z_j(f_j) \mid \ell^H$. Equation (21.5) gives $\ell^H d(\tau_j) = 0$. In an induction one must include the earlier Frobenius conditions as well as the earlier inertial ones; their separation into types is part of the triangle, not an optional strengthening of a weaker statement.

The remaining unramified class. Suppose finally that a class d of positive type e has no auxiliary ramification, and vanishes on every Frobenius of that type. Put $\xi(g) = \varepsilon_\nu(g)^{-e}$ and $\Delta(g) = \xi(g) - 1$. The function

$$B(g) = \Delta(a)\, 2d(g) - \Delta(g)\, 2d(a)$$

is a cocycle of the same twist, with the same required unramified local conditions. For the first term this is linear stability of the Selmer module, and the second term is a coboundary.

For a detector i the cocycle formula gives

$$d(ai) - d(ia) = \Delta(a)d(i) - \Delta(i)d(a),$$
$$d(b_i) = \xi(ai)^{-1}\big(d(ia) - d(ai)\big).$$

By (21.6), $2d(b_i) = 0$. Thus B vanishes at every $i \in D_0$, so the detector assertion gives $B = 0$. This is precisely the coboundary relation

$$\Delta(a)\, 2d(g) = \Delta(g)\, 2d(a).$$

We also have $4\Delta(a)^2 \mid \ell^H$ by the same factor estimates. The relation controls the remaining class with a scalar denominator; no division modulo a possibly noninvertible character difference has been assumed. The patching argument can now work with the whole even group, whereas a claim that it was uniformly small would require a much stronger arithmetic input.

Combining the two consequences with the detector bound gives

$$(|J| + |D_0|)H = t'|D_0|yh \le y^{22}h.$$

If the separation bound has $h = 0$, the entire cost is zero; no estimate of the product of the number of detectors and 0 needs the cardinal bound for small primes. The precise package is `IsPrimeTriangle` in `New/Selmer/Corepresentable.lean`, and its construction is `exists_isPrimeTriangle` in `New/Selmer/TrianglePrimes.lean`. Its field hypothesis, the local conditions in the law and detector assertions, and the genericity of the Frobenius set are needed for the arithmetic existence of the primes; the calculation of H is elementary given those assertions. This explains both the dimension-independent arrangement of the auxiliary levels and the particular power y^{22} that appears later in the bound.

21.6 The algebra of adjoining relations

We turn to the part of the determinant estimate in which lengths of modules of functions must be converted into products of congruence ideals. It is useful to separate this purely integral calculation from bounds on Galois cohomology. Let R be a commutative coefficient algebra, let $N \subset R^d$ be the submodule of rows expressing a system of relations, and denote the ideal generated by the determinants of d rows of N by $\mathfrak{F}(N)$. This is the zeroth Fitting ideal of R^d/N (with the evident determinant convention when $d = 0$). The inclusion of a larger relation module gives a larger Fitting ideal. Thus one cannot enlarge the relation system freely: the determinant gained at the end of an argument has to be brought back into the original ideal with the correct cost.

Lemma 21.12 (Cost of one relation)**.** *If $v \in R^d$ and $qv \in N$, then $q\mathfrak{F}(N + Rv) \subset \mathfrak{F}(N)$. More generally, suppose that $q_j v_j \in N + Rv_1 + \cdots + Rv_{j-1}$ for $1 \le j \le L$. Then*

$$\left(\prod_{j=1}^{L} q_j\right) \mathfrak{F}(N + Rv_1 + \cdots + Rv_L) \subset \mathfrak{F}(N).$$

Proof. Expand a determinant of rows $n_i + a_i v$. Every term using v in two or more rows vanishes by alternation. The term containing no occurrence of v belongs to $\mathfrak{F}(N)$, and after multiplication by q each remaining term is a multiple of the determinant with row $qv \in N$ in the appropriate position. This proves the first assertion. The second is obtained by iteration in descending order of j. □

In the formal source these are `fittingIdeal.mul_mem_of_smul_mem` and `fittingIdeal.triangle` in `Lib/General/Algebra/FittingDuality.lean`. A point of the triangle argument is that a relation can already hold after the previous relations have been added; Lemma 21.12 does not require them all to hold in the first quotient. This avoids paying separately for the unknown number of classes in each twist. If $q_j = \ell^H$, the explicit cost is $|J|H$ in normalized valuation.

There is a related use of lengths which applies when relation forms cannot be specified as a finite ordered list in advance. We give an algebraic model reproducing the induction in `fittingIdeal.exists_cost_le_relLength`. Let $\mathcal{O}$ be a DVR, $p \neq 0$, $A = \mathcal{O}/p$, R an $\mathcal{O}$-algebra. Let $\phi : \mathcal{O}^u \to R^d$ be linear over $\mathcal{O}$, and $z(g) \in \mathcal{O}^u$ be prescribed vectors indexed by a set G. Assume $p\phi(\mathcal{O}^u) \subset N$. For any relation module $L \supset N$, set

$$D_L = \{\lambda \in \mathrm{Hom}_{\mathcal{O}}(\mathcal{O}^u, A) : \lambda(\phi^{-1}(L)) = 0\}, \qquad F_L = \{g \mapsto \lambda(z(g)) : \lambda \in D_L\}.$$

Fix a submodule V of A-valued functions on G. The relative length $\mathrm{length}(F_L/(F_L \cap V))$ is finite; denote it by b_L. No finiteness assumption on the set G is needed because F_L is an image of a finite module.

Proposition 21.13 (Relative-length cost)**.** *There exists $L \supset N$ with $\pi^m \mathfrak{F}(L) \subset \mathfrak{F}(N)$ for some integer $m \leq b_N$, such that every $v \in \mathcal{O}^u$ annihilated by all $\lambda \in D_L$ with $\lambda(z(\cdot)) \in V$ satisfies $\phi(v) \in L$.*

Proof. If the conclusion already holds at N, take $m = 0$. Otherwise choose a bad v: the selected functionals annihilate v but $\phi(v) \notin N$. Let $j \geq 1$ be minimal such that $\pi^j \phi(v) \in N$; the existence of such j follows from the hypothesis on p. Adjoin the row $\pi^{j-1}\phi(v)$ to the relation module. The new module N' costs one power of π, by Lemma 21.12. The annihilator $D_{N'}$ is the additional restriction $\lambda(\pi^{j-1}v) = 0$. It is a proper submodule of D_N: over $\mathcal{O}/p$, the dual maps into A separate points of a finite quotient annihilated by p. This follows at once by writing that quotient as a sum of modules $\mathcal{O}/\pi^j$ with $j \leq v_\pi(p)$ and embedding each cyclic summand into A.

Now $F_{N'} \subset F_N$. If their relative lengths to V were equal, their images in the quotient by V would be equal. For every $\lambda \in D_N$ choose $\lambda' \in D_{N'}$ with $\lambda(z(\cdot)) - \lambda'(z(\cdot)) \in V$. By the defining property of v, $(\lambda - \lambda')(v) = 0$. It would follow that $\lambda(\pi^{j-1}v) = 0$ for every $\lambda \in D_N$, contrary to the properness just proved. Thus $b_{N'} < b_N$. Repeat the construction at N'; there can be at most b_N such steps, since the relative length is a nonnegative integer. Multiplying the one-step ideal inclusions proves the claimed total cost. □

Two caveats about this proposition are relevant to its use with cocycles. The module V must be fixed during the induction (in the application it expresses the permitted surviving cocycles), and the condition on v concerns the annihilator after the enlargement, not only the annihilator at the starting module. The proof works because any bad vector decreases the remaining relative length. It would be incorrect simply to identify the normalized length of the space of cocycles with the valuation of a single congruence element without proving such an inclusion of Fitting ideals.

21.7 Passing from scalar rows to matrix blocks

For the distinguished block one has an additional advantage: the integral span of the representation is the full matrix algebra. This makes the ordinary determinant bound for scalar lattices apply to a matrix of group-algebra elements. The matrix units from Lemma 21.5 explain the mechanism behind the following statement.

Lemma 21.14 (Block determinant from a dual bound)**.** *Let $\Psi : B \to M_n(\mathcal{O})$ be a surjective homomorphism of $\mathcal{O}$-algebras, $L \subset B^u$ a left B-submodule, and $p \neq 0$. Suppose that the space of $w \in ((\mathcal{O}/p)^n)^u$ annihilated by every $t \in L$ under the pairing*

$$(t, w) \longmapsto \sum_{i=1}^{u} \Psi(t_i) w_i$$

has ordinary length at most $b < v_\pi(p)$. There is a square matrix $D \in M_u(B)$ with all rows in L such that its flattened block image has $v_\pi(\det \Psi(D)) \leq b$.

Proof. Let $K \subset \mathcal{O}^{nu}$ be generated by the scalar rows of the horizontal block matrices $(\Psi(t_1), \ldots, \Psi(t_u))$ for $t \in L$. By surjectivity, left multiplication by any matrix unit is induced by an element of B. Consequently an arbitrary scalar row of K can be placed in any desired position of a block row with all other scalar rows zero; addition lets us choose the n scalar rows independently. The annihilator in $(\mathcal{O}/p)^{nu}$ is the ordinary orthogonal complement of K there. Apply Smith normal form to $K' = K + p\mathcal{O}^{nu}$. The length of its annihilator equals the length of $\mathcal{O}^{nu}/K'$, by duality over $\mathcal{O}/p$. Indeed its elementary divisors are powers π^a with $a \leq v_\pi(p)$, and the corresponding kernel of multiplication has length a. Take a basis matrix D' for K' with determinant valuation at most b. Each scalar row of this matrix is congruent modulo p to a row in K. Replace each by that row; the determinant is unchanged modulo p, and because $b < v_\pi(p)$ its valuation is unchanged. Divide the chosen list of nu rows into successive groups of n. The independent-row observation produces a block row of L for each group. Taken together these give D with the required flattened determinant. □

This is `exists_rows_addVal_det_le` in `Lib/General/Algebra/BlockFittingBound.lean`. Note the strict inequality against the modulus. If the determinant's valuation were at least $v_\pi(p)$, a lift of its class modulo p might have larger valuation, or even be zero. The source can choose very deep Taylor–Wiles moduli, and this strict inequality is therefore natural. The lemma requires integral, not merely rational, surjectivity; it is here that the primitive residue-field labels pay for themselves a second time.

21.8 Two applications of Cramer's rule over a family

The determinant of a block relation matrix at a distinguished Galois representation is not yet a Hecke operator acting on every eigensystem. The construction must take place in a common algebra of functions on the eigensystems. Moreover the Hecke algebra need not contain the individual matrix entries of the representations. It may contain their traces and characteristic-polynomial coefficients only. We give the common-algebra argument, which is a useful safeguard against incorrectly treating the representations as a simultaneously diagonalizable family of matrices.

Let $S \subset \prod_{x \in Y} L$ be the commutative algebra of coefficients on a finite family Y, and let $u_x : B \to M_n(L)$ be representations of an algebra B containing a group algebra over $\mathcal{O}$. Suppose the coefficient functions of the characteristic polynomials of the group elements belong to S. We need the following integrality input, which the formal source proves from Amitsur's formula: the coefficients of the characteristic polynomial of any element of the group algebra with coefficients extended to S belong to S, and the assertion remains true for matrices of such group-algebra elements after flattening the blocks. This is `amitsur_charpoly_mem_matrix` in `New/Selmer/RibetCramer.lean`, together with `amitsur_charpoly_mem_of_span` used there. The conclusion is genuinely integral. Newton identities over a field of characteristic zero alone do not give it over an integral ring because of division by integers. In the following argument we use the Amitsur integrality statement as formal-source input.

It follows by Cayley–Hamilton that each square group-algebra matrix D has a polynomial "adjugate" W with coefficients in S, satisfying after application of every u_x

$$u_x(W)u_x(D) = \det(u_x(D))1.$$

Here a group-algebra determinant symbol on the right is not defined; it is the determinant of the flattened image, and the family of its values belongs to S. The polynomial is the signed characteristic-polynomial expression for the adjugate, so the displayed identity also holds at singular matrices. The same reasoning supplies the adjugate of any element of the group algebra separately. Both lift to the algebra extended by S.

Let $\mathfrak{J}$ be an S-submodule of the ambient product of coefficient fields, and form the left ideal spanned by elements r whose trace against every group element belongs to $\mathfrak{J}$. Then the trace of any such left-ideal element against any group-algebra element with coefficients in S belongs to $\mathfrak{J}$. This follows from the cyclic invariance of trace and linearity: for br paired with h, use $\operatorname{tr}(brh) = \operatorname{tr}(rhb)$, and expand hb as a linear combination of group elements. Let V be a left group-algebra-linear map on rows, let w be a row, and let D be square with

$$V(D_i) = r_i + m_i V(w)$$

for elements r_i of this trace ideal and m_i of a designated submodule $\mathfrak{a}$. We write r and m as columns. Define the function β by $\beta_x = \det u_x(V(w))$ and write $F_x = \det u_x(D)$. Both $\beta, F \in S$. The group-algebra expression over S

$$b = F\,1 - wWm$$

satisfies $bV(w) = wWr$ after application of each u_x: indeed the identity for the adjugate gives $wWV(D) = FV(w)$. For a group-algebra element h set $f_x = \operatorname{tr} u_x(bh)$. The coefficient integrality implies $f \in S$. Multiply the trace by β, using the adjugate of $V(w)$ and trace cyclicity. We obtain

$$\beta f = \operatorname{tr}\big(u(bV(w))\operatorname{adj}(u(V(w)))u(h)\big) \in \mathfrak{J}.$$

For clarity, this equality of functions uses the adjugate identity before substituting $bV(w) = wWr$: in each component the product $u(V(w))\operatorname{adj}(u(V(w)))$ is the scalar β_x. Finally, at any component x where $u_x(mh) = 0$ for all $m \in \mathfrak{a}$, the second term of b contributes no trace and

$$f_x = F_x \operatorname{tr} u_x(h).$$

This proves the assertion of `exists_ribet_element_cramer` in `New/Selmer/RibetCramer.lean` given the Amitsur input. The integral cost of the extra normalizing determinant β is visible here: a small valuation for F_x alone does not give a desired integral trace relation. The arithmetic construction spends part of its budget ensuring a bound for the determinant of $V(w)$ and a separate bound for F_x. The distinguished component permits the annihilation by $\mathfrak{a}$, not necessarily the other components; the relation $\beta f \in \mathfrak{J}$, however, is valid simultaneously at all components. This is why the construction can supply an operator that isolates a row in the Hecke algebra.

21.9 A transfer calculation at one auxiliary level

We finish with the first-order operation underlying the local patching statement. It concerns a line stabilizer in a finite general linear group. Although the global statement also requires a comparison of automorphic forms at different levels, the algebra of this local operation can be checked directly.

Lemma 21.15 (Logarithmic transfer)**.** *Let k be a finite field of size q, let $G = \mathrm{GL}_N(k)$, and let H be the stabilizer of a line. Suppose $m \mid q-1$ and fix a generator u of $k^\times$. Define $a : H \to \mathbb{Z}/m$ by the logarithm of the scalar action on the line, taking u to 1. The transfer of a is of the form $d \circ \det$ for a character $d : k^\times \to \mathbb{Z}/m$. It satisfies*

$$N!\, d(u) = N! \quad \text{in } \mathbb{Z}/m.$$

Proof. Choose one representative x for every left coset of H. For $g \in G$ let $\overline{gx}$ denote the representative of gxH, and define $h(g,x) = \overline{gx}^{-1} gx \in H$. Transfer is the homomorphism $g \mapsto \sum_x a(h(g,x))$; multiplicativity follows by composing the permutations of the cosets and adding the corresponding logarithms. On any elementary transvection its value is zero because the transvection has order equal to the characteristic of k, prime to m. Since G is generated by elementary matrices and invertible diagonal matrices, and diagonal entries in different positions are conjugate, the transferred character factors through determinant. At the scalar matrix $u1$, every coset is fixed and its term in the sum is $a(u1) = 1$. There are $(q^N-1)/(q-1) = 1 + q + \cdots + q^{N-1}$ lines, which is congruent to N modulo m. Since the determinant of $u1$ is u^N, we get $Nd(u) = N$. Multiplying by $(N-1)!$ proves the formula (if $N = 1$ it is immediate). □

The source uses precisely this calculation in `exists_line_stabilizer_log` in `New/Patching/OnePrime.lean`. In its local-field formulation the line stabilizer is the reduction modulo the prime of the corresponding block parahoric. Both characters are trivial on the indicated pro-Iwahori group, because there the determinant and the scalar action on the line reduce to 1. The factor $N!$ is deliberately not inverted; in particular the proof works when $\ell \mid N!$. In characteristic zero it would be possible to divide the scalar identity, but doing so would change the integral denominators.

Here is the ring calculation with logarithmic weights. We give it in a form that also applies to noncommutative operator rings; the distinction between operator multiplication and scalar multiplication is retained throughout. Let R be such a ring with identity, $s : R \to R_0$ a ring map killing m, and $\rho : G \to R^\times$ the group action. Let B be an additive map from integer-valued functions on $\mathbb{Z}/m$ to R, with

$$B(f)\rho(p) = B(f(\cdot - a(p))) \qquad (p \in H).$$

The source takes B to be obtained from level-changing operators; for the moment it is only an additive map obeying this translation law. Choose the function $v : \mathbb{Z}/m \to \mathbb{Z}$ of least absolute representatives. Any lift of the residue class would give the same image modulo m. For fixed left coset representatives x as above define

$$T = \sum_x \rho(x), \quad U_1 = \sum_x \rho(x^{-1}), \quad \mathcal{N} = \sum_{g \in G} \rho(g), \quad F = TB(-v)U_1.$$

Assume $\rho(H)$ fixes $B(1)$ on the left, $LB(1) = 0$, $L\rho(g) = L$ for every $g \in G$, and an element Θ satisfies

$$\sum_{p \in H} \rho(p^{-1})\Theta\rho(p) = n\,1.$$

Suppose multiplication by $|G|$ is injective in R. These hypotheses imply

$$\mathcal{N}\Theta B(1) = nTB(1), \qquad L\Theta B(1) = 0. \tag{21.7}$$

For the first identity decompose $\mathcal{N} = T\sum_{p \in H} \rho(p)$ and use $\rho(p^{-1})B(1) = B(1)$ to insert the inverse factor on the right of Θ. The second follows on multiplying the first by L on the left, because the left-hand side then equals $|G|L\Theta B(1)$ and the right-hand side vanishes. Injectivity of multiplication by $|G|$ is required before reduction modulo m, not in R_0.

Proposition 21.16 (First-order weighted operator)**.** *Under these hypotheses, for each $g \in G$ there exists $Z_g \in R$ such that*

$$s(nF(\rho(g)-1)) = s(\mathcal{N} Z_g), \qquad LZ_g = 0.$$

In addition, write $S_H(f) = \sum_{p\in H} f(a(p))\rho(p)$ and $S_G(f) = \sum_{g\in G} f(\mathrm{trf}(a)(g))\rho(g)$, where $\mathrm{trf}(a)$ is transfer. Suppose for some fixed $\eta \in G$ and integer k, $k\,\mathrm{trf}(a)(\eta) = k$ in $\mathbb{Z}/m$. Suppose operators $A_t, A_{t'}, r$ satisfy $A_t, A_{t'}$ commuting with the operators T, U_1 and, for every integer-valued f,

$$B(f(\cdot - k) - f)A_t = S_H(f(\cdot - k) - f)A_{t'},$$
$$S_G(f)A_{t'} = (\rho(\eta) - 1)S_G(f)r.$$

Then

$$2k\, s(FA_t) = 2k\, s(\mathcal{N} r).$$

Proof. Congruent integer weights modulo m give the same result after s, by additivity. Reindex the sum defining $U_1\rho(g)$ using the coset representative of gxH. The translation law for B and $-v(b-a) \equiv -v(b) + v(a) \pmod m$ show that

$$s(F(\rho(g)-1)) = s\left(TB(1)\sum_x v(a(h(g,x)))\rho(x^{-1})\right).$$

Multiplying by n and using (21.7), take

$$Z_g = \Theta B(1)\sum_x v(a(h(g,x)))\rho(x^{-1}).$$

The second part of (21.7) gives $LZ_g = 0$.

For the additional assertion, expand $TS_H(v)U_1$ by its three indices. For each fixed coset representative x, each $g \in G$ occurs exactly once in this expansion with final factor $\rho(x^{-1})$ and middle weight $v(a(h(g,x)))$. Summing the weights over x gives a representative of the transferred logarithm, hence

$$s(TS_H(v)U_1) = s(S_G(v)).$$

Now take $f(b) = -v(b)(v(b)+k)$. Modulo m its difference $f(b-k) - f(b)$ is $2kv(b)$. Apply the first additional operator identity to this f, multiply by T, U_1, and reduce modulo m. Since F uses the weight $-v$, this yields

$$2k\, s(FA_t) = -2k\, s(S_G(v)A_{t'}),$$

where the commutation of A_t and $A_{t'}$ with U_1 is used, as in the source's commuting family. Inserting the second operator identity gives $-2k\, s((\rho(\eta)-1)S_G(v)r)$. But $\rho(\eta)$ multiplies the weighted sum by a translation of its weight through $\mathrm{trf}(a)(\eta)$, by reindexing over G. The difference of these two weights, after multiplication by $-2k$, is congruent to the constant $2k$; this uses $k\,\mathrm{trf}(a)(\eta) = k$. The result is $2k\, s(\mathcal{N} r)$ as claimed. □

This is the algebraic content of `one_prime_first_order` in `New/Patching/OnePrime.lean`. The local logarithm gives one factor $k = N!$, while the factor 2 appears because the polynomial used to extract the first-order term is quadratic. Notice that the argument did not divide by k, by 2, or by the finite index of the line stabilizer. The operator F is built by sums and products in the ring of integral level operators before any reduction modulo m. This integrality is the reason for introducing least absolute representatives rather than an abstract logarithm of an operator over a fraction field.

The passage from this ring identity to the automorphic comparison is not merely an application of group theory. In the source, the operators B are defined from level changes, the two weighted-sum identities arise from Hecke operators acting on a line stabilizer, and the normalizing element satisfies a finite averaging identity. The identifications of the spaces, integrality of the operator on the prescribed lattice of forms, and the multiple-prime construction of compatible diagonal operators are the formal-source inputs `BaseLevel.exists_tie` and `BaseLevel.exists_diagOp_mem_ellIntegralForms` in `New/Patching/OnePrimeTie.lean`, and the Hecke-string assertions cited in Section 20.11. Proposition 21.16 reconstructs the transfer calculation used in these inputs; it does not establish their global automorphic hypotheses.

21.10 Separation of the arithmetic and algebraic budgets

The preceding arguments explain why four superficially different finiteness problems lead to four different costs in the final exponent of the integral projection.

1. Residual irreducibility is formulated as full integral span, proved with a tame pair. It enables a bound on a single row test space to be turned into a bound for a flattened determinant (Lemma 21.14). It incurs no extra linear algebra denominator.

2. The balanced row condition concerns the global dimension of local restrictions. The mirror estimate turns an apparently many-variable isotropy constraint into a scalar defect. This is why the estimate measures the length of a kernel of a scalar, not the rank of the entire dual Selmer module.

3. In nonzero Tate twists the odd class bounds and the reciprocity laws serve different purposes. The law makes the cost of adjoining a relation in triangular order uniform, even without bounding all even classes. Lemma 21.12 turns each such addition into a multiplicative loss of at most ℓ^H. The scalar character differences and the detector cost have already been included in $y^{22}h$; one should not multiply this cost by the number of all possible residual characters.

4. A flattened determinant is converted into a trace function in the common eigenvalue algebra by two instances of Cramer's rule. The determinant of the normalizer is a further explicit cost. The local first-order group-ring calculation, in turn, can be executed without inverting the finite group orders; where an integer factor is lost it is recorded in the denominator budget.

These calculations do not make the arithmetic hypotheses in the previous chapter dispensable. In particular the existence and local comparison of attached Galois representations, the calculation of crystalline local lengths in the balance proposition, the global duality and odd-class bounds, the simultaneous twist selection for all rows, and the level-comparison assertions are specialized inputs whose exact boundaries were stated there. The estimates assembled from them control the valuations uniformly in the coefficient field, the individual twist and the deep finite Taylor–Wiles modulus. This uniformity, rather than an equality for one specially chosen automorphic eigensystem, is what permits the arguments of Sections 20.12 and 20.13 to yield the fit family.

Chapter 22

Arithmetic of the congruence argument: cyclotomic classes, descent and simultaneous twists

We reconstruct a further portion of the arithmetic underlying the congruence bound. In particular, we prove the cyclotomic estimates that permit the determinant argument to use an explicit integer controlling only odd classes. A comparable bound on the whole class group of a cyclotomic tower would be an unjustified strengthening. We also develop finite coordinates for cocycles, the construction of pencils of bounded cardinality, and selected steps in the transport from Hecke algebras to Galois representations and back. The deepest geometric compatibility and integral comparison assertions are separated from the finite arguments they imply.

In the first half of this chapter let N be the conductor of a cyclotomic field containing the CM field E, and let e be a positive integer. This N is the conductor, not the varying split prime of the unitary chain. Write ℓ for the coefficient prime, $m_0 = \ell^\nu$ for a torsion modulus, and O for the integers in a finite extension of $\mathbb{Q}_\ell$. The ordinary O-module length is denoted by λ_O and the normalized length by $\lambda = \lambda_O/\lambda_O(O/\ell)$. We use arithmetic Frobenius for reciprocity, so the cyclotomic character ϵ at an unramified Frobenius is the residue cardinality. Let ψ be a Dirichlet character modulo N and write

$$\chi(g) = \epsilon(g)^e \psi(g \bmod N),$$

where the second factor means evaluation via the action on the N-th roots of unity. At the odd components we impose $\chi(c) = -1$, or equivalently $\psi(-1) = -(-1)^e$.

22.1 The integer in the odd-class budget

The first task is to find an integer whose divisors control the relevant Stickelberger expressions at all coefficient primes. The coefficient ring is allowed to vary, and the summation of the local losses uses the factorization of an integer. The source implementation is `IsStickelbergerBound` and `exists_isStickelbergerBound` in `Lib/Cyclotomic/Voronoi/StickelbergerInteger.lean`. Here is a construction with the same bound, using a determinant of an integral convolution matrix.

Let $B_j(X)$ be the Bernoulli polynomials, given by

$$\frac{te^{Xt}}{e^t - 1} = \sum_{j \geq 0} B_j(X) \frac{t^j}{j!}.$$

For every $j \geq 0$, $(j+1)!B_j(X)$ has integral coefficients. Indeed the recurrence $B_0 = 1$, $B_j = -\frac{1}{j+1}\sum_{i<j}\binom{j+1}{i}B_i$ proves integrality of $(j+1)!B_j$ by induction, and the formula $B_j(X) = \sum_{i=0}^{j}\binom{j}{i}B_i X^{j-i}$ proves the polynomial assertion. We use

$$P(X) = (e+3)!N^{e+2}B_{e+2}(X/N), \quad D = (e+3)!(e+2)N, \qquad b(X) = P'(X).$$

The generating function gives the identities

$$P(X+N) - P(X) = DX^{e+1}, \qquad b(X) = DN^e B_{e+1}(X/N). \tag{22.1}$$

Lemma 22.1 (A twisted Voronoi congruence)**.** *Let $f = Ns$ with $s > 0$ and let $d > 0$ satisfy $d \equiv 1 \pmod N$ and $(d, s) = 1$. For a Dirichlet character ψ modulo N in a characteristic-zero domain O, extended by zero on nonunits, set*

$$S_f = \sum_{0 \leq u < f} \psi(u)u^{e+1}, \qquad T_f(d) = \sum_{0 \leq u < f} \psi(u)u^e \left\lfloor \frac{du}{f} \right\rfloor.$$

Then

$$(d^{e+1} - 1)S_f \equiv (e+1)d^e f T_f(d) \pmod{f^2 O}. \tag{22.2}$$

Moreover, if $B(\psi) = \sum_{0 \leq u < N} \psi(u)b(u)$, then

$$DS_f \equiv fB(\psi) \pmod{f^2 O}. \tag{22.3}$$

Proof. Multiplication by d permutes residues modulo f and preserves ψ. Writing $v = du - f\lfloor du/f \rfloor$ and expanding gives $v^{e+1} \equiv d^{e+1}u^{e+1} - (e+1)d^e u^e f\lfloor du/f \rfloor \bmod f^2$. The weighted sum of the left side is S_f. For the second assertion, group the terms by their residues modulo N and apply (22.1). For a residue u, the s successive differences of P telescope to $P(u+f) - P(u)$, congruent to $fb(u)$ modulo f^2 because P is integral. □

Put $A = (\mathbb{Z}/N)^\times$ and write $[v]_N \in [0, N)$ for the representative of v. Define

$$h(v) = 2b([v]_N) + (-1)^e \mathbf{1}_{v=1} + \mathbf{1}_{v=-1}, \qquad C_{u,v} = h(v/u).$$

For every character ψ of A, the vector $(\psi(v))_{v \in A}$ is an eigenvector of C with eigenvalue $c(\psi) = 2B(\psi) + (-1)^e + \psi(-1)$. Over $\mathbb{C}$ the character vectors form a basis. The formula for the Dirichlet value at a negative integer is

$$L(-e, \psi) = -\frac{N^e}{e+1}\sum_{u=0}^{N-1} \psi(u)B_{e+1}(u/N).$$

This follows by expressing L as a sum of Hurwitz zeta functions, each of which has value $-B_{e+1}(u/N)/(e+1)$ at $s = -e$. By the functional equation of a primitive Dirichlet L-function and its nonvanishing Euler product at $1 + e$, $L(-e, \psi) = 0$ if and only if $\psi(-1) = (-1)^e$; for an imprimitive character the additional Euler factors at $-e$ do not vanish because $p^e > 1$. In particular, $B(\psi)$ vanishes in this parity and is nonzero in the other. Thus $c(\psi)$ equals $2(-1)^e$ or $2B(\psi)$ respectively. It follows that $D_0 = |\det C| > 0$.

Proposition 22.2 (One bounded integer for the Stickelberger elements)**.** *There is an integer X_e with $1 \leq X_e \leq 2^{(N+e+3)^6}$ such that:*

1. *For each prime ℓ some $d > 0$, $d \equiv 1 \bmod N$, prime to ℓ, has $(d^e - 1)(d^{e+1} - 1) \mid X_e$.*

2. *If $d \equiv 1 \bmod N$ and $\psi(-1) = -(-1)^e$, then in any characteristic-zero domain containing the values of ψ there exists an element w independent of s such that for $(s,d) = 1$*

$$(d^{e+1} - 1)X_e \equiv 2wT_{Ns}(d) \pmod{sO}. \tag{22.4}$$

Proof. Take $X_e = D_0 \prod_{d \in \{1+N, 1+2N\}} (d^e - 1)(d^{e+1} - 1)$. The two choices of d are coprime. The adjugate identity for C applied to the character vector over O shows that $c(\psi)$ divides $\det C$: use the coordinate indexed by 1, which equals 1. In the desired parity $c(\psi) = 2B(\psi)$, so $X_e = 2B(\psi)u_\psi$ for some $u_\psi \in O$. Combining the two congruences of the lemma gives

$$f(d^{e+1} - 1)B(\psi) \equiv f(e+1)d^e DT_f(d) \pmod{f^2 O}.$$

Cancel f in the domain, multiply by $2u_\psi$, and take $w = u_\psi(e+1)d^e D$.

For the size, let $s_0 = N + e + 3 \geq 5$. The Fourier series for B_{e+1} on $[0,1]$ gives $|B_{e+1}(x)| \leq 2(e+1)!\zeta(e+1)/(2\pi)^{e+1} \leq (e+1)!$. Thus with $M_0 = DN^e(e+1)!$ one has $|h(v)| \leq 4M_0$. Hadamard's inequality yields $D_0 \leq (4M_0\sqrt{|A|})^{|A|}$. By $n! \leq n^n$,

$$\log_2 M_0 \leq 4s_0 \log_2 s_0, \quad \log_2 D_0 \leq |A|(2 + (4s_0 + \tfrac{1}{2}) \log_2 s_0).$$

The four other factors contribute at most $2(2e+1)\log_2(1+2N) \leq 8s_0^2$. Since $\log_2 s_0 \leq s_0$, the total is bounded by $5s_0^3 + 10s_0^2 \leq s_0^6$. □

For $x \in \mathbb{Q}/\mathbb{Z}$ write $\{x\} \in [0,1)$ for its fractional representative. The function in the Gauss-sum source is $f_d(x) = \lfloor d\{x\} \rfloor - \lfloor d\{-x\} \rfloor$.

Lemma 22.3 (Removing multiples of the coefficient prime)**.** *Let $f = N\ell^\nu$, $d \equiv 1 \bmod N$, $(d, \ell) = 1$, and $\psi(-1) = -(-1)^e$. In $O/\ell^\nu O$ the element*

$$A_\nu(d) = \sum_{a \in (\mathbb{Z}/f)^\times} f_d(a/f)[a]_f^e \psi(a)$$

divides $(d^{e+1} - 1)X_e$.

Proof. For $\nu = 0$ this is automatic. Put $f_j = N\ell^j$ and $K = (d^{e+1} - 1)X_e$. The sum of the unit terms of T_{f_ν} equals

$$U_\nu = T_{f_\nu}(d) - \ell^e \psi(\ell) T_{f_{\nu-1}}(d).$$

This follows by the substitution $a = \ell b$, using the zero-on-nonunits convention even when $\ell \mid N$. As $e \geq 1$, (22.4) at the two moduli gives $2wU_\nu \equiv (1 - \ell^e\psi(\ell))K \bmod \ell^\nu O$. Substituting the negative of a unit changes its weight by a minus sign modulo ℓ^ν, by parity and binomial expansion. Thus $A_\nu(d) \equiv 2U_\nu \bmod \ell^\nu$. The factor $1 - \ell^e\psi(\ell)$ is a unit, proving the assertion (`sum_floorDiff_dvd` in `Lib/Cyclotomic/Kolyvagin/ComponentBound.lean`). □

22.2 The Jacobi family and its distribution relation

Let p be an auxiliary rational prime and let B be a place of $\overline{\mathbb{Q}}$ above p. Put $n = p - 1$, choose the character of $\mathbb{F}_p^\times$ with values in μ_n whose reduction at B is the identity, and choose a primitive p-th root ζ_p. For $x \in \mathbb{Q}/\mathbb{Z}$ annihilated by n, set

$$G_x = \sum_{a \in \mathbb{F}_p^\times} \chi_x(a)\zeta_p^{[a]_p}, \qquad \overline{\chi_x(a)} = a^{-n\{x\}} \text{ in } \mathbb{F}_p^\times.$$

Here the bar denotes reduction, not complex conjugation. The characters are unique by lifting prime-to-p roots of unity. In particular $G_0 = -1$. The finite geometric-series identity gives, for $x \neq 0$,

$$G_x G_{-x} = \chi_x(-1)p, \qquad |G_x|_{\mathbb{C}} = \sqrt{p}. \tag{22.5}$$

For the first equality, write the double sum in terms of the ratio a/b: the sum over b is $p-1$ if the ratio is -1 and -1 otherwise. The second follows by complex conjugation.

Lemma 22.4 (Tame Gauss valuation)**.** *Normalize $v_B(p) = 1$. Then $v_B(G_x) = \{x\}$, and G_x is a unit at every finite place off p.*

Proof. In $\mathbb{Q}_p(\zeta_p, \mu_n)$ the extension by μ_n is unramified and $\pi = \zeta_p - 1$ has valuation $1/n$. For $x \neq 0$, the reduction of G_x modulo π is zero, as it is the sum of a nontrivial character of $\mathbb{F}_p^\times$. Similarly G_{-x} is a nonunit. If $G_x = \pi^k u$ with u a unit, (22.5) gives $0 < k < n$. For $v \in \mathbb{F}_p^\times$ let σ_v fix the n-th roots and send ζ_p to $\zeta_p^{[v]_p}$. It acts trivially on residue fields and $\sigma_v(\pi)/\pi$ reduces to v. By substituting $a \mapsto va$, $\sigma_v(G_x) = \chi_x(v)^{-1} G_x$. Reducing the quotient gives $v^k = v^{n\{x\}}$, so $k = n\{x\}$ using its range. The case $x = 0$ is immediate. Integrality of the Gauss sums and (22.5) show that they are units off p. □

For $d > 0$ put $j_d(x) = G_x^d / G_{dx}$ and $J_d(x) = j_d(x)/j_d(-x)$. If σ acts on the n-th roots by multiplication by a and on the exponents of ζ_p by multiplication by v, then $\sigma G_x = \chi_{ax}(v)^{-1} G_{ax}$. The scalar cancels in j_d, hence $\sigma J_d(x) = J_d(ax)$.

Lemma 22.5 (Jacobi-family identities)**.** *Suppose $M \mid p-1$ and $Mx = 0$. Then $J_d(x)$ belongs to $\mathbb{Q}(\mu_M)^\times$, is a unit off p, has valuation $f_d(x)$ at B, and has absolute value 1 at every complex embedding. If $q \mid p-1$ is a prime and $qb = 0$, then at every place above q,*

$$J_d(x + b) \equiv J_d(x) \pmod{\mathfrak{m}_q}.$$

Finally, if $r \mid p-1$ and $(r, d) = 1$, then

$$\frac{\prod_{b=0}^{r-1} J_d(x + b/r)}{J_d(rx)} \quad \text{is a root of unity.} \tag{22.6}$$

Proof. Galois covariance shows the field assertion. The valuation follows from Lemma 22.4 and $d\{x\} - \{dx\} = \lfloor d\{x\} \rfloor$. The unit statement follows from the same lemma, and $|j_d(x)| = |j_d(-x)|$ follows from (22.5), also when $dx = 0$.

At $q \neq p$ all roots of unity of q-power order reduce to 1. Thus the terms of the sums defining G_{x+b} and G_x have the same reductions; the denominators of J_d are units. For the distribution relation note that $\sum_{b=0}^{r-1} \{x + b/r\} = (r-1)/2 + \{rx\}$. Since $(r, d) = 1$, the same formula applied to $d(x + b/r)$ gives

$$\sum_{b=0}^{r-1} \lfloor d\{x + b/r\} \rfloor = (d-1)(r-1)/2 + \lfloor d\{rx\} \rfloor.$$

Replace x by $-x$ and subtract to cancel the constant, giving $\sum_b f_d(x + b/r) = f_d(rx)$. This holds after every automorphism of the cyclotomic field, so the quotient in (22.6) is an algebraic unit, of absolute value one at every embedding. It is a root of unity by Kronecker's theorem (its powers have minimal polynomials of bounded degree with integral coefficients bounded by the corresponding binomial coefficients, so they range over a finite set). These arguments reconstruct `exists_jacobi_family` in `Lib/Cyclotomic/Stickelberger/GaussSums.lean`. □

22.3 The derivative cocycles from the Jacobi family

It is not legitimate to take arbitrary Kummer classes of (22.6) and ignore the root-of-unity ambiguity. The normalization that eliminates it is the following. Fix $m_0 = \ell^\nu > 1$ and $\gamma \in G_{\mathbb{Q}}$. In this subsection write the multiplicative group of $\overline{\mathbb{Q}}^\times$ additively, so multiplication by m_0 means taking an m_0-th power. Choose y_x satisfying $m_0 y_x = J_d(x)$, and put

$$F(g,x) = \gamma y_x - g\gamma g^{-1} y_x \in \mu_{m_0}.$$

It belongs to μ_{m_0} because the two automorphisms act identically on the abelian field containing $J_d(x)$. It is independent of the chosen root y_x, because γ and $g\gamma g^{-1}$ agree on all roots of unity. Identify μ_{m_0} with $\mathbb{Z}/m_0$ via a fixed primitive root, so its Galois action is multiplication by ϵ.

Write a_g for the exponent of the action of g on the relevant roots of unity. We have

$$F(gh,x) = F(g,x) + \epsilon(g) F(h, a_g^{-1} x), \tag{22.7}$$

$$F(\gamma,x) = 0, \qquad \sum_{b=0}^{r-1} F(g, x + b/r) = F(g, rx) \quad ((r,d)=1). \tag{22.8}$$

In the first identity insert $g\gamma g^{-1} y_x$ into the difference and use Galois covariance of J_d. The element $g^{-1} y_x$ differs from $y_{a_g^{-1}x}$ by an m_0-th root of unity, which cancels. For the last equality, $\sum_b y_{x+b/r} - y_{rx}$ is itself a root of unity by (22.6); applying $\gamma - g\gamma g^{-1}$ annihilates it.

Put $M = N m_0$ and let $\mathcal{P}$ be a finite set of rational primes $q \equiv 1 \bmod dM$. Choose $p \equiv 1 \bmod M \prod_{q\in\mathcal{P}} q$ different from ℓ and the primes of $\mathcal{P}$. For each $q \in \mathcal{P}$ choose a surjective homomorphism $L_q : \mathbb{F}_q^\times \to \mathbb{Z}/m_0$. For $S \subseteq \mathcal{P}$ define

$$\begin{aligned}\kappa_S(g) = \sum_{a \in (\mathbb{Z}/M)^\times} [a]_M^{e} \psi(a) \sum_{(u_q) \in \prod_{q\in S} \mathbb{F}_q^\times} \left(\prod_{q\in S} L_q(u_q) \right) \\ \cdot F\left(g, \frac{a}{M} + \sum_{q\in S} \frac{u_q}{q} \right) \quad \text{in } O/m_0.\end{aligned} \tag{22.9}$$

The finite logarithmic derivative here is a group-algebra operation, not a derivative of an analytic function.

Lemma 22.6 (Independence of logarithm shifts)**.** *Replacing any $L_q(u_q)$ in (22.9) by $L_q(u_q) - d_q$ with constant $d_q \in \mathbb{Z}/m_0$ does not change the sum.*

Proof. Induct on $|S|$. To change the logarithm at q while allowing arbitrary shifts for the others, compute the difference: the constant d_q times the sum with L_q omitted. For $x = a/M + \sum_{r\in S\setminus\{q\}} u_r/r$, (22.8) gives

$$\sum_{u_q \in \mathbb{F}_q^\times} F(g, x + u_q/q) = F(g, qx) - F(g,x).$$

As $q \equiv 1 \bmod M$, multiplication by q preserves a and permutes the u_r. Reindexing changes their logarithms only by constants $L_r(q)$, and the induction hypothesis shows the sums agree. Repeat to handle all shifts. □

Proposition 22.7 (Finite derivative cocycles)**.** *The maps (22.9) are locally constant cocycles with coefficient character $\epsilon^{e+1}\psi$ and vanish at γ. They vanish on inertia over $\mathbb{Q}(\mu_N)$ away*

from ℓ, p and S. For $q \in \mathcal{P} \setminus S$ and an arithmetic Frobenius φ at a prime over q, there is a tame inertia element τ over the same prime with

$$\kappa_{S\cup\{q\}}(\tau) = \kappa_S(\varphi). \tag{22.10}$$

If $\kappa_\varnothing$ vanishes on inertia at the place B above p, then

$$(\epsilon(\gamma)^{e+1}\psi(\gamma) - 1) \sum_{a\in(\mathbb{Z}/M)^\times} f_d(a/M)[a]_M^e\psi(a) = 0 \quad \text{in } O/m_0. \tag{22.11}$$

Proof. The sums are finite and factor through a finite extension. In (22.7), multiplication of the arguments by a_g^{-1} induces permutations of the units modulo the denominators. Reindexing changes the weight $[a]_M^e\psi(a)$ by the factor $\epsilon(g)^e\psi(g)$, and changes the logarithms by constants, which cancel by Lemma 22.6. This proves the cocycle relation. Vanishing at γ follows from the definition. Outside the permitted set the fields defined by the denominators of the x are unramified over $\mathbb{Q}(\mu_N)$, and the Jacobi values are units. Taking roots of order prime to the residue characteristic of such a unit gives unramified local extensions (lift the separable roots of its residue polynomial by Hensel's lemma). This applies also to every conjugate of every y_x, proving the inertia vanishing.

For (22.10) fix $q \notin S$ and $x = a/M + \sum_{r\in S} u_r/r$. Select τ in inertia at q inducing $v \in \mathbb{F}_q^\times$ on the q-th roots of unity with $L_q(v) = -1$. This is possible by the total tame ramification of the q-th cyclotomic field at q. It fixes the roots of unity of orders prime to q. Choose all roots $y_{x+u/q}$, $0 \le u < q$, to have the same reduction as y_x at the selected place above q. This is possible by Lemma 22.5 and injectivity on reduction of μ_{m_0}. It follows that

$$\tau^{-1}y_{x+u/q} = y_{x+v^{-1}u/q}.$$

In fact the two sides are roots of the same J_d and have the same reduction. Every element here is a local unit. The definition of F and triviality of the residue action of inertia give

$$F(\tau, x + u/q) \equiv \gamma(y_{x+u/q} - y_{x+v^{-1}u/q}) \pmod{\mathfrak{m}_q},$$

where reduction is still written additively in the unit group. There is no requirement that γ itself preserve the chosen prime.

Take integer lifts of the logarithms and reindex the second sum. The coefficient differences equal 1 modulo m_0 and sum to zero as integers. By Galois covariance, $\gamma(J_d(x + u/q) - J_d(x))$ is a principal unit at q, so after subtracting y_x from every term we can replace these coefficient differences by 1. Thus

$$\sum_{u\in\mathbb{F}_q^\times} L_q(u)F(\tau, x + u/q) \equiv \gamma\left(\sum_{u=0}^{q-1} y_{x+u/q} - qy_x\right) \pmod{\mathfrak{m}_q}. \tag{22.12}$$

The difference $\eta = \sum_{u=0}^{q-1} y_{x+u/q} - \varphi y_x$ is a root of unity: take its m_0-th power and use (22.6) and $\varphi J_d(x) = J_d(qx)$. On reduction the sum equals qy_x, which is the reduction of φy_x. Thus η is a root of unity reducing to 1 and has q-power order. Its conjugate $\gamma\eta$ also reduces to 1. Hence the right side of (22.12) is congruent to $\gamma\varphi y_x - q\gamma y_x$. This is the reduction of $F(\varphi, qx)$, as one may use φy_x as the root in its root-independent definition. Both sides of the desired identity lie in μ_{m_0}, on which reduction is injective. Thus the congruence is equality. Summing over a and the u_r, $r \in S$, yields the sum of $F(\varphi, qx)$ weighted by their logarithms. Reindex the arguments using $q \equiv 1 \bmod M$ and apply Lemma 22.6 to obtain (22.10).

For (22.11), choose τ in tame inertia at B whose action on an m_0-th root of p is translation by a generator ξ of μ_{m_0}. The Jacobi values of denominator M lie in a field unramified at

p, and $v_B(J_d(x)) = f_d(x)$. Dividing by a suitable power of p^{1/m_0} leaves a root of a unit, unramified at p, so

$$(\tau - 1)y_x = f_d(x)\xi, \qquad (\tau - 1)\gamma y_x = f_d(a_\gamma x)\xi.$$

As τ fixes μ_{m_0}, the definition of F gives $F(\tau, x) = (\epsilon(\gamma)f_d(x) - f_d(a_\gamma x))\xi$. Reindex the unit sum by $a \mapsto a_\gamma a$. Multiplication by the unit $\epsilon(\gamma)^e\psi(\gamma)$ then gives (22.11) whenever $\kappa_\varnothing(\tau) = 0$. □

This constructs the derivatives and proves their norm relation and bottom formula as used in `exists_kolyvagin_cocycles` in `Lib/Cyclotomic/Kolyvagin/ComponentBound.lean`. Notice why the helper parameter γ is necessary: roots of unity arising from the distribution relation need not be divisible in the finite cyclotomic field, but their contribution cancels in the commutator formula for F.

22.4 The odd component, local vanishing and reciprocity

We now fix the Dirichlet character of the desired parity, and use the component called `component` in `Lib/Cyclotomic/Kolyvagin/Components.lean`. It consists of locally constant cocycles $x : G_\mathbb{Q} \to O/\ell^\nu$ for $\chi^{-1} = (\epsilon^e\psi)^{-1}$ such that:

1. $x(c) = 0$;
2. x vanishes on inertia over $\mathbb{Q}(\mu_N)$ at all finite primes away from ℓ;
3. x vanishes on the decomposition groups over $\mathbb{Q}(\mu_N, \mu_{\ell^\infty})$ above ℓ.

We write $\mathfrak{C}_\nu(\psi, e)$ for this module. Conjugation of a cocycle of character λ obeys

$$x(hgh^{-1}) = \lambda(h)x(g) + (1 - \lambda(g))x(h). \tag{22.13}$$

This follows by applying the cocycle equation twice and using $x(h^{-1}) = -\lambda(h)^{-1}x(h)$. In particular, the restriction to a subgroup acting trivially on the coefficients is additive and transforms by λ on conjugation.

Let $d \equiv 1 \bmod N$ be prime to ℓ, and set $F = \mathbb{Q}(\mu_{Nm_0})$. For a place of $\overline{\mathbb{Q}}$ above ℓ, choose δ in inertia with $\epsilon(\delta) = d$ and fixing μ_N. This choice is possible by the surjectivity of local inertia on the pro-ℓ cyclotomic extension and the fact that it acts trivially on the prime-to-ℓ part of μ_N; the congruence $d \equiv 1 \bmod N$ gives triviality on the remaining part. If $g \in G_F$ belongs to the decomposition group at this place, then its commutator with δ lies in the decomposition group over $\mathbb{Q}(\mu_N, \mu_{\ell^\infty})$. Thus $x(\delta g\delta^{-1}) = x(g)$, since the restrictive character on G_F is trivial. By (22.13),

$$(d^e - 1)x(g) = 0. \tag{22.14}$$

The source allows an additional factor of 2, and we keep it in the next statements. This argument explains how the local condition at ℓ controls the ordinary, not merely the cyclotomic-kernel, Artin value.

Lemma 22.8 (Artin-law calculation)**.** *Let $x \in \mathfrak{C}_\nu(\psi, e)$, let $q \equiv 1 \bmod Nm_0$ be a rational prime distinct from ℓ, and let φ be an arithmetic Frobenius at a place B above q. For $u \in (\mathbb{Z}/Nm_0)^\times$ choose $s_u \in G_\mathbb{Q}$ acting on the Nm_0-th roots by u. Suppose $y \in F^\times$ has the following property at every place over a prime other than q, ℓ: either x is zero on the*

decomposition group, or y is a unit times an m_0-th power of an element of a local unramified extension. Then

$$\left(\sum_u v_B(s_u y)[u]^e\psi(u)\right)\, 2(d^e-1)x(\varphi)=0. \tag{22.15}$$

Here $v_B(q)=1$ and the valuations in the sum are integers.

Proof. The homomorphism $g\mapsto 2(d^e-1)x(g)$ on G_F is additive into the finite module O/m_0. By (22.14) its restriction to local decomposition groups above ℓ is zero. Since F is totally complex, its finite quotients are evaluated by global Artin reciprocity without real-place terms. Its Artin value on the principal idele y is zero. Off ℓ it is unramified and depends only on the local valuation. The valuation at a place of a local unramified extension has integer values on that extension for the base normalization, so divisibility by m_0 of the valuation of y in the second alternative makes the local contribution zero. The first alternative also gives zero. At q, full splitting in F makes the valuations integers and the primes correspond to the $s_u^{-1}B$. Frobenius at $s_u^{-1}B$ is conjugate to $s_u^{-1}\varphi s_u$, and (22.13) shows that its value of x is $[u]^e\psi(u)x(\varphi)$. Substitution into the sum of local Artin values proves (22.15). Equivalently one can apply every additive character of O/m_0 to the Artin sum and use separation of points by the characters of a finite abelian group. This is the calculation `artin_law_component` in `Lib/Cyclotomic/Kolyvagin/ComponentBound.lean`. □

Proposition 22.9 (Odd-component annihilation)**.** *Every member of $\mathfrak{C}_\nu(\psi,e)$ is killed by $4X_e^2$, for X_e of Proposition 22.2.*

Proof. Choose d as in that proposition. If $4X_e^2x\neq 0$, then $2X_e^2x$ does not vanish identically on G_F. Indeed, if it did, that cocycle would descend to the abelian quotient $\mathrm{Gal}(F/\mathbb{Q})$, since its coefficient character modulo m_0 also descends. It is normalized to zero at c, whose character is -1, and a cocycle with this property on an abelian group is killed by 2 by comparing its values at cg and gc. This would contradict the assumption. By Chebotarev we can find $q\equiv 1 \bmod Nm_0$ different from ℓ and a Frobenius φ over q with $2X_e^2x(\varphi)\neq 0$. Here and below Chebotarev is applied to a finite quotient through which the relevant cocycles and cyclotomic exponents factor, with conjugation allowed in the selection of the place. Apply the Jacobi family with $p=q$ and $y=J_d(1/(Nm_0))$. Its off-q unit property gives the hypothesis of Lemma 22.8; its valuations by Galois covariance are $f_d(u/(Nm_0))$. Thus $A_\nu(d)\,2(d^e-1)x(\varphi)=0$, using the notation of Lemma 22.3. That lemma gives

$$2(d^{e+1}-1)(d^e-1)X_e x(\varphi)=0.$$

But $(d^e-1)(d^{e+1}-1)$ divides X_e, so this implies $2X_e^2x(\varphi)=0$, a contradiction. □

At large primes a substantially stronger conclusion about the *number of detecting coordinates* is needed. Merely applying Proposition 22.9 to all cocycles cannot give that conclusion: the number of coordinates might grow with ℓ. For this reason the finite derivative classes are accompanied by a Kummer-duality calculation.

Lemma 22.10 (Reciprocity modulo the coefficient prime)**.** *Assume $\phi(N\ell)$ is a unit in O. Let $x\in\mathfrak{C}_\nu(\psi,e)$ be killed by ℓ. Let T be a finite set of rational primes, all $\equiv 1 \bmod N\ell$, and let κ be a locally constant cocycle for $\epsilon^{e+1}\psi$ with values in a finite O-module killed by ℓ, unramified on inertia over $\mathbb{Q}(\mu_N)$ away from ℓ and T. Suppose $q\in T$, x vanishes on the decomposition groups above $T\setminus\{q\}$, and*

$$2(d^e-1)x(\varphi)\neq 0$$

at an arithmetic Frobenius φ over q. Then κ vanishes on inertia at the chosen place B over q.

Proof. Set $F_1 = \mathbb{Q}(\mu_{N\ell})$. On G_{F_1}, κ is additive, since its character acts trivially on ℓ-torsion. Fix an inertia element τ at B and let $v = \kappa(\tau)$. The cyclic O-module Ov is a finite vector space over $\mathbb{F}_\ell$. Choose an $\mathbb{F}_\ell$-basis b_i of this vector space, and let L_i be extensions of its coordinate homomorphisms to the target ℓ-torsion module of κ. Choose a primitive ℓ-th root used to identify $\mathbb{F}_\ell$ with μ_ℓ. Kummer theory over F_1 associates to $L_i\kappa|_{G_{F_1}}$ an element $y_i \in F_1^\times/(F_1^\times)^\ell$, since F_1 contains μ_ℓ. If κ vanishes on inertia at a prime off ℓ, T, an ℓ-th root of y_i is unramified there. Thus the unit-times-ℓ-th-power condition in Lemma 22.8 holds for y_i. At primes of $T \setminus \{q\}$ the first alternative holds by hypothesis. The Artin-law calculation applies with modulus ℓ for the additive homomorphism $2(d^e-1)x|_{G_{F_1}}$; the local equation (22.14) remains true in its ℓ-torsion target O/ℓ^ν.

Let $a_{i,u} = v_B(s_u y_i)$, $u \in (\mathbb{Z}/N\ell)^\times$, and let $j \in \mathbb{F}_\ell$ be the action of τ on a fixed ℓ-th root of q. All the following equalities of coordinates are in $\mathbb{F}_\ell$. By local Kummer theory at $q \neq \ell$, the value on τ of the Kummer character of $s_u y_i$ is $j a_{i,u}$. Indeed after taking out the valuation factor the remaining root is a root of a unit and is fixed by inertia. On the other hand Galois action on Kummer characters gives

$$j a_{i,u} = L_i(u^{-e}\psi(u)^{-1} v).$$

To check this identity, conjugate τ by s_u^{-1} and then apply s_u to the roots of unity. The cocycle relation (22.13) multiplies $\kappa(\tau)$ by $\epsilon(s_u)^{-e-1}\psi(u)^{-1}$, and the action on μ_ℓ multiplies the result by u. This proves the identity and shows why the two powers in the dual character differ by one.

By the Artin-law calculation, $r_i = \sum_u [u]^e \psi(u) a_{i,u}$ annihilates the nonzero element $2(d^e-1)x(\varphi)$. Therefore each r_i belongs to the maximal ideal $\mathfrak{m}_O$. Using the preceding coordinate identity and summing over u, we obtain in the cyclic O-module Ov

$$\sum_i j r_i b_i = \sum_u [u]^e \psi(u) \sum_i j a_{i,u} b_i = \phi(N\ell) v.$$

Here integer lifts of j and the coordinates can be used since the module is killed by ℓ. The left side belongs to $\mathfrak{m}_O Ov$. Since $\phi(N\ell)$ is a unit, $v \in \mathfrak{m}_O Ov$. Nakayama's lemma gives $Ov = 0$. This holds for every τ, proving the assertion. This argument is integral even when the ℓ-torsion module is not a vector space over the residue field of O: the intermediate coordinates are deliberately taken over $\mathbb{F}_\ell$. □

22.5 The Kolyvagin–Rubin detection argument

We first isolate the finiteness and Chebotarev facts used by the induction. For a continuous action on a finite module V, continuous cocycles unramified outside a fixed finite set of primes form a finite set. In fact a cocycle is equivalent to a lift of the action to the semidirect product $V \rtimes \mathrm{Aut}(V)$. The degree of the finite Galois extension cut out by the lift is bounded by $|V||\mathrm{Aut}(V)|$, and the extensions in question are unramified outside the fixed set (including the set of ramification of the action). Hermite–Minkowski gives finiteness of extensions of bounded degree unramified off that set; one can obtain the necessary discriminant bound either from the finiteness of local extensions of bounded degree or the usual local different bound in terms of degree. A finite Galois group has only finitely many homomorphisms to the fixed target.

Another useful elementary fact is that a cocycle for a scalar character λ on $G_\mathbb{Q}$, zero at an element γ where $\lambda(\gamma) - 1$ is a unit, cannot factor nontrivially through an abelian quotient on which the action is defined. To prove this, compare the cocycle at γg and $g\gamma$. In particular, if a nonzero such cocycle is zero on $G_{\mathbb{Q}(\mu_L)}$ for a modulus L divisible by the conductor of its

coefficient character, we would get a contradiction. Given two such cocycles, with possibly different actions or different normalizing elements, we may consequently find an element of $G_{\mathbb{Q}(\mu_L)}$ at which both are nonzero. Indeed the two zero subgroups of the restricted cocycles are proper, and the union of two proper subgroups cannot cover a group. Chebotarev realizes the selected pair of nonzero values at an arithmetic Frobenius at a prime $\equiv 1 \bmod L$. We will use this with arbitrarily large but finite additional factors in L. If the first cocycle is normalized at c with odd character, then -2 is a unit as soon as $\ell > 2$.

Theorem 22.11 (Bounded number of detecting points for odd components)**.** *If $\ell > N+e+3$, there are at most $v_\ell(X_e)+1$ elements of $G_{\mathbb{Q}}$ such that evaluation at these elements is injective on $\mathfrak{C}_\nu(\psi, e)$.*

Proof. Write $b = v_\ell(X_e)+1$. We may use b positions, repetitions allowed. Suppose no such list detects all the elements. Given a list of at most b elements, there is then a nonzero member of the component vanishing at all of them. By multiplying it by a suitable power of ℓ, we can arrange it is killed by ℓ and remains nonzero. Denote the set of nonzero ℓ-torsion classes of the component by Z. This set is finite by the preceding argument.

Since $\ell > N+e+3$, the Chinese remainder theorem gives $d \equiv 1 \bmod N$ with reduction modulo ℓ generating $\mathbb{F}_\ell^\times$. In particular $\ell \nmid d$ and d^e-1, $d^{e+1}-1$, 2, and $\phi(N\ell)$ are all units in O. Choose an element γ of inertia over ℓ fixing μ_N with cyclotomic value d, as in (22.14). The scalar $\epsilon(\gamma)^{e+1}\psi(\gamma) - 1 = d^{e+1}-1$ is a unit.

For a finite set T of rational primes $\equiv 1 \bmod N\ell$ not dividing ℓ, let $Y(T)$ denote the module of locally constant cocycles with values in O/ℓ^b, of character $\epsilon^{e+1}\psi$, vanishing on inertia over $\mathbb{Q}(\mu_N)$ off ℓ, T. Write $Y = Y(\varnothing)$. Say that T is good if each member of $Y(T)$ killed by ℓ belongs to Y. The empty set is good, and all the sets below are finite.

First we claim that whenever T is good and $|T| \le b$, some $z \in Z$ has the property that every $q \equiv 1 \bmod N\ell$ for which z does not vanish on a decomposition group over q lies outside T and makes $T \cup \{q\}$ good. Choose an arithmetic Frobenius over each prime of T. By our supposition we can take $z \in Z$ vanishing at them. Since those primes split in $\mathbb{Q}(\mu_{N\ell})$, all inertia is in $G_{\mathbb{Q}(\mu_N)}$ and the coefficient action on their decomposition groups is trivial on ℓ-torsion. The inertia condition on z gives vanishing at each such entire decomposition group, and (22.13) gives vanishing also at its conjugates (the action on ℓ-torsion is trivial there). Hence q as in the claim is outside T and $z(\varphi_q) \neq 0$ at a Frobenius over q. Lemma 22.10 applies to any ℓ-torsion member of $Y(T \cup \{q\})$: it forces the cocycle to vanish on inertia at q. Conjugation makes the same true at all places over q. This member belongs to $Y(T)$ and hence to Y, proving the claim.

The set Y is finite; its ramification over $\mathbb{Q}$ is confined to the primes dividing $N\ell$. For every pair

$$(y, z) \in \{y \in Y : y \neq 0,\ y(\gamma) = 0\} \times Z$$

choose a prime $q(y,z) \equiv 1 \bmod dN\ell^b$ and an arithmetic Frobenius $\varphi(y,z)$ above it at which both values $y(\varphi(y,z))$ and $z(\varphi(y,z))$ are nonzero. This is possible by the two-cocycle Chebotarev argument above, since y is normalized at γ and z at complex conjugation. Let $\mathcal{P}$ be the resulting finite set of primes. It is not important how large $\mathcal{P}$ is. Only the cardinality of the final detecting list is being bounded.

Apply the claim with $T = \varnothing$, and select $p \equiv 1 \bmod N\ell^b \prod_{q\in\mathcal{P}} q$, outside $\mathcal{P}$, at which that z has a nonzero value at Frobenius. The singleton $\{p\}$ is then good. Construct the cocycles κ_S, $S \subseteq \mathcal{P}$, using Proposition 22.7 at modulus $m_0 = \ell^b$, with the chosen γ, d, p.

We prove by descending induction on $i = b, b-1, \dots, 0$ the assertion

$$|S| \le i, \quad S \subseteq \mathcal{P}, \quad \{p\} \cup S \text{ good} \quad \Longrightarrow \quad \ell^i \kappa_S \in Y. \tag{22.16}$$

At $i = b$ the left-hand cocycle is zero. Suppose the assertion known at $i+1$, and take S at stage i. If $y = \ell^{i+1}\kappa_S$ were nonzero, it would belong to Y by induction, and be normalized to zero at γ. By the claim with $T = \{p\} \cup S$, of size at most $i+1 \le b$, choose $z \in Z$ as above. The prime $q = q(y,z)$ belongs to $\mathcal{P}$ and its selected Frobenius has both y and z nonzero. The claim shows that $q \notin T$ and $T \cup \{q\}$ is good. By (22.16) at $i+1$, $\ell^{i+1}\kappa_{S\cup\{q\}}$ is unramified at q. But the derivative relation (22.10) shows that its value on a tame generator is $y(\varphi(y,z)) \neq 0$, a contradiction. Therefore $\ell^{i+1}\kappa_S = 0$. Since $\ell^i\kappa_S$ is a member of $Y(T)$ killed by ℓ and T is good, it belongs to Y. This finishes the induction.

At the bottom we obtain $\kappa_\varnothing \in Y$, so it is unramified at p. Equation (22.11) shows $A_b(d) = 0$ in O/ℓ^b because $\epsilon(\gamma)^{e+1}\psi(\gamma) - 1$ is a unit. Lemma 22.3 now implies that $(d^{e+1}-1)X_e = 0$ in O/ℓ^b, which is impossible by $v_\ell(X_e) = b-1$. The contradiction proves the theorem. □

This is `exists_detecting_points_component` in `Lib/Cyclotomic/Kolyvagin/ComponentBound.lean`. The descending induction uses the good-set condition only for ℓ-torsion members of $Y(T)$. This matters: a claim that every member of a module at an auxiliary level is already unramified could be false, and would not be needed. Likewise the finite set $\mathcal{P}$ is not the eventual detecting set; its size may be very large. The proof only uses it to arrange implications simultaneously over the finite coefficient ring.

22.6 Uniform detection for the small primes

The preceding component argument works when ℓ is large compared with the conductor and exponents. For small ℓ we use a deliberately crude bound, independent of the torsion depth. We spell out how it can be made independent of the actual coefficient ring too.

Let $E \subseteq \mathbb{Q}(\mu_N)$ and set $K_0 = \mathbb{Q}(\mu_{N\ell^2})$. Write $D_0 = [K_0 : \mathbb{Q}] \le N\ell^2$ in this subsection. The group of continuous homomorphisms $G_{K_0} \to \mathbb{Z}/\ell$ unramified off ℓ factors through the maximal elementary abelian ℓ-quotient of the abelian Galois group unramified off ℓ. By ray class theory, its dimension over $\mathbb{F}_\ell$ is at most

$$\log_\ell h_{K_0} + D_0 + 2\,\#\{w : w \mid \ell\}. \tag{22.17}$$

Indeed the maximal pro-ℓ quotient of the ray class exact sequence is a quotient of an extension of the ℓ-part of the ordinary class group by the product of the pro-ℓ unit groups at the places above ℓ. The elementary generator count for the class group is at most $\log_\ell h_{K_0}$. For a finite extension $K_{0,w}/\mathbb{Q}_\ell$, the pro-ℓ unit group has at most $[K_{0,w} : \mathbb{Q}_\ell] + 2$ generators: a group of finite index is via the logarithm a $\mathbb{Z}_\ell$-lattice of rank the local degree, and the torsion of the unit group is cyclic; allowing two instead of one covers all small-prime conventions. Equivalently use the classification of finitely generated abelian pro-ℓ groups. Summing the local degrees proves (22.17).

We give one explicit estimate on the class-number term. The cyclotomic discriminant formula gives $|D_{K_0}| \le (N\ell^2)^{D_0}$. Minkowski's bound provides an integral ideal in each class of norm at most $\sqrt{|D_{K_0}|}$, since the remaining factors in the usual bound are at most 1. If L is any number field of degree d and $R \ge 1$, the number of ideals in L of norm at most R is bounded by $(\lfloor R \rfloor + 1)^{d+1}$: at each rational prime there are at most d prime ideals over it, and the coefficient of norm n is no larger than the d-fold divisor-counting function $d_d(n)$, by inserting the possible exponents with their residue degrees. The sum of $d_d(n)$ counts tuples of positive integers each at most R, hence is at most $(\lfloor R \rfloor + 1)^d$; the extra power is harmless. Consequently

$$\log_2 h_{K_0} \le (D_0 + 1)\bigl(1 + \tfrac{1}{2} D_0 \log_2(N\ell^2)\bigr).$$

For $y \geq 3$, $N \leq y$ and $\ell \leq y^2$, one has $D_0 \leq y^5$. The last display, plus $3D_0$ from the local terms of (22.17), is bounded by $y^{20} - D_0$, using $\log_2 y \leq y$ and $y \geq 3$. For example the terms are bounded by $(y^5+1)(1+(5/2)y^6)+3y^5 \leq y^{20}-y^5$ for $y \geq 3$.

Choose a list of s Galois elements of G_{K_0} whose images generate the exponent-ℓ quotient just described, and add a representative of every coset of G_{K_0} in G_E. There are at most D_0 such representatives. The resulting list has at most y^{20} elements at the stated sizes.

Lemma 22.12 (A uniform detecting set for Tate cocycles)**.** *The list just constructed detects all locally constant cocycles $z : G_E \to O/\ell^\nu$ for the character ϵ^{-e} unramified outside ℓ, for all integers $e \geq 0$ and all depths $\nu \geq 1$. The list can be chosen independently of e, ν, O.*

Proof. Suppose a cocycle vanishes on the list. If it were nonzero, some power of ℓ times it would be nonzero and killed by ℓ. Replace it by that multiple. On G_{K_0} the coefficient action on ℓ-torsion is trivial. Its restriction is then an additive homomorphism unramified outside ℓ, with values in a finite elementary abelian ℓ-group. Since all its scalar coordinates in $\mathbb{Z}/\ell$ kill the generator list, the restriction is zero. For any $g \in G_E$ write $g = hu$ with $h \in G_{K_0}$ and u a chosen representative. The cocycle equation gives $z(g) = z(h) + \epsilon(h)^{-e} z(u) = 0$, contradicting nonzero. This proves the assertion. □

The exact uniformity of this proof is useful also in the triangular construction, where no bound on the even class group is invoked. This is the arithmetic behind `exists_detecting_points_cocycles` in `Lib/Cyclotomic/OddSha/OddShaBound.lean`; use of the stronger ℓ^2 cyclotomic field in the detecting list is harmless at every depth, including $\ell = 2$. For the count we used elementary class field theory and the Minkowski bound, not an Iwasawa-theoretic estimate.

22.7 Returning from Dirichlet components to the CM field

We now combine the preceding results without assuming that the CM field has degree prime to ℓ. Let $G = G_{\mathbb{Q}}$, $H = G_E$, and $K = G_{\mathbb{Q}(\mu_N)} \subset H$. All are viewed using an embedding into $\overline{\mathbb{Q}}$; H is normal in G because E is contained in the cyclotomic field. For $e > 0$, let x be an element of the odd group `shaOdd` of the source, i.e. a locally constant cocycle on H for ϵ^{-e} with values in O/ℓ^ν, unramified off ℓ, satisfying the prescribed vanishing on cyclotomic-kernel decomposition groups at ℓ and

$$x(cgc^{-1}) = -x(g) \quad (g \in H).$$

Here the local condition is read on decomposition elements whose full cyclotomic character is trivial. At $e = 0$ the source instead uses the class-group convention discussed below.

For $u, g \in G$ write $b_u(g) = u(cgc^{-1}g^{-1})u^{-1} \in K$ (a commutator has trivial action on all roots of unity). Let $\mathcal{S}$ be any set of representatives of $K\backslash G$ with 1 representing the identity coset. Define for each Dirichlet character ψ modulo N

$$z_\psi(g) = \sum_{u \in \mathcal{S}} \epsilon(u)^e \psi(u) x(b_u(g)).$$

There is no division by $|G/K|$ here. This will keep the argument valid also at small primes.

Lemma 22.13 (Integral decomposition by commutators)**.** *Each z_ψ is a locally constant cocycle of the character $(\epsilon^e\psi)^{-1}$, zero at c, and satisfying the component local conditions of Section 22.4. For $g \in H$ one has*

$$4\phi(N)x(g) = \sum_{\psi} \big(z_\psi(cgc^{-1}) - z_\psi(g)\big). \tag{22.18}$$

In the sum on the right only characters with $\epsilon^e(c)\psi(c) = -1$ can contribute.

Proof. If u is replaced by hu for $h \in K$, the value of $x(b_u(g))$ is multiplied by $\epsilon(h)^{-e}$. This is (22.13), since $b_u(g)$ has trivial cyclotomic character. The weight $\epsilon(u)^e\psi(u)$ is multiplied by the reciprocal factor, proving independence of representatives. Next

$$b_u(gh) = b_u(g)b_{ug}(h).$$

The action of each factor on coefficients is trivial, so substituting into the sum, changing representatives in the second part, gives the cocycle equation of the claimed character. We have $b_u(c) = 1$, so the value at c is zero. For g in the specified inertia or decomposition groups of the component, $b_u(g)$ is a product of conjugates of g and of g^{-1} under elements of G. These factors belong to H (indeed $g \in K$) and are in the corresponding local vanishing sets of the cocycle x: away from ℓ they are inertia elements, while above ℓ they have trivial full cyclotomic action. Thus all the z_ψ meet the local conditions. Finite sums and continuity give local constancy.

Assume O contains a primitive root of unity of order the exponent of $(\mathbb{Z}/N)^\times$, as in the source. The character orthogonality identity $\sum_\psi \psi(u) = \phi(N)\mathbf{1}_{u=1}$ holds integrally over O. Consequently

$$\sum_\psi z_\psi(g) = \phi(N)x(cgc^{-1}g^{-1}) = -2\phi(N)x(g) \qquad (g \in H).$$

The last equality uses the odd condition and the cocycle formula at the product. Apply it also to cgc^{-1}, subtract, and obtain (22.18). As $z_\psi(c) = 0$, the cocycle formula gives $z_\psi(cgc^{-1}) = \epsilon(c)^{-e}\psi(c)^{-1}z_\psi(g)$, so terms of even parity vanish exactly. The sum-and-difference identity, without division by 2, proves the assertion also over coefficient rings above 2. □

For each odd character, Proposition 22.9 says that $4X_e^2 z_\psi = 0$. Lemma 22.13 therefore gives

$$Y_e x = 0, \qquad Y_e = 4\phi(N) \cdot 4X_e^2. \tag{22.19}$$

This annihilates the entire odd group. We next control its number of coordinates. Fix $t \geq 1$, $e < t$, and $y_0 = N + t + 3$. For $\ell \leq y_0^2$, Lemma 22.12 gives at most y_0^{20} detecting elements. For $\ell > y_0^2$, Theorem 22.11 gives a list $\gamma_{\psi,i}$ of at most $v_\ell(X_e) + 1$ elements detecting z_ψ at each odd parity. Evaluate x on the elements $b_u(\gamma_{\psi,i})$ for the $\phi(N)$ choices of u, the $\phi(N)$ or fewer choices of ψ, and the indicated list of i. If all such evaluations are zero, every odd z_ψ is zero, so $4\phi(N)x = 0$. For the large primes under consideration $4\phi(N)$ is a unit, hence this implies $x = 0$. The number of evaluations is bounded by

$$\phi(N)^2(v_\ell(X_e) + 1) \leq y_0^2(y_0^6 + 1) \leq y_0^{20}.$$

Thus a bound on the number of detecting coordinates holds at every prime.

A finite-length submodule of q copies of O/ℓ^ν killed by Y_e has normalized length at most $q\, v_\ell(Y_e)$, by the elementary divisor description of the kernel of multiplication by Y_e. Applying this to the injective evaluation map and (22.19) gives $\lambda(\mathrm{shaOdd}) \leq y_0^{20} v_\ell(Y_e)$ for $e > 0$. At $e = 0$ the source's local definition gives a subgroup of the homomorphisms from the ordinary ideal class group of E to O/ℓ^ν: the cocycles for the trivial character are homomorphisms and the local conditions impose unramifiedness also at ℓ. A cyclic decomposition of the finite ideal class group shows that this subgroup has normalized length at most $v_\ell(h_E)$.

Theorem 22.14 (The global integer bounding odd classes)**.** *There exists a positive integer*

$$X = h_E \prod_{1 \leq e < t} Y_e^{y_0^{20}} \leq h_E\, 2^{4(N+t+3)^{27}}$$

such that at every prime ℓ, every torsion depth and every $0 \le e < t$, the normalized length of the odd module used in the congruence argument is at most $v_\ell(X)$. The coefficient ring is assumed to contain enough roots of unity of order the exponent of $(\mathbb{Z}/N)^\times$.

Proof. The length assertion was verified above. By Proposition 22.2,

$$\log_2 Y_e \le 4 + \log_2 \phi(N) + 2(N+e+3)^6 \le 3y_0^6.$$

There are at most $t-1 < y_0$ factors, so $\log_2(X/h_E) \le 3y_0^{27} \le 4y_0^{27}$. This proves the statement with the same exponent as the source. □

Theorem 22.14 reconstructs the bound `exists_length_shaOdd_le` in `Lib/Cyclotomic/OddSha/OddShaBound.lean`, as used in Proposition 20.16. Its proof uses the usual classical results of cyclotomic and global class field theory, Chebotarev, the elementary Kummer correspondence, and Dirichlet functional equations; it does not invoke the cyclotomic main conjecture or a bound on the even class groups. The relatively large exponent is a convenient consequence of the difference between small-prime class-number counting and large-prime Kolyvagin detection. The argument deserves emphasis: a coarse annihilator plus a count of detecting coordinates is enough for the sum of normalized lengths. Counting the derivative primes or estimating the full ramification degree of the Euler system is neither needed nor compatible with this uniform budget.

22.8 Uniform cardinality of Selmer pencils

We return to the notation of Chapter 20: E is the CM field of the datum, and m is the number of arrows. This section proves the uniform finiteness part of the simultaneous-twist theorem, conditional on the arithmetic mirror criterion (the difficult global reciprocity step isolated in Proposition 20.9). The earlier chapter explains the tilting lemma; we now give the precise induction on the torsion depth and the subsequent specialization that uses it.

Fix a finite collection of rows R_i of ranks n_i over a DVR O, balanced at depth M, and write π for a uniformizer. We retain the source notation β_i for their ordinary length allowances. For a divisor a of ℓ^M, for a finite set of pencil places J, and for each row let

$$X_i(J,a) \subset Z^1(G_E, (O/(a))^{n_i})$$

denote the cocycle module with the old crystalline relations, vanishing at the normalizer γ_i, unramified outside the union of the old set S and $J \cup J^c$, and with the ties at all places of J. The ties use values at tame generators τ_P, conjugate inverses $(\tau_P^c)^{-1}$, Frobenii φ_P, and $(\varphi_P^c)^{-1}$, as in `IsPencil` and `tieRelations` of `New/Selmer/Corepresentable.lean`. For a place where the slope of row i is nonzero, the row representation on both Frobenius elements is trivial modulo ℓ^M; all represented auxiliary places are unramified for the original row.

The following observations concern the hypotheses of the induction. At a new place of residue characteristic different from ℓ, a torsion cocycle for an unramified representation is a homomorphism on inertia. Its restriction factors through the pro-ℓ tame quotient; indeed wild inertia is pro-p for the residue prime $p \neq \ell$. Thus requiring its value at the chosen generator of this quotient to be zero makes it unramified. At conjugate places we use the other generator as well. This proves that, in a relaxed module with an extra pair of ramification places permitted, the additional conditions of zero tame evaluations are exactly the old unramified conditions. If zeros at all auxiliary tame generators are imposed, the module is contained in the fixed module with only the old ramification set. That latter module has finite length for finite coefficients by the finiteness argument in Section 22.5. Its length is independent of the depth

N of the auxiliary characters. All linear relations, including the cocycle equations and the old relations, commute with the injections

$$O/\pi^{k-1} \longrightarrow O/\pi^{k}, \qquad x \longmapsto \pi x.$$

In particular this map identifies the smaller-depth cocycle module with the subset of the larger module whose values all lie in $\pi O/\pi^{k}$.

Proposition 22.15 (Construction of a bounded pencil from the mirror criterion)**.** *Suppose the mirror criterion of Proposition 20.9 holds for the row modules just described, with the one-prime compatibility and conjugate-inverse conventions of* `IsPencil`*. Assume that at every depth N one can select a helper place with its Frobenius in the required neighborhood and with all its slopes zero, and that the residue field has more than $4n_i + 2$ elements for all i. Then there exists an integer J_0 independent of N such that, for each N, there is a pencil J of at most J_0 places with*

$$\lambda_O(X_i(J,a)) \le \beta_i \qquad \text{for all } i.$$

This holds at every fixed divisor a of ℓ^M.

Proof. Induct on $k = v_\pi(a)$. For $k = 0$ the modules vanish, and the helper place gives a pencil of size one for every N. Suppose a uniform bound J_1 obtained at $k-1$. For each i let S_i^k be the cocycle module modulo π^k with original ramification set S, old relations, and normalization but without new ties. Put

$$C_k = J_1 + \sum_i (\lambda_O(S_i^k) + 2n_i k J_1).$$

By induction, for every N there is a pencil with $|J| \le J_1$ and all the required bounds at $k-1$. Evaluate the module $X_i(J,\pi^k)$ on both tame generators at every new place. Its kernel lies in S_i^k, and its image has length at most $2n_i k|J|$. Hence this pencil satisfies

$$|J| + \sum_i \lambda_O(X_i(J,\pi^k)) \le C_k. \tag{22.20}$$

This is a bound independent of N, even though the two pencils for different N may have completely different places.

Among pencils satisfying the $k-1$ bounds and (22.20), choose one minimizing the sum of the lengths at depth k. Such a minimum exists as the set is nonempty and the sum takes nonnegative integer values. Suppose some i has $\lambda_O(X_i(J,\pi^k)) > \beta_i$. There must be a cocycle in this module with at least one unit coordinate: otherwise all cocycles have values in $\pi(O/\pi^k)^{n_i}$, and the injective identification with $X_i(J,\pi^{k-1})$ above contradicts the strict inequality.

Apply the contrapositive of the mirror criterion at depth k. It supplies a place $P \notin J$, admissible for the pencil for every choice of its new slope at row i and slopes zero at the others, with the following failure of inclusion. For $a = \pi^k$ let Y_i be the module relaxed at P, P^c, with the old ties unchanged; let $\Phi_0 : Y_i \to (O/(a))^{2n_i}$ evaluate the two tame generators, and let Φ_1 evaluate the two Frobenius elements. Then

$$X_i(J,a) = Y_i \cap \ker \Phi_0, \qquad \Phi_1(X_i(J,a)) \not\subseteq \Phi_0(Y_i).$$

For a new slope $t = s^{-1}$, the additional ties are the kernel of $s\Phi_0 - \Phi_1$. Use Lemma 20.8 both at $a = \pi^k$ and at $a = \pi^{k-1}$. The prohibited residue classes of s at the two depths together have cardinality at most $4n_i$. Since zero must also be excluded, the size assumption on the

residue field permits a unit choice of s. At both depths the length of the new tied module is at most the length of

$$Y_i \cap \ker \Phi_0 \cap \Phi_1^{-1}(\Phi_0 Y_i).$$

At the larger depth the failed inclusion makes this a proper submodule of $X_i(J, \pi^k)$ and so the ordinary length drops by at least one. At the smaller depth it is a submodule of $X_i(J, \pi^{k-1})$, preserving the required bound. For the other rows choose slope zero. Then the new ties simply impose the original unramified condition at the new places and hence do not increase either length.

The new pencil increases $|J|$ by one and decreases the sum of lengths at depth k by at least one. Thus (22.20) remains valid and the smaller-depth bounds remain valid, contrary to minimality. It follows that every row has length at most β_i at the larger depth. In particular the new uniform bound on the cardinality can be taken to be C_k. Induction proves the proposition. □

This is the length argument of `exists_isPencil` in `New/Selmer/TwistTheorem.lean`. It uses the mirror criterion, rather than an unconditional assertion that adjoining a generic prime always lowers a Selmer rank. It also uses the helper- and enlarging-prime compatibility for all choices of the new slope; if the new place depended on the slope, Lemma 20.8 would not have fixed linear maps to which to apply it. The existence of the helper itself follows from the auxiliary anticyclotomic construction of Section 21.4, applied with no prior nonhelper primes. The compatibility of an enlarging prime with the mirrors is part of the stated criterion and remains a specialized global reciprocity input in this reconstruction.

22.9 Dropping ties along a line of approximate characters

We make the intermediate approximate-character calculation explicit. At a pencil of auxiliary depth N, denote the helper by P_0 and the characters of its places by $h_P : G_E \to \mathbb{Z}/\ell^N$. If $y = (y_P)$ is an actual tuple of ℓ^N-th roots, the associated anticyclotomic character is

$$\xi_y(g) = \prod_{P \in J} y_P^{[h_P(g)]_{\ell^N}}.$$

The character takes the value y_P on both chosen tame generators at P and the conjugate place. At each nonhelper Frobenius, on the two sides, it takes the helper coordinate y_{P_0}. This follows directly by multiplying the character values: the matrix of tame evaluations $h_Q(\tau_P)$ is the identity; for nonhelper P the column $(h_Q(\varphi_P))_{Q \in J}$ is the indicator of $Q = P_0$. Anticyclotomicity changes conjugation into inversion, giving the same values at the conjugate inverses. This is `IsPencil.pointChar_idOrConjInv` in `New/Selmer/TiltingLemma.lean`.

The special polynomial point for a single row i has

$$y_{P_0} = u, \qquad y_P = 1 + t_{P,i}(u-1) \quad (P \neq P_0), \qquad u^{\ell^k} = 1. \tag{22.21}$$

In general only the helper coordinate is a true root of unity. By the elementary power calculation of (20.11), each $y_P^{\ell^{k+M}}$ is 1 modulo ℓ^M. Thus for $N \geq k + M$ the pointwise twisted matrices give a representation on O'/ℓ^M over any sufficiently large coefficient ring O' in which these points are defined. All original relations in the cocycle problem still make sense after evaluation.

Lemma 22.16 (An affine form of the tame relation)**.** *Let R be an unramified representation at a new place P with residue cardinality q. Let B be the coefficient ring modulo ℓ^M over O', and twist R by the polynomial point (22.21). Assume $N \geq k + 2M$, $q \equiv 1 \bmod \ell^N$. For every*

cocycle c on the twisted representation T, at a selected tame generator τ and Frobenius φ of this auxiliary place, one has

$$(T(\varphi)-1)c(\tau) = (T(\tau)-1)c(\varphi). \tag{22.22}$$

The same formula holds on the conjugate-inverse side.

Proof. The cocycle corresponds to the affine action $A(g) : v \mapsto T(g)v + c(g)$ on B^{n_i}. On inertia, its linear part has ℓ-power order, and wild inertia maps trivially to the affine group: all its linear images and translational images lie in ℓ-groups, whereas wild inertia is pro-p for a different prime p. At the tame generator the linear part is the scalar y_P. Its ℓ^{k+M}-th power is 1 by the approximate-root calculation, so $A(\tau)^{\ell^{k+M}}$ is a translation. The additive group B^{n_i} is killed by ℓ^M, hence $A(\tau)^{\ell^{k+2M}} = 1$. The tame local Galois relation is $\varphi\tau\varphi^{-1} = \tau^q$ in the tame quotient. It therefore gives commuting affine maps, since $q \equiv 1 \bmod \ell^N$. Comparing their values at zero gives (22.22). At the conjugate place the inverses obey the same commuting relation, which is all that is needed in the calculation. □

Proposition 22.17 (Length cost of dropping ties)**.** *At the row-specific point (22.21) with $N \geq k + 2M$, remove the ties at a subset I of the pencil places. If the helper belongs to I, assume $u = 1$. For $P \in I$, choose matrices $R(w_{P,b})$ agreeing modulo ℓ^M with $R(\varphi_{P,b})$ at the two sides b (with inverse conjugation as above). The increase in normalized length, measured relative to the module with these ties, is bounded by*

$$\sum_{P\in I} \sum_{b} v_\ell \det(uR(w_{P,b}) - 1).$$

Proof. For P nonhelper the difference defining the tie is $d_c = c(\tau) - t_{P,i}c(\varphi)$. The scalar of $T(\tau)$ is $1 + t_{P,i}(u-1)$ and $T(\varphi) = uR(\varphi)$ modulo ℓ^M. If $t_{P,i} = 0$, (22.22) immediately gives $(T(\varphi)-1)d_c = 0$. If $t_{P,i} \neq 0$, the pencil condition gives $R(\varphi) = 1$ modulo ℓ^M, and (22.22) gives $(u-1)d_c = 0$ again. Thus in both cases d_c takes values in the kernel of $uR(\varphi) - 1$ on B^{n_i}. If $P = P_0$, the assumption $u = 1$ makes all the parameters 1, and the same identity holds for the expression of the original tie: when the slope is nonzero $R(\varphi) = 1$, while when it is zero we use unramifiedness at τ. Consequently d_c takes values in the stated kernel in all cases. Imposing its vanishing decreases length by at most the length of that kernel. Replacing $R(\varphi)$ by $R(w_{P,b})$ makes no difference modulo ℓ^M. By Smith normal form the normalized length of the kernel is at most the valuation of the determinant (with an infinite bound allowed when the determinant vanishes). Summing these estimates over P, b proves the proposition. □

The last result is the precise estimate `IsPencil.length_le_of_drop_ties`. We used an affine action because the action on coefficients alone does not justify commuting the cocycle values at the tame generator and at Frobenius. The extra M in the requirement $N \geq k + 2M$ kills the translations in the affine group.

22.10 The two specializations in the common-twist theorem

We now give the quantitative deduction of Proposition 20.14, conditional on the mirror criterion, the existence of balanced row lattices, and the integral measure assertion for the value families, as in the compact source. The relevant extra inputs are specified rather than replaced by a qualitative twisting theorem.

First note the uniform finite-coordinate argument for the cocycle modules. Choose an open normal subgroup U on which the reductions of the row matrices, and the finitely many

characters defining the pencil, are trivial. Let $V \subset U$ be the intersection of the kernels of all additive homomorphisms from U to $\mathbb{Z}/\ell^M$ that are unramified at the places of its field outside the stipulated finite set, and take its normal core in G_E. There are finitely many such homomorphisms, by the class-field or bounded-degree argument of Section 22.5. The resulting open normal subgroup V is independent of the point y and the coefficient ring used for its evaluation. Indeed on U the action is trivial on the torsion module, so the restriction of any cocycle is additive there; the additive characters into the subgroup of $\mathbb{Q}/\mathbb{Z}$ killed by ℓ^M detect every nonzero value, showing the restriction kills V. As $c(gv) = c(g) + T(g)c(v)$, it follows that all such cocycles are determined by their values on the finite set G_E/V.

Write these coordinates as a finite column over O'/ℓ^M. The cocycle relations and the desired inertia and other linear conditions form a relation module in a finite free polynomial module over $O[Y_P]$, using least nonnegative representatives of $h_P(g)$ in the exponents. Include ℓ^M times the free module. This presentation is valid for every tuple for which the twisted matrix action is defined modulo the torsion modulus, including (22.21). In particular it is noetherian even if the old linear conditions are expressed as an arbitrary set of finitely supported relations. Lemma 20.11 supplies a determinant polynomial which realizes the normalized length at one selected point exactly and bounds the length from above at all other specializations. The coefficient ring is allowed to enlarge finitely after the selected point has been found: normalized lengths of finite modules and valuations of determinant polynomials are preserved by finite extension of DVRs at a point over the smaller ring. This follows either from faithful flatness or from multiplying the elementary-divisor lengths and the normalization factor by the same ramification index. This explains the compatibility of `exists_corepresenting_coords` with `exists_sharp_fitting_poly` in the source.

Write β_i for the ordinary balanced length bound and

$$B_i \geq \frac{\beta_i}{\lambda_O(O/\ell)} + \sum_b v_\ell \det(R_i(z_b) - 1), \qquad z_1 = z, \quad z_2 = (z^c)^{-1},$$

the assumed real upper bounds in Proposition 20.14. Refine the allowed neighborhood of z so that the matrices of the helper Frobenius on both sides agree with these matrices modulo ℓ^M. For the finite family of Euler-factor characters μ_j, choose the neighborhood also so that $v_\ell(\mu_j(\varphi_{P_0}) - 1) \leq c$; this is possible because the analogous inequality holds at z, and this is an open condition when the upper bound is finite. The bound J_0 in Proposition 22.15 now depends on that neighborhood but not on N.

Let $\eta > 0$ be the desired final error, and choose $\delta > 0$ with $(8J_0 + 2)\delta < \eta$.

First selection. Choose k by Lemma 20.13 in one variable, with the finite row and value bounds described below, and with an upper bound on the total degree of the monic polynomials to be avoided. These upper bounds depend only on the row ranks, the number of fixed value families and J_0, not on the number or degrees of the cocycle relations. The requisite k can consequently be fixed before selecting the pencil. Next take $N \geq k + 2M$ sufficiently large for the second selection below. Choose the pencil with at most J_0 places and the tied length bounds of Proposition 22.15.

For row i take a sharp Fitting polynomial $G_i(Y)$ at $\mathbf{1}$, retaining all ties except the helper tie. At $\mathbf{1}$ the normalized length without the helper tie is at most B_i: apply Proposition 22.17 with $u = 1$ and $w_{P_0,b} = z_b$. Therefore $v_\ell(G_i(\mathbf{1})) \leq B_i$. Restrict G_i to the polynomial line (22.21) and include this as a polynomial in the variable u. Its value at $u = 1$ gives the needed lower bound for root selection.

For every value family V include also the polynomial $V(\{P_0\}, \xi_u)$, where ξ_u is the character of the helper alone at value u. This polynomial exists and is integral by the finite-measure

formula (20.13). At $u = 1$, the insertion inequality (20.14) and the helper neighborhood imply

$$v_\ell V(\{P_0\}, 1) \le v_\ell V(\varnothing, 1) + 8c.$$

This gives a finite initial bound independent of the pencil. Finally include the monic polynomials obtained from $\det(uR_i(\varphi_{P,b}) - 1)$ for every nonhelper P and both sides b. Their leading coefficients $\det R_i(\varphi_{P,b})$ are units, so one can divide by them. Include the monic linear factors for the possible bad Euler values $\mu_j(\varphi_{P,b})u^{\pm 1}$: up to multiplication by a unit at a root u, their distance from 1 is the distance from a specified unit. The number and degrees of all these monic polynomials have the required uniform bound.

Select a common ℓ^k-th root of unity u so that every one-variable row and value polynomial has valuation increased by at most δ relative to the upper bounds at 1, and every stated monic factor has valuation at most δ. This is Lemma 20.13, with $m = 1$ and initial point 1.

At the row-specific point $y_i(u)$, the sharp polynomial bounds the length with all nonhelper ties by $B_i + \delta$. Dropping those ties by Proposition 22.17 costs at most $2J_0\delta$. Hence the unrestricted normalized module at this point has length at most $B_i + (2J_0 + 1)\delta$.

For each value family use the actual root tuple $y^V_{P_0} = u$ and $y^V_P = 1$ at the other places. Every nonhelper Frobenius evaluates this character at u or the prescribed conjugate-inverse value; the Euler-factor polynomials selected above imply the required valuation bound δ. Successively adding the nonhelper places to the declared modulus and applying the insertion inequality (20.14) gives

$$v_\ell V(J, \xi_{y^V}) \le v_\ell V(\varnothing, 1) + 8c + (8J_0 + 1)\delta.$$

The conductor of the character at this intermediate point can be smaller than the declared auxiliary modulus; this is allowed in the value-family distribution law. No equality between y^V and any of the row-specific points is claimed.

Second selection. For the unrestricted module of row i, take its sharp Fitting polynomial $F_i(Y)$ at the row-specific point $y_i(u)$; this point is a legitimate torsion specialization since $N \ge k + 2M$. Then

$$v_\ell F_i(y_i(u)) \le B_i + (2J_0 + 1)\delta.$$

For each value family take its polynomial on the full declared set J, evaluated at y^V as just bounded. Every coordinate of every one of these tuples satisfies the initial power condition $\ell \mid (y-1)^{\ell^k}$: this holds for u and for $1+t(u-1)$ by binomial expansion in residue characteristic ℓ. All initial valuation bounds, and the bound $|J| \le J_0$, are independent of the original degree of the polynomials and of N.

Apply Lemma 20.13 with the first m powers of one tuple of roots, the common power parameter $a = \ell^k$, and loss δ. Include the monic polynomials $Y^j - 1$ for $1 \le j \le m$. Thus a threshold for N was available before the pencil was chosen, since it depends only on the listed bounds and not the polynomial degree. This is the precise reason it was possible to choose N in the first paragraph. The selection yields one tuple ζ of ℓ^N-th roots such that all its first m powers satisfy:

$$\begin{aligned} \lambda(\text{normalized module of } \xi_{\zeta^j} R_i) &\le B_i + (2J_0 + 2)\delta, \\ v_\ell V(J, \xi_{\zeta^j}) &\le v_\ell V(\varnothing, 1) + 8c + (8J_0 + 2)\delta, \\ \sum_{P \in J} v_\ell(\zeta_P^j - 1) &\le J_0\delta. \end{aligned}$$

For the last inequality use the selected monic factors; finiteness of their valuations ensures no ζ_P^j is 1. The character ξ_ζ constructed from the pencil is anticyclotomic, locally trivial at the

required old places, and all its first m powers have the same auxiliary conductor. The defining finite characters come from $U(1)$ by the anticyclotomic class-field construction of Section 21.4; taking their products and powers preserves this assertion. Integral coefficient rings can be enlarged to contain all the selected roots and coefficients, without changing the normalized upper bounds. Since $(8J_0 + 2)\delta < \eta$, these are precisely the inequalities of Proposition 20.14. They reconstruct the two-stage argument of `exists_anticyclotomic_twist` in `New/Selmer/TwistTheorem.lean`. In particular, the new conductor introduces no uncontrolled term beyond the sum of tame valuations.

To obtain `exists_mFoldTwist_selmerLengthLe` in the source one now chooses the family of rows from the finitely many eigensystems and rank-one pairs. The balance bound of `IsRow.eventually_isBalanced` is $10fn_ic_\ell$, and the auxiliary-character value bound at the old distinguished inertia is c_ℓ; use $\eta = c = c_\ell$. The result is the $(10fn_i+1)c_\ell$ normalized Selmer length and the $9c_\ell$ inverse-ratio valuation bounds of (20.17) and (20.18). These deductions use the actual value-family and balanced-lattice comparison results, not an assumption that twisting always makes a Galois representation generically acyclic.

22.11 Weak attachment: the trace and descent reductions

The global representation used in the congruence argument is not obtained by assuming an equality of local Euler factors at every place. The definition `Attached` in `Lib/GaloisRep/Main/AttachedDefs.lean` asks for a continuous semisimple representation whose characteristic polynomials approximate the specified Hecke polynomials arbitrarily closely on a particular filter of geometric Frobenius data. A set belongs to the filter if every nonempty open set of Galois elements contains a nonempty open subset such that all the Frobenii in that subset at places of prime norm, except for finitely many places, belong to the set. The filter localized near a fixed group element is proper by the existence statement of Chebotarev, so one can take limits there. This is enough to identify traces of two attached representations by continuity. We give the algebraic reconstruction from trace limits, the quadratic descent device, and the even-rank cancellation. The geometric construction of the seed, and the supply of actual automorphic approximations and matching local types, are still separate inputs.

Lemma 22.18 (Reconstruction of a limiting trace)**.** *Let k be an algebraically closed field of characteristic zero and G a group. Suppose $T : G \to k$ satisfies all polynomial identities among finitely many values of traces of n-dimensional representations of G over k. Then T is the trace of a semisimple representation of dimension n over k. If k is a topological field and T is continuous for a topological group G, the constructed representation is continuous.*

Proof. Extend T linearly to $\tau : k[G] \to k$. It is a cyclic trace, since the identity $\mathrm{tr}(AB) = \mathrm{tr}(BA)$ holds for all matrices. Put

$$I = \{a \in k[G] : \tau(ab) = 0 \text{ for all } b \in k[G]\}.$$

Cyclicity shows I is a two-sided ideal, and τ descends to $A = k[G]/I$ by taking $b = 1$. The induced bilinear pairing $\tau(ab)$ on A is nondegenerate. For any list of $n^2 + 1$ elements a_i and b_j of $k[G]$, the determinant of the matrix $\tau(a_ib_j)$ vanishes: it is a polynomial combination of trace values, which vanishes for an actual n-dimensional representation by the rank bound on the space of matrices. If A had dimension greater than n^2, one could select n^2+1 independent a_i and choose b_j with independent pairings, contrary to the determinant identity. Hence $\dim A \le n^2$.

We check the constraints on the possible trace functionals of A. For any $a \in A$, define a monic polynomial p_a of degree n from $\tau(a), \dots, \tau(a^n)$ by Newton identities, and let its

roots be $\lambda_1, \dots, \lambda_n$. For every $j \geq 1$ one has $\tau(a^j) = \sum_s \lambda_s^j$, because this identity (expressed in the coefficients by Newton identities) holds for the traces of all actual matrices. If a is nilpotent, the right side vanishes for all sufficiently large j. Distinct nonzero roots and their multiplicities would give a nonzero solution of a Vandermonde system with invertible matrix, using as many consecutive large values of j as there are distinct nonzero roots. Thus all roots are zero and $\tau(a) = 0$. If a belongs to the Jacobson radical of A, then ab is nilpotent for every b because the radical of the finite-dimensional algebra is nilpotent. The preceding argument gives $\tau(ab) = 0$ for all b and therefore $a = 0$. Consequently A is semisimple.

Over k it is a finite product of full matrix algebras $M_{d_i}(k)$. On each summand cyclicity makes τ a scalar multiple $c_i \operatorname{tr}$ of the usual trace, as is seen by applying it to products of matrix units. For a rank-one idempotent a in that summand, $\tau(a^j) = \tau(a) = c_i$ for all $j \geq 1$. The root power-sum identity just proved, or the recurrence of those moments, then forces the roots of p_a to consist only of 0 and 1. Hence c_i is the multiplicity of 1, a nonnegative integer. Nondegeneracy forces $c_i > 0$, and $\tau(1) = n$ gives $\sum_i c_i d_i = n$. The direct sum of the usual modules of the matrix summands with multiplicities c_i is the desired representation of A, and restriction along $G \to A^\times$ gives the representation of G.

For continuity, choose a basis of A and its dual basis for the trace pairing. The coordinate functions of the image of g in A are linear combinations of finitely many $T(gh)$, because the dual basis lifts to finite group-algebra sums. They are continuous, so the endomorphisms of the described representation vary continuously with g. □

This is the argument of `exists_rep_of_pseudo` and the continuity-of-trace lemmas in `Lib/GaloisRep/Pseudocharacter/FibreSum.lean` and `SemisimpleReps.lean`. In particular, pointwise limits of traces in a fixed complete coefficient field contained in $\overline{\mathbb{Q}}_\ell$, when the limiting function is continuous, satisfy the lemma. The completeness stipulation is needed: arbitrary limits in $\mathbb{C}_\ell$ need not be algebraic over $\mathbb{Q}_\ell$.

For completeness, a continuous semisimple representation of a profinite group into $\mathrm{GL}_n(\overline{\mathbb{Q}}_\ell)$ can be realized in a finite coefficient extension and stabilizes an integral lattice. Its compact image is a union of its intersections with $\mathrm{GL}_n(L)$ over finite subextensions L of $\overline{\mathbb{Q}}_\ell$. There are countably many such L (there are finitely many isomorphism classes in each finite degree, and each has finitely many embeddings). Each intersection is closed in the compact image. By the Baire property one has nonempty interior, and, being a subgroup, is open. It has finite index, so adjoining the coefficients of finitely many coset representatives to L puts the entire image over one finite extension. Over its integer ring, take the O-span of the orbit of a standard lattice. Compactness bounds its entries, so it lies between two homothetic standard lattices and is finitely generated. By construction it is stable. Enlarging the coefficient field afterward by any finite extension preserves the lattice constructions in the appropriate scalar extensions.

We next describe the construction that avoids an uncontrolled patching of Galois representations over varying extensions. The source assumes the CM field is `StablyUnram`, meaning it contains a primitive 5040-th root of unity. In particular it contains i with $i^2 = -1$ and an algebraic integer ω with $\omega + \bar{\omega} = 1$ and $\omega\bar{\omega} = 2$, since it contains the appropriate quadratic subfield $\mathbb{Q}(\sqrt{-7})$ of the seventh cyclotomic field. We use these two elements in the following elementary construction, which is the number-theoretic part of `exists_quadratic_step` in `Lib/GaloisRep/Main/AttachedDeterminant.lean`.

Lemma 22.19 (Quadratic steps with prescribed splitting)**.** *Let E be a CM field containing these i, ω and let n_0 be a positive integer. There is a quadratic character $\xi : G_E \to \{\pm 1\}$ such that, putting*

$$E' = \overline{\mathbb{Q}}^{\ker(\xi\xi^c)}, \qquad U = \ker(\xi\xi^c), \qquad H = \ker(\xi) \cap U,$$

the field E' contains a quadratic imaginary field K' in which every rational prime dividing n_0 splits and is unramified. At a place above a prime dividing n_0, its decomposition group lies in U, its inertia group lies in H, and its decomposition group lies in H if the place of E differs from its conjugate. Also $G_E = HV$ whenever a subgroup V contains inertia at all places of E over a certain rational prime not dividing n_0.

Proof. Set $k = |N_{E/\mathbb{Q}}(4)|$, write $R = 4n_0(2n_0)^k = 2^j b$ with b odd, and let $D_1 = 1 + 2b + 8b^2$. The integer D_1 is prime to R and congruent to 3 modulo 4. By Dirichlet's theorem choose a rational prime $q \equiv -D_1^{-1} \bmod R$, greater than $D_1 n_0$ and the absolute discriminant of E. Then $q \equiv 1 \bmod 4$, so $q = s^2 + t^2$ for integers s, t. The algebraic integer

$$\alpha = (1 + 2b\omega)(s + ti)$$

satisfies $\alpha\bar{\alpha} = D = D_1 q$. In particular $D \equiv -1 \bmod R$. The prime q is unramified in E with $v_P(D) = 1$ for all $P \mid q$. At each such place the valuations of α and its conjugate are nonnegative and sum to one; hence $P \neq P^c$, and in the conjugate pair one valuation is one and the other zero.

Use the Chinese remainder theorem for the pairwise coprime ideals w^k at the places $w \mid n_0$. For each pair w, w^c choose an ordering and choose u integral with $u \equiv 1 \bmod w^k$ and $u \equiv \alpha \bmod (w^c)^k$. At the self-conjugate places choose $u \equiv 1$. If there are no conditions take $u = 1$. Put $a = \alpha u\bar{u}$. It is nonzero and $\bar{a}a = D(u\bar{u})^2$. At a split conjugate pair above n_0, a is congruent to D modulo w^k on the first side and to α^2 on the other side. In the first case multiplication by i^2 makes it congruent to 1; in the second multiplication by α^{-2} does so. More generally all elements used here are local units, since D is prime to n_0. The exponent $k = 4^{[E:\mathbb{Q}]}$ is greater than $2v_w(2)$ at every such place. Therefore the unit congruence to 1 implies a local square by Hensel's lemma applied to $X^2 - a$ after the adjustment. At a self-conjugate place, the element a is still a local unit. The residue characteristic is not 2: because E contains $\mathbb{Q}(\sqrt{-7})$, and 2 splits in that quadratic field, conjugation exchanges the places of E above its two places. Thus $X^2 - a$ gives a split or unramified quadratic local extension at any self-conjugate place of the prescribed set.

Let ξ be the quadratic character of $\sqrt{a}$. Then $\xi\xi^c$ cuts out $E(\sqrt{-D})$ because $i\sqrt{a}\sqrt{\bar{a}}/(u\bar{u})$ has square $-D$. The square class of $-D$ is nontrivial (valuation at every $P \mid q$ is odd). For rational primes dividing n_0, the congruence $D \equiv -1 \bmod R$ shows that $-D$ is a square in the corresponding $\mathbb{Q}_p$ by Hensel's lemma; the exponent of 2 is sufficient since R is divisible by 8 if $2 \mid n_0$. It follows that $K' = \mathbb{Q}(\sqrt{-D})$ has the claimed splitting. Thus the local decomposition groups lie in U, and the local square and unramifiedness assertions above for a give the claimed restrictions to H.

Finally, for a pair P, P^c above q, the valuations of $u\bar{u}$ have the same parity on the two sides, while those of $\alpha, \bar{\alpha}$ have opposite parity. Since the residue characteristic is odd, an odd valuation gives a ramified quadratic extension, and an even valuation gives a split or unramified quadratic extension. The restrictions of (ξ, ξ^c) on inertia at these two places consequently have the two independent nonzero images $(1, -1)$ and $(-1, 1)$ in $\{\pm 1\}^2$, in some order. The kernel of the pair is H; hence every subgroup containing these inertia groups maps onto the quotient G_E/H and gives $G_E = HV$. This proves the lemma. □

The detailed theorem `exists_quadratic_step` also asserts the normalized Weil and base-change conditions `IsStep`. We have reconstructed the splitting and quotient-control part here; those additional automorphic conditions are used as a compatibility input.

Lemma 22.20 (Gluing on a common quotient)**.** *Let G be profinite, H_0, H_T open normal subgroups, and r_0, r_T continuous semisimple representations on these subgroups over an*

algebraically closed field of characteristic zero. Suppose their restrictions to $H_0 \cap H_T$ have equal traces for every sufficiently large T. Suppose a closed normal subgroup K of G is contained in $\bigcap_T \ker r_T$ (for those sufficiently large T), and $G = H_0 K$. Then there is a continuous semisimple representation $r : G \to \mathrm{GL}_n$ extending r_0, whose restriction to each such H_T is conjugate to r_T.

Proof. The restrictions to $H_0 \cap H_T$ are semisimple. One elementary form of Clifford's argument is to note that the sum of the translates of an irreducible submodule for the finite-index normal subgroup is a nonzero module for the original group, and hence is the whole irreducible representation. As a sum of simple submodules it is semisimple. Restriction of a semisimple representation then gives semisimplicity. Equality of traces identifies semisimple representations by the argument of Lemma 21.3. In particular r_0 is trivial on $H_0 \cap K$. Therefore define $r(hk) = r_0(h)$ for $h \in H_0$, $k \in K$. This is a well-defined representation of G/K and hence of G, as $H_0/(H_0 \cap K) \simeq G/K$. The isomorphism is a homeomorphism of compact Hausdorff groups, giving continuity. Semisimplicity follows from the image equality with $r_0(H_0)$. For an element of H_T write it as hk with $h \in H_0$, $k \in K$. Since $K \subseteq H_T$ and is killed by r_T, $h \in H_0 \cap H_T$ and the fixed conjugacy on this intersection gives the desired relation on H_T. □

Applied as follows, this proves the group-theoretic descent assertion used in `sphericalCompat_of_odd` of `Lib/GaloisRep/Main/AvatarCompatibility.lean`. For an eigenform of odd rank with finite level set S, choose a sequence of integers n_T divisible by ℓ and all primes under S and eventually divisible by every given rational prime. Apply Lemma 22.19 to obtain H_T. The weak base-change theorem and odd-rank split-place attachment supply the restricted representations r_T on H_T, with approximation of the original Hecke polynomials there, and unramifiedness at the places dividing n_T outside S, ℓ. Equality of their traces on $H_0 \cap H_T$ follows from weak attachment and Chebotarev. In Lemma 22.19 for H_0, the relevant inertia prime lies outside n_0 and hence outside S, ℓ. For all sufficiently large T, it divides n_T, so its inertia groups lie in H_T and are killed by r_T. Put K equal to the normal core of the intersection of these kernels. It contains all the indicated inertia groups, and the quotient-control assertion gives $G_E = H_0 K$. Lemma 22.20 gives r.

To verify the weak approximation on G_E from its restrictions, take any nonempty open subset of G_E and a geometric Frobenius of a place of split prime norm there, outside the finite exceptional set (Chebotarev). Choose T so large that this place divides n_T; its decomposition group lies in H_T by Lemma 22.19. The intersection of the open set with H_T is nonempty and open. The restricted weak attachment at this T supplies the required subwindow of Frobenius approximation, which is also valid for r by conjugacy. This gives weak attachment at G_E in the precise filter sense. The local comparison requires the additional assertion of the source that finite type and jumps descend from the finite extension (`IsWeakBaseChange.exists_attached_descent` and its local comparison lemmas). At a split spherical place one takes T for which its decomposition group is in H_T; the killing-polynomial identity is then inherited directly.

The geometric part not supplied by these elementary arguments is the existence of the split-place odd-rank seed with free weights and the approximation by such seeds, together with the normalized weak base change. The source constructs it by PEL Shimura varieties and cohomological representations, with representations extracted from traces and local type information. Principal anchors are `exists_rep_of_cuspidal_type` in `Lib/GaloisRep/Main/SeedRepresentation.lean`, `sphericalCompat_of_split` in `Main/AvatarCompatibility.lean`, and `isApproximatingMonoid_of_mem_closure` in the weight-congruence development. They are substantial known-theory inputs in the present reconstruction. The

point of the quotient calculation is that it does not require an independent modularity-lifting or eigenvariety patching theorem once those inputs are available.

Proposition 22.21 (Two theta strings remove the parity restriction)**.** *Suppose a Hecke system of rank n admits two odd-rank theta lifts of rank $n+t_0$, whose normalized Hecke polynomials off a finite set are $P_w D_{j,w}$ for $j = 0, 1$. Assume the auxiliary factors are cyclotomic or algebraic character strings, with attached continuous character lists $\lambda_{j,1}, \dots, \lambda_{j,t_0}$, disjoint between the two lists. If the odd-rank systems have continuous semisimple attached representations R_j, then there is a continuous semisimple n-dimensional representation with trace*

$$\operatorname{tr} R_0 - \sum_{i=1}^{t_0} \lambda_{0,i} = \operatorname{tr} R_1 - \sum_{i=1}^{t_0} \lambda_{1,i}.$$

It is weakly attached to P_w. At a spherical place where both theta systems have the required local comparison, if $D_{0,w}, D_{1,w}$ are coprime, their two killing-polynomial identities imply that P_w kills the corresponding Frobenius on this representation.

Proof. The trace equality follows by taking the coefficient of degree $n + t_0 - 1$ in the two factorizations and applying weak attachment, taking limits near each Galois element. Finite-character twists used in normalizing the polynomials are already accounted for in R_j. By equality of traces of semisimple representations, every character in the first list occurs with at least its listed multiplicity in R_0 (use disjointness and compare the multiplicity on the other side of the equality). Thus the virtual trace is the trace of a direct summand of rank n, with equal multiplicities for the analogous direct summand of R_1. These are actual continuous semisimple representations. They satisfy the stated trace equality and, by Newton identities on the powers of Galois elements, their characteristic polynomials factor the actual characteristic polynomials of both R_j by the corresponding character polynomial. At the generic Frobenii, the norm of the error after multiplication by the monic auxiliary polynomial bounds the norm of the original error, by multiplicativity of the non-Archimedean polynomial Gauss norm and the fact the auxiliary polynomial has norm at least one. This proves weak attachment to P_w, including arbitrary approximation precision. Direct summands of unramified or crystalline representations have the corresponding local property; their multiset of Hodge–Tate jumps is obtained by removing the character jumps. For the killing statement, on the common summand both $P_w D_{0,w}$ and $P_w D_{1,w}$ annihilate the operator, and the Bezout identity for the two D gives the result. □

In the source, `sphericalCompat_std` in `Lib/GaloisRep/Main/SphericalAllRanks.lean` uses $t_0 = n + 1$ and two unramified $U(1)$-characters with different weights. Their powers and finite values are arranged so that their auxiliary characters differ at every split place. This gives the disjointness and the Bezout step of the proposition. The specialized auxiliary theta-lift existence `exists_theta_lift_heckePoly` and the exact jump and local character normalizations remain part of the automorphic input. Weak attachment for the rows of the congruence argument needs specifically the tame parahoric label compatibility, not just this spherical construction. The metacyclic calculation in Chapter 21 applies once these label and inertia polynomials are supplied; in particular full matrix span cannot be inferred from spherical attachment alone.

22.12 Hecke-polynomial continuations

The false lifts in the denominator estimate need not themselves be theta lifts, but their local unitarity may force a fixed Hecke root string. We record the exact polynomial arguments.

Lemma 22.22 (An integral normalized multiple)**.** *Let A, A', B, C be polynomials over $\overline{\mathbb{Q}}_\ell$ with B monic and C integral. If*

$$\|AB - C\|_{\mathrm{G}} < 1, \quad \|A'B - C\|_{\mathrm{G}} < 1, \quad |A'(1)|_\ell \geq 1$$

for the polynomial Gauss norm, then there is an integral polynomial Q divisible by A with $Q(1) = 1$.

Proof. The Gauss norm is multiplicative with $\|B\|_{\mathrm{G}} \geq 1$. Consequently $\|A\|_{\mathrm{G}} \leq 1$ and, subtracting the two errors, $\|A - A'\|_{\mathrm{G}} < 1$. Evaluation at 1 implies $|A(1)|_\ell = |A'(1)|_\ell \geq 1$. The first norm bound implies equality with one. Set $Q = A/A(1)$. □

Lemma 22.23 (Eliminating eigenvalue one)**.** *Let T be a continuous n-by-n matrix function of a group element over $\overline{\mathbb{Q}}_\ell$. Suppose $T(g_i) = s^i A$ for $i \geq 0$, with s a unit and $|s^j - 1|_\ell = 1$ for $1 \leq j \leq n$. Some $i \in \{0, \ldots, n\}$ has a neighborhood on which $|\det(1 - T(g))|_\ell \geq 1$.*

Proof. If $|\det(1 - s^i A)|_\ell < 1$, some eigenvalue λ of A has $|1 - s^i\lambda|_\ell < 1$. The same eigenvalue cannot do so for two distinct i, j in the range: otherwise $|s^{i-j} - 1|_\ell < 1$. If the strict inequality held for all $n + 1$ indices there would be at least $n + 1$ distinct eigenvalues. At the selected point the nonzero determinant has locally constant absolute value by continuity. □

These prove the declarations `exists_integral_multiple_eval_one` and `exists_eventually_no_eigenvalue_one` in `New/Selmer/HeckeString.lean`. The scalar inertia of a row allows application of the second lemma. The possible false lifts on which every element has eigenvalue one are handled differently.

Proposition 22.24 (String from a root of exceptional magnitude)**.** *Suppose P is the spherical Hecke polynomial of an irreducible unramified unitary representation of GL_n over a local field of residue cardinality q, in the normalization in which the roots are $q^{(n-1)/2}$ times normalized Satake parameters. If c is a root and $|c|^2 q^t = q^n$ for an integer $t > 0$, then cq^i is a root for every $0 \leq i < t$.*

Proof. The classification of the spherical unitary dual expresses the parameters as centered strings

$$\{uq^{-(a-1)/2}, uq^{-(a-1)/2+1}, \ldots, uq^{(a-1)/2}\}, \qquad |u| = 1,$$

and pairs of such strings shifted by the powers $q^\delta, q^{-\delta}$ for $0 < \delta < 1/2$. This is the spherical form of Tadić's classification (M. Tadić, *Spherical unitary dual of general linear group over non-Archimidean local field*, Annales de l'Institut Fourier 36 (1986), no. 2, 47–55, DOI: 10.5802/aif.1046). The normalized magnitude of c is $q^{-(t-1)/2}$. Its exponent is a half-integer, so cannot occur in strictly shifted strings. In a centered string of length a it occurs only if $a \geq t$ and the parities agree. Starting at exponent $-(t-1)/2$, the string continues through all exponents up to $(t-1)/2$ with the same phase. Thus the next $t - 1$ roots obtained by multiplication by q are present. □

For a unitary cyclic local representation generated by a spherical eigenvector use its corresponding irreducible spherical constituent. Alternatively the elementary double-coset formula `isRoot_string_of_norm` in `Lib/LocalGL/Tadic/ChainLemma.lean` applies directly to its invariant pairing. The rescaling by `chiOneInv` is essential: the root 1 of the normalized Galois representation corresponds to a root of exactly this exceptional magnitude. If a normalized attached representation has eigenvalue 1 at every group element, the local polynomial-killing property thus implies at a split spherical place all roots cq^i for $0 \leq i < t = n - r$, where r is the row rank. Divide the Hecke polynomial by these distinct factors

and rescale the remaining factor. On generic Frobenii its product with the rescaled string is congruent coefficientwise to the characteristic polynomial of the attached representation. The latter is integral by the lattice argument of Section 22.11. If another system has the same characteristic-polynomial reduction and a distinguished row polynomial A' whose value at 1 has norm at least one (as ensured by Lemma 22.23), then the remaining factor A satisfies Lemma 22.22 with the common monic string B. It therefore divides an integral polynomial equal to 1 at 1. This is the algebraic implication `eventually_heckePoly_eq_string_mul` in `New/Selmer/HeckeString.lean`.

This distinguishes two kinds of false lift. If some Galois point has no eigenvalue one, finite polynomial avoidance as in Section 20.10 makes a normalizer determinant nonzero on that system. If every group element has eigenvalue one, the Hecke string is forced and such systems can be tested at new primes. One cannot conclude solely from the string that these systems are actual theta lifts. The specialized patching estimate in `New/Patching/OnePrimeTie.lean` is still needed to transport the first-order relations into integral operators at the original level. Its finite transfer algebra was studied in Section 21.9.

22.13 Rank-one normalization of the projection bound

A finite averaging calculation accounts for the class factor in the rank-one pairing. The test condition in `AuxPair.IsTestForm` of `New/Selmer/FitTwists.lean` involves $\ker\Theta$ intersected with the label-$X-1$ idempotent image and the test Hecke kernel; here Θ is the descent map to forms on the rank-one group. By contrast, for the line Z of the trivial lift, the condition in `DenominatorLe` involves $Z^{\perp}$ in place of $\ker\Theta$. A form satisfying the first condition need not satisfy the second. The averaging bridges the difference.

Lemma 22.25 (Averaging onto the rank-one theta line)**.** *Let C_1 be the finite group of double classes of $U(1)$ at the chosen level and assume these classes have trivial stabilizers, writing $h' = |C_1|$. Let the kernel vectors K_g, $g \in C_1$, give $\Theta(v)(g) = \langle K_g, v\rangle$ and $L = \sum_{g\in C_1} K_g$. Suppose the twisted central action $T(s)$ satisfies*

$$T(s)^{-1}K_g = K_{gs}, \quad T(s)L = L, \quad T(s)^* = T(s)^{-1},$$

and commutes with the label projection e_0 and every test Hecke polynomial. Assume the class representatives s_g can be chosen integral above ℓ with $T(s_g)^{-1}$ preserving the integral lattice at the chosen embedding. If `DenominatorLe` holds for the projection to $Z = \mathbb{C}L$ with rank-one allowance β, then for two integral rank-one test forms v, v' one has

$$|\iota(h'\langle v, P_Z v'\rangle)|_\ell \le e^{\beta}.$$

Proof. The zero-line case is immediate. Put $\widetilde{v}' = \sum_{g\in C_1} T(s_g)^{-1}v'$, integral at the chosen embedding. For $w \in Z^{\perp}$ lying in the image of e_0 and in the kernel of a test polynomial, $A(w) = \sum_g T(s_g)w$ belongs to the same image and kernel. Also

$$\Theta(A(w))(h) = \sum_g \langle T(s_g)^{-1}K_h, w\rangle = \sum_g \langle K_{hs_g}, w\rangle = \langle L, w\rangle = 0.$$

This uses the permutation of the classes by multiplication and their equal weights from trivial stabilizers. Thus $A(w)$ lies in the test intersection containing the kernel of Θ. Since v' is a rank-one test form, $\langle w, \widetilde{v}'\rangle = 0$. This is the stronger orthogonality required by the rank-one denominator bound. Finally $P_Z T(s_g)^{-1} = P_Z$ because $T(s_g)$ is unitary and fixes L. Hence $P_Z\widetilde{v}' = h' P_Z v'$. Applying the rank-one denominator bound to $v, \widetilde{v}'$ gives the assertion. □

The hypotheses on the actions and integral representatives are verified in the source by the scalar Weil identity, the algebraic character norm in the Weyl lattice, and integral representatives at the chosen level. The finite calculation above is `DenominatorLe.rank_one_bound` in `New/Selmer/RankOneDenominator.lean`. It explains why the ordinary class factor multiplies the pairing, rather than entering as an unexplained volume normalization.

22.14 Dependency status of the congruence argument

The odd-class bound of Proposition 20.16 has been supplied by an actual conventional calculation in Theorem 22.14. The uniform detecting-coordinate theorem of Section 22.6 also supplies an input to the even-part triangular construction. We gave the two-stage common-twist proof from the mirror criterion, balanced-lattice estimate and value-distribution formula rather than treating it as a single black box. We also proved the trace-reconstruction and quotient-descent parts of weak attachment, the elementary quadratic steps, the even-rank cancellation once two appropriate lifts are available, the Hecke-string algebra, and the rank-one projection averaging.

These developments do not prove the entire denominator theorem `denominatorLe_rowBudget` by themselves. Its remaining obligations include the geometric seed and exact local-global type compatibility in the attachment, the global-duality construction of the mirrors and the odd reciprocity laws used for the even-part triangle, the precise integral value measures of the ramified pairs, and the arithmetic Ribet-to-Hecke lifting step at the actual patched auxiliary levels. The integral comparison of crystalline lattices in the following chapters remains essential. An explicit uniform bound on the odd modules does not control the even modules, any more than qualitative local-global compatibility controls the torsion lattice losses.

Chapter 23

Local crystalline defects and the integral cup-product estimate

Chapters 24 and 25 supply further proofs of the classical cohomological rank and index calculation, the chart algebra, depth and long-gap lifting and integral cokernel estimates, subject to the precise comparison hypotheses specified there. The remaining geometric integral comparison hypotheses are not replaced by a general rational comparison theorem.

The congruence argument for the lower bound on the determinant needs considerably more than the vanishing of the ordinary Bloch–Kato pairing on two crystalline representations. The two classes to be paired arise from lattices which need not be dual to one another in characteristic zero; only a pairing of their reductions is given. Moreover the lattices occurring in the congruences have weights and ranks increasing with the main parameter x. An error term depending exponentially on a rank, or even a loss of order b on every line of a rank r lattice at every small prime, would not suffice. This section formulates the local integral problem, explains the Kummer-chart calculation used to solve it, and shows exactly how the resulting loss enters the global congruence budget. The integral period-ring assertions quoted here as *formal-source inputs* are statements of the compact Lean development, not assertions for which we claim an independent verification of the proof. In contrast, we give the length, valuation, and budget deductions explicitly.

23.1 Lattices, crystalline cocycles and imperfectly dual reductions

Write $p = \ell$ for the rational prime throughout the local discussion, $G_p = \mathrm{Gal}(\overline{\mathbb{Q}}_p/\mathbb{Q}_p)$, and let $j : \Gamma \hookrightarrow G_p$ identify Γ with an open subgroup. Write F for the corresponding finite extension of $\mathbb{Q}_p$ and $d_F = [F : \mathbb{Q}_p]$. The coefficient ring O denotes the integers of a sufficiently large finite extension of $\mathbb{Q}_p$, with fraction field K_O. For its uniformizer π_O, put $e_O = v_{\pi_O}(p)$; our valuation v_p on K_O has $v_p(p) = 1$. In particular the normalized length of a torsion O-module M is

$$\mathrm{len}_p M = e_O^{-1} \,\mathrm{length}_O M.$$

It is additive on short exact sequences and $\mathrm{len}_p(O/p^aO) = a$. Infinite lengths can be understood as $+\infty$. For submodules S, T of a common module, we use the notation

$$\mathrm{rl}(S,T) = \mathrm{length}_O(S/(S \cap T)), \qquad \mathrm{rl}_p(S,T) = e_O^{-1}\,\mathrm{rl}(S,T).$$

These correspond to `relLength` and `CoeffRing.normLength` in the formal source. The normalization is essential: passing to a more ramified coefficient field must not artificially

enlarge the estimates. For a finite-length module over O, its length over $\mathbb{Z}_p$ is $[O/\pi_O : \mathbb{F}_p]$ times its length over O. Consequently

$$\mathrm{rl}_p(S,T) = \frac{\mathrm{length}_{\mathbb{Z}_p}(S/(S\cap T))}{[K_O:\mathbb{Q}_p]}. \tag{23.1}$$

The statements below are uniform in O once the necessary finite coefficient fields and roots of unity are included. This is the meaning of the eventual quantifier on coefficient rings in the source.

A Galois lattice T is finite free over O, with continuous O-linear Γ-action. Put $T^*(1) = \mathrm{Hom}_O(T, O(1))$, where the cyclotomic character acts on $O(1)$. The jump convention is that $O(1)$ has jump -1. At an embedding $\gamma : O \to \mathbb{C}_p$ let $j_T(\gamma, u)$ denote the multiplicity of jump $u \in \mathbb{Z}$; thus

$$j_{T^*(1)}(\gamma, -u-1) = j_T(\gamma, u). \tag{23.2}$$

The local development defines $j_T(\gamma,u)$ by the dimension over $\overline{\mathbb{Q}}_p$ of the space of γ-semilinear functionals $f : T \to \mathbb{C}_p$ satisfying $f(gv) = \chi_{\mathrm{cyc}}(g)^{-u} g f(v)$ for g sufficiently close to 1. In particular its weights are Hodge–Tate weights in the above sign convention, not the oppositely signed integers sometimes used for cyclotomic twists.

The O-module of crystalline cocycles $Z_f(T)$ consists of maps $c : \Gamma \to T$ such that there is $y \in B_{\mathrm{cris}} \otimes_{\mathbb{Z}_p} T$ with

$$1 \otimes c(g) = gy - y \qquad (g \in \Gamma).$$

The tensor product here is over $\mathbb{Z}_p$, and the action is diagonal. The definition is useful also for potentially crystalline T. This is the inverse image in $Z^1(\Gamma, T)$ of the usual kernel of the map to $H^1(\Gamma, B_{\mathrm{cris}} \otimes_{\mathbb{Q}_p} T[1/p])$. The cocycle identity follows by applying $gh - 1 = (g-1) + g(h-1)$ to y and using injectivity of the natural inclusion. The module includes all integral coboundaries. It is saturated with respect to multiplication by a nonzero element of O: if ac is a crystalline cocycle, divide its period primitive by a in the tensor product with B_{cris}, where K_O acts and a is invertible. In particular

$$Z_f(T)/p^a Z_f(T) \hookrightarrow Z^1(\Gamma, T/p^a T). \tag{23.3}$$

The formal source also proves continuity, including the eventual local constancy of reductions of such cocycles. The relevant definitions and results are `crysCocycles`, `jump` and `HasJumps` in `Lib/PadicHodge/Crystalline/Defs.lean`, `crysCocycles.of_smul` in `Lib/PadicHodge/Crystalline/Jumps.lean`, and `crysCocycles.eventually_eq_pow_smul` in `Lib/PadicHodge/BlochKato/CocycleContinuity.lean`.

Proposition 23.1 (Finite local cocycle rank)**.** Formal-source input. *Let T be potentially crystalline. Then $Z_f(T)$ is finite free over O. Rank over O is at most $(1+d_F)\,\mathrm{rk}_O T$. More precisely, in the situation where the jumps are indexed by embeddings of the completion of E in $\overline{\mathbb{Q}}_p$ and hence by a set Σ_v of complex embeddings of E determining a finite place $v \mid p$,*

$$\mathrm{rk}_O Z_f(T) = \mathrm{rk}_O T + \sum_{\sigma\in\Sigma_v} \#\{u \in J_T(\sigma) : u < 0\}.$$

The same formula with a local field of degree d_F and with summation over its embeddings applies after extending coefficients so that all the Hodge jumps can be read there.

This is the Bloch–Kato rank calculation on cocycles (as opposed to on classes), expressed by `crysCocycles.finrank_eq` in `Lib/PadicHodge/BlochKato/BlochKato.lean` and the embedding form `finrank_crysCocycles_eq_add_neg_jumps` in `New/Selmer/Crystalli`

`neRank.lean`. The rank formula on classes comes from the fundamental period-ring exact sequence and the de Rham filtration; adding the dimension of the coboundaries cancels the contribution of invariants. The formal-source statement on $\mathbb{Z}_p$-ranks uses d_F times the number of negative jumps in the multiset of a restriction of scalars. In the indicated embedding setting, decomposing that multiset and dividing by $[K_O : \mathbb{Q}_p]$ gives the asserted formula. Freeness follows from torsion-freeness and finite generation over a DVR. In particular the rank bound requires no extra hypothesis on inertial eigenvalues.

We describe the class of lattices against which the defect is measured. Let U be a family of finite subsets of $\mathbb{Z}$ indexed by coefficient embeddings. The notation $\mathcal{Q}(F_0, \mathcal{K}, U)$ will mean the class `Lift.InClass` of the formal source; here F_0 is the filter on Γ that specifies the required fixed crystalline periods, and $\mathcal{K}(g, Q)$ specifies a family of polynomials annihilating the operator on potentially crystalline periods associated with g. More explicitly a member Y of the class is admissible for periods fixed almost everywhere for the supremum of the near-identity filter and F_0, has $j_Y(\gamma, u) = 0$ for $u \notin U(\gamma)$, and, for g mapping to inertia, each polynomial Q with $\mathcal{K}(g, Q)$ annihilates the linear operator $\beta_{g,0}$ on $B_{\mathrm{cris}} \otimes_{\mathbb{Z}_p} Y$. This operator is the B_{cris}-linearization of g on the potentially crystalline periods (the source calls it `linearFrob`). For $F_0 = \top$, $\mathcal{K} = \bot$ we just require crystallinity and the specified weights. All members are potentially crystalline. The jump and polynomial conditions are preserved when a lattice is a subrepresentation of a direct sum of the given representations. The dual of such a class has jump sets $-U - 1$ and inertial polynomial condition at inverse elements. These facts are expressed in `InClass.of_jointly_injective`, `InClass.dual` and `InClass.tensor`, `Lib/PadicHodge/Crystalline/PairingDefs.lean`.

An imperfectly dual pairing of lattices Y, T is specified by a positive modulus p^a, an O-linear form $\psi : Y \otimes_O T \to O(1)$, a $\mathbb{Z}_p$-linear functional $\lambda : O \to \mathbb{Z}_p$, and a crystalline cocycle $q \in Z_f((Y \otimes T)^*(1))$ satisfying

$$g\psi - \psi = p^a q(g). \tag{23.4}$$

Thus the induced bilinear form $\bar{\psi}$ with values in $\mathbb{Z}/p^a(1)$, obtained by applying λ and reducing modulo p^a, is equivariant; the unreduced pairing need not be equivariant. The cup product of $c \in Z_f(Y)$ and $c' \in Z_f(T)$ under $\bar{\psi}$ is represented by

$$(g, h) \longmapsto \bar{\psi}(c(g), gc'(h)) \in \mathbb{Z}/p^a(1).$$

The vanishing required in the source means that, after the inclusion of the cyclotomic module $\mathbb{Z}/p^a(1)$ into $\overline{\mathbb{Q}}_p^{\times}$, the displayed cocycle is a continuous coboundary. One should not replace this assertion by an assertion of vanishing in $H^2(\Gamma, \mathbb{Z}/p^a(1))$ without an argument. The chosen codomain is the one used by the index form of local Tate duality below. Both the imperfect equivariance and the third class q are necessary: it would be illegitimate simply to apply the pairing of $D_{\mathrm{cris}}(T)$ with its exact dual to a representation merely congruent modulo p^a to that dual.

Define the submodule $T^{\perp_{\mathcal{Q}}}$ of maps $c' : \Gamma \to T$ by imposing this cup-product vanishing for every $Y \in \mathcal{Q}$, every $c \in Z_f(Y)$, and every pairing of the above form, including every modulus. Define its crystalline defect

$$\Delta_p(T; \mathcal{Q}) = \mathrm{rl}_p(Z_f(T), T^{\perp_{\mathcal{Q}}}).$$

We use the symbol a in pairing modulus p^a only in this local definition: in estimates below it is the inherited global derivative-order parameter unless specifically renamed. The Lean definitions of pairings, their cup product, and the orthogonal submodule are, respectively, `Lift.Pairing`, `Pairing.CupBounds`, and `cupOrthogonal` in `Lib/PadicHodge/Herr/GaloisLattices.lean`.

Lemma 23.2 (Exponent-to-length conversion)**.** *If $p^s Z_f(T) \subseteq T^{\perp_{\mathcal{Q}}}$, then*

$$\Delta_p(T;\mathcal{Q}) \le 2d_F\, s\, \mathrm{rk}_O\, T.$$

Proof. There is a surjection from $Z_f(T)/p^s Z_f(T)$ onto the defect quotient, because its denominator contains $p^s Z_f(T)$. The former has normalized length $s\,\mathrm{rk}_O\, Z_f(T)$. Proposition 23.1 yields $s(1+d_F)\,\mathrm{rk}_O\, T \le 2sd_F\,\mathrm{rk}_O\, T$, since $d_F \ge 1$. This is precisely the numerical content of `normLength_defect_le` in `New/Selmer/CrystallineRank.lean`. □

23.2 Why this defect bounds the local discrepancy of congruent cocycles

Take a finite list of global integral lattice representations $u_i : G \to \mathrm{GL}_{n'}(O)$ and a block representation $\rho : G \to \mathrm{GL}_{r'}(O)$. Fix a modulus p^M and write $R = O/(p^M)$. The global image of the reducibility ideal gives, on the block, cocycles $c : G \to R^{r'}$ satisfying not only the cocycle identity, but also all linear relations among the matrices $(u_i(g)-1)_i$. Let $\mathcal{Z}$ denote the module of cocycles satisfying these relations. Call its local crystalline-property submodule at $v \mid p$ the intersection with the module of cocycles whose restriction to G_v obeys the relations of $Z_f(T_v)$, where $T_v = (O^{r'}, \rho|_{G_v})$. Here "the relations of a module" $Z \subseteq \prod_g O^{r'}$ are finitely supported O-linear relations vanishing on every element of Z; they are imposed modulo p^M. This is the definition in `obeyingCocycles` and `crysRels`, in `New/Selmer/BlocksAndCocycles.lean` and `New/Selmer/DeterminantBound.lean`. This local condition is stronger than continuity and the cocycle identity, and must not be assumed automatically from the construction of a congruence. The entire point of the local estimate is to bound how far $\mathcal{Z}$ is from imposing it.

Here is the algebraic reason these cocycles are amenable to crystalline estimates. Let P be the finite free module of the tuple of matrices, with the action by left multiplication by $(u_i(g))_i$. Put $D(g) = (u_i(g)-1)_i$ and $A = \sum_g OD(g) \subseteq P$. Then $D(gh) = D(g) + u(g)D(h)$, so A is stable under G. Since O is Noetherian, A is finite. If c obeys the displayed relations then a linear map $q : A \to R^{r'}$ is defined by $q(D(g)) = c(g)$; it is equivariant by the two cocycle identities. By the Artin–Rees lemma, for some j large enough $A \cap p^{Mj}P \subseteq p^M A$. Consequently q extends to a well-defined equivariant map

$$q' : Y = A + p^{Mj}P \longrightarrow R^{r'}, \qquad q'(a + p^{Mj}z) = q(a).$$

Indeed if two such expressions have equal first component then their difference in A belongs to $p^M A$, on which q vanishes. The action on $p^{Mj}P$ is stable. The map satisfies $q'(D(g)) = c(g)$, and Y is a finite free lattice in P. Equivalently the graph of q' projects injectively into P. This is the constructive role of the graph subrepresentation in `obeyingCocycles.exists_graph_subrepresentation`. In particular after restriction to G_v, Y belongs to the class $\mathcal{Q}$ of lattices inherited from the u_i, whenever those representations have the requisite admissibility, jump sets and annihilating polynomials. The map $g \mapsto D(g)$ is itself a crystalline cocycle whenever the ambient representations are potentially crystalline: over B_{cris} its primitive is the tuple of identity matrices, and multiplication by p^{Mj} puts this primitive in Y after extending scalars. Thus every obeying cocycle has the local form of a reduction of a linear image of a crystalline cocycle on Y. It need not be a reduction of a crystalline cocycle on T_v.

For the duality step, lift q' to an O-linear map $q_1 : Y \to T_v$ (possible by freeness). The reduction is equivariant, so $gq_1 - q_1 = p^M q_2(g)$ with $q_2(g)$ integral. Under the perfect pairing $T_v \otimes T_v^*(1) \to O(1)$ this is exactly (23.4) on $Y \otimes T_v^*(1)$. Its third class q_2 is crystalline: $p^M q_2$ is a coboundary, and saturation shows that q_2 lies in $Z_f((Y \otimes T_v^*(1))^*(1))$. Hence every reduction in $\mathcal{Z}|_{G_v}$ annihilates the reduction of the orthogonal module $(T_v^*(1))^{\perp_{\mathcal{Q}}}$ under the local Tate

pairing. The same conclusion holds for reductions from $Z_f(T_v)$ itself, since $T_v \in \mathcal{Q}$ and the usual perfect pairing is permitted.

Proposition 23.3 (Local congruence loss)**.** The local Tate-duality index estimate, derived in Lemma 24.2 under the classical duality and crystalline orthogonality hypotheses. *Under the preceding hypotheses,*

$$\mathrm{rl}_p\left(\mathcal{Z}, \{c : c|_{G_v} \text{ obeys all relations of } Z_f(T_v)\}\right) \le \Delta_p(T_v^*(1); \mathcal{Q}).$$

For clarity, the content of the index estimate is the following. A finite module $V = T_v/p^M T_v$ has the Cartier dual provided by the reduction of the $\mathbb{Z}_p$-dual Tate lattice. A $\mathbb{Z}_p$-linear functional on O can be chosen to make the pairing perfect: the $\mathbb{Z}_p$-dual of O is a rank-one O-module (the inverse different). The images A and A' of $Z_f(T_v)$ and $Z_f(T_v^*(1))$ have orders whose product is

$$p^{Md(2+d_F)}, \qquad d = \mathrm{rk}_{\mathbb{Z}_p} T_v. \tag{23.5}$$

Indeed, by (23.3) their orders are the powers of p with exponents M times their $\mathbb{Z}_p$-ranks. The dual jump relationship (23.2) pairs each nonnegative jump of T_v with a negative jump of its dual; their ranks therefore sum to $(2 + d_F)d$. The size bound from local Tate duality says that the number of extra classes of the restriction $\mathcal{Z}|_{G_v}$ modulo its intersection with A divides the index of the orthogonal part of $Z_f(T_v^*(1))$. It is applied at the cocycle level, not after omitting the coboundaries; the count (23.5) includes those terms. Every member of A satisfies the indicated crystalline relations, so passing from the condition of belonging to A to that of satisfying its relations can only decrease the loss. Pulling the resulting estimate back along restriction proves the inequality from the cited index estimate. The exact formal implementation is `IsPotCrys.relIndex_dvd_of_orthogonal` in `Lib/ClassField/LocalDuality/CrystallineDuality.lean` and `relLength_obeyingCocycles_le` in `New/Selmer/RowBalanced.lean`. A separate proof of the former would require developing the local Tate pairing and its index formula; here we have isolated where that input is used. Note especially that the right side is independent of M. This uniformity is what makes the first-order global congruence construction effective.

23.3 The Kummer chart, letter decomposition and height of a frame

We continue the local discussion, writing p for ℓ and reserving E without an additional subscript for the global CM field. In particular, the Eisenstein element in the chart will be denoted by $\mathfrak{e}$ (the Lean notation for it is E). Let $A_{\mathrm{inf}} = W(\mathcal{O}_{\mathbb{C}_p^\flat})$ and B_{dR}^+ be Fontaine's usual period rings, with the map $\vartheta : A_{\mathrm{inf}} \to \mathcal{O}_{\mathbb{C}_p}$ and Frobenius φ on A_{inf}. Write t_{dR} for the uniformizing parameter of B_{dR}^+, so $\mathrm{Fil}^n B_{\mathrm{dR}} = t_{\mathrm{dR}}^n B_{\mathrm{dR}}^+$ for $n \in \mathbb{Z}$. This avoids confusion with the rank difference $t = 5r$ of the global proof.

The local chart fixes an unramified coefficient ring $\Lambda = W(\mathbb{F}_{p^{f_0}})$, a ring $\mathfrak{S} \cong \Lambda[[z]]$ embedded in A_{inf}, and an element

$$\mathfrak{e} = z^{e_0} A + p \quad (A \in \mathfrak{S}^\times)$$

associated with a generator of $\ker \vartheta$. Here z is the Teichmüller lift of the element of the tilt arising from a tower of roots of a local uniformizer. For the completion-group case, $e_0 f_0 = d_F$. The chart also contains an open subgroup Γ' of index prime to p in Γ, preserving the indicated constants and the untilt of z, with the required transitivity on roots. Thus local cup product vanishing may be tested after restriction to Γ', since restriction followed by

corestriction multiplies by the prime-to-p index, and the relevant classes are p-power torsion. The unramified chart has $e_0 = 1$, $\mathfrak{e} = z + p$. For the quantitative depth steps the relevant number is not the full discriminant of the global CM field but

$$\delta_0 = \begin{cases} 0, & e_0 = 1, \\ v_p(e_0) + 1, & e_0 > 1. \end{cases} \tag{23.6}$$

In particular δ_0 is a specified upper bound used in the derivative estimate of the chart and is not the actual different exponent in an arbitrary normalization. The existence of such a chart in the completion-group case, including the subgroup and the compatibility when the extension is unramified, is a *formal-source input*, `exists_kummer_chart` in `Lib/PadicHodge/WachKisin/KummerChartExists.lean`. See also the definition of `Lift.Chart` in `Lib/PadicHodge/WachKisin/HodgeFrame.lean`.

Let $q_0 = p^{f_0} - 1$, and choose a primitive q_0-th root of unity ζ in O (hence the initial enlargement of the coefficient ring). The source indexes the full set of units of $\mathbb{F}_{p^{f_0}}$ by the term *letters*; it is convenient to write $I = \mathbb{F}_{p^{f_0}}^{\times}$. A letter $i \in I$ indexes the summand of the Witt-vector ambient module on which multiplication by ζ on the lattice factor agrees with multiplication by the Teichmüller root $\omega(i)$ on the Witt-vector factor. The projector to this summand on $\mathcal{V}_T = W(\mathbb{C}_p^{\flat}) \otimes_{\mathbb{Z}_p} T$ is

$$e_i = \frac{1}{q_0} \sum_{j=0}^{q_0 - 1} \left(\omega(i)^{-1} \zeta \right)^j . \tag{23.7}$$

Because q_0 is prime to p this operator is integral. Orthogonality and idempotence follow by summation of a geometric progression over distinct roots of unity; $\omega(i) - \omega(j)$ is a unit for $i \neq j$. The sum of the projectors is the identity, and Frobenius sends the summand i to the summand i^p. Denote the permutation $i \mapsto i^p$ by σ_0 and put $i[-a] = \sigma_0^{-a} i$ for $a \geq 0$. The coefficient embeddings of O belonging to letter i are the $\gamma : O \to \mathbb{C}_p$ with $\gamma(\zeta) = \vartheta(\omega(i))$. Define

$$U_i = \{0\} \cup \bigcup_{\gamma(\zeta) = \vartheta(\omega(i))} U(\gamma), \qquad m_T(i, u) = \sum_{\gamma(\zeta) = \vartheta(\omega(i))} j_T(\gamma, u).$$

The zero in U_i permits crystalline extensions by a trivial lattice. The sum of $m_T(i, u)$ over all u and all i is $[K_O : \mathbb{Q}_p] \operatorname{rk}_O T$. In the arguments below it matters that the sum is over letters, whereas f_0 measures the size of the residue field extension: these are different integers. The projector formula, its Frobenius compatibility, and the definitions of the jumps attached to a letter are `letterProj`, `phi_pow_smul_mem_letterPart_iff`, and `jumpMultiplicity` in `Lib/PadicHodge/WachGamma/KisinModuleDefs.lean` and `Lib/PadicHodge/WachGamma/Letters.lean`; the enlarged weight sets are `KummerChart.letterJumps` in `New/Lift/ExponentTheorems.lean`.

We explain the choice of integral lattice used for this calculation. The source works with $W(\mathbb{C}_p^{\flat})$ over A_{inf} and an auxiliary element $\tau = [\epsilon^{1/p}] - 1$ of A_{inf}, where ϵ is the compatible tower of p-power roots of unity. Put $A_T = A_{\mathrm{inf}} \otimes_{\mathbb{Z}_p} T$ inside $\mathcal{V}_T$. An invariant form on T is an A_{inf}-linear functional $\mu : \mathcal{V}_T \to B_{\mathrm{dR}}$ satisfying $\mu(1 \otimes gv) = g\mu(1 \otimes v)$ for $g \in \Gamma'$; write $\mathcal{H}_T$ for the family of such forms. Let M_T be the source's submodule `kisinAinf`, consisting of elements in the τ-saturation of A_T such that

$$\mu(\varphi^a v) \in B_{\mathrm{dR}}^{+} \quad (\mu \in \mathcal{H}_T,\ a \geq 1).$$

For a more intrinsic way of specifying the finite-height module $\mathfrak{N}_T$ in M_T, impose the additional value-field condition in the source: $\mu(v)$ must belong to the chart's value field for all invariant

forms. This defines `kisinModule`; the value-fields and their Frobenius transforms are as in `Lib/PadicHodge/WachKisin/DeRhamValues.lean`. Let $M_T^{\mathrm{loc}} = M_T[1/\mathfrak{e}_0, 1/\mathfrak{e}_1, \ldots]$ inside $\mathcal{V}_T$, where $\mathfrak{e}_a = \varphi^a \mathfrak{e}$, and let $\mathfrak{N}_T^{\mathrm{loc}}$ be its points satisfying the same value-field condition. We do not identify the latter with an arbitrary scalar extension of an abstract Breuil–Kisin module; the proof uses the above integral model and its retraction. In particular it is important that the period-frame relations are preserved by certain $\mathfrak{S}$-linear retractions on its scalar coefficients.

Proposition 23.4 (Integral frame package)**.** Formal-source input. *Let T be admissible for the chart and let H_T bound the absolute values of its jumps. Then*

$$M_T = A_{\mathrm{inf}}\mathfrak{N}_T, \qquad \tau^{H_T} M_T \subset A_T, \quad \tau^{H_T} A_T \subset M_T, \quad \mathfrak{e}^{H_T}\varphi(\mathfrak{N}_T) \subset \mathfrak{N}_T.$$

Relations among value-field points remain valid after applying an $\mathfrak{S}$-linear map $A_{\mathrm{inf}} \to \mathfrak{S}$ to the coefficients. Moreover let $l, A, \delta \geq 0$ with $3e_0 \leq p^l$, $e_0 A \leq p^l \delta$. Then the value-field points of $\varphi^l(M_T)$, after multiplication by $p^{4\delta}$, are spanned over $\mathfrak{S}$ by value-field points invariant modulo $\mathrm{Fil}^A B_{\mathrm{dR}}$ under all Γ' when tested by all invariant forms. The invariant crystalline forms also test divisibilities in the frame.

These assertions are components of a single theorem, `Lift.KummerChart.JumpLattice.frame_theorem`, in `Lib/PadicHodge/WachKisin/FrameTheorem.lean`, with preparation in `PeriodMatrixFrame.lean`. The underlying results concern Breuil–Kisin lattices and comparison periods, but the last estimate is used with varying l, A, δ; it is not enough to cite finite height or crystalline comparison with unspecified constants. In the formal proof the period matrix is used to compare the natural A_{inf}-lattice with its Frobenius transports, and functionals arising from the invariant period frame bound orders on each Hodge-graded line. The package including the $p^{4\delta}$ estimate is a specific input whose concrete proof would be part of a longer expositional reconstruction. For the current section it will be important to retain this exact exponent. The other uses of the chart involve only the parameters e_0, δ_0 and the Hodge jump sets, rather than the matrix size in the period frame.

23.4 Exponent lattices and the cost of increasing depth

For integer-valued functions $c(a, u)$ on nonnegative integers a and allowed jumps $u \in U_{i[-a]}$, define the exponent lattice $M_T(i, c)$ to be the subset of M_T^{loc} in the letter i such that

$$\mu(\varphi^{-a} v) \in \mathrm{Fil}^{c(a,u)} B_{\mathrm{dR}} \tag{23.8}$$

for every a and every invariant form μ of type u at $i[-a]$. Here the type condition is $\mu(e_{i[-a]}(1 \otimes T)) \subseteq \mathrm{Fil}^{-u} B_{\mathrm{dR}}$. This is `withExponents` in `New/Lift/ExponentLattices.lean`. We call c the exponents of the point and call c' its depth if, for all $g \in \Gamma'$, $(g-1)v \in M_T(i, c')$. Multiplication by $\mathfrak{e}_a$ raises the exponent at a by one, and Frobenius transports indices $i \mapsto i^p$. In particular the exponents of level l are

$$\mathrm{st}_l(a, u) = \begin{cases} 0, & a < l, \\ -u, & a \geq l. \end{cases} \tag{23.9}$$

Larger exponents mean a smaller module. A useful operation on exponents is

$$\mathrm{raised}_s(a, u) = -u + s\mathbf{1}_{a=0}.$$

Thus raised_s increases only the exponent at the original Eisenstein element. Depth is distinguished from the exponents of a point itself; the former can be improved without requiring arbitrarily large orders on the nonzero constants to which the points project.

The basic numerical cost attached to a step (a,n) is

$$c_p(a,n) = v_p(n) + \delta_0 + 3a, \tag{23.10}$$

where we put $v_p(0) = 0$ in this formula only. One point of the construction is that $c_p(a,n)$ is independent of the rank. Its origin is the following assertion underlying the finite-difference operation.

Lemma 23.5 (One-step moving operator)**.** Formal-source input. *For $n \neq 0$, write $\mathfrak{e}_a = \varphi^a \mathfrak{e}$. There exist $g \in \Gamma'$, $\omega \in A_{\mathrm{inf}}^{\times}$ and $\alpha, X \in A_{\mathrm{inf}}$ with*

$$g\mathfrak{e}_a = \omega \mathfrak{e}_a, \quad \mathfrak{e}_a \mid gs - s \ (s \in \mathfrak{S}), \quad z\tau X = \omega^n - 1, \quad \mathfrak{e}_a \mid \alpha X - p^{c_p(a,n)}.$$

If $n = 0$, a sum of p^{δ_0+3a} elements of Γ' has invariant action on the constants modulo $\mathfrak{e}_a$. After the chart retraction these identities supply an additive endomorphism of the localized module raising depth from n to $n+1$, modulo invariant constants, with cost $p^{c_p(a,n)}$ on the quotient of an extension by its kernel.

The identities used here are `Lift.KummerChart.exists_depth_step`, `Lift.KummerChart.exists_depth_step_zero` in `New/Lift/ExponentLattices.lean`, and `exists_reaches_depth_succ` in `New/Lift/HodgeExpansion.lean`. To explain the cancellation for $n \neq 0$, suppose that the leading contribution of an invariant form at the divisor $\mathfrak{e}_a = 0$ is $\mathfrak{e}_a^n h$ modulo higher order, and that h is invariant to order zero. Taking the difference $g-1$ gives $(\omega^n - 1)\mathfrak{e}_a^n h$ modulo the next order. The congruence for αX then implies that the combination

$$p^{c_p(a,n)} v - \frac{\alpha}{z\tau}(gv - v)$$

cancels this leading contribution. When $n = 0$, an averaging operator is used instead. The formal identity for the step includes precisely the hypotheses which allow the quotient by $z\tau$ and the application of the coefficient retraction: it is not legitimate to divide by $z\tau$ in an arbitrary integral submodule. The estimate for $c_p(a,n)$ is obtained by lifting a unit under ϑ, the bound for the differentiated Eisenstein equation and the valuation of $\omega^n - 1$. The latter contributes the factor $v_p(n)$ because lifting the exponent in a multiplicative congruence gives a factor dividing the valuation of n. The additional term $3a$ is the allowance for the Frobenius divisor. This explains why subsequent costs consist of sums of valuations of ordinary integers.

Consider an exact sequence of integral lattices $0 \to T \to \widetilde{T} \to C_0 \to 0$ in which C_0 is trivial, described by an O-linear map $\alpha : C_0 \to Z_f(T)$. The underlying O-module of $\widetilde{T}$ is $T \oplus C_0$, with action

$$g(v,c) = (gv + \alpha(c)(g), c).$$

It is crystalline over the chart because each $\alpha(c)$ is a coboundary after extending to B_{cris}. Denote the natural projection by π. In letter i a point over $p^{\kappa}c$ means $v \in \mathfrak{N}^{\mathrm{loc}}_{\widetilde{T}}$ satisfying $\pi(v) = p^{\kappa} e_i(1 \otimes c)$ and having prescribed exponents and depth. The relation $\pi(gv - v) = 0$ means that its Galois differences lie in a kernel lattice. Strictness of the kernel for the exponent conditions, namely every point in the kernel of π satisfying the exponent conditions on the extension comes from a point of $M_T(i,c)$ and similarly for localized value-field points, is a *formal-source input*: `exists_inclKer_eq_of_projTriv_eq_zero` in `New/Lift/LongGap.lean`. Its proof uses extension of invariant forms from the sublattice to the extension while retaining their filtration degree, via `exists_extend_form` of the same file.

The following consequence of the moving operators is the useful form for bounds. Let S be a finite set of steps (a,n) containing those required to raise depth from ρ to ρ': for each a, u it contains every pair (a,n) with $\rho(a,u) \leq n < \rho'(a,u)$. Starting with a point over $p^{\kappa}c$

of depth ρ, one obtains a point over $p^{\kappa+\sum_S c_p(a,n)}c$ of depth ρ'. Its exponents improve to $\rho'(a,u)$ wherever its old exponents were positive, and to at least $\min(0,\rho'(a,u))$ on the other lines. In the precise statement the integers may be raised further by paying more powers of p. This is *formal-source input* `pointsOver_nonempty_of_costs`, `exists_reaches_deep` and `mem_withExponents_of_deep` in `New/Lift/LongGap.lean` and `New/Lift/HodgeExpansion.lean`. The induction is on the largest step index to be increased and its depth. Compose the operators of Lemma 23.5; when operator D has cost p^e, the difference $Dv - p^e v$ belongs to the kernel of π, so the next operator preserves the accumulated scalar. Thus the cost of a composition is the sum of the integer exponents. The crucial difference from a dimension argument is that the same operator raises the specified depth on every line of the pertinent jump, so multiplicities do not multiply this cost. The exponent improvement statement uses that a de Rham value whose Galois differences lie in a deep filtration can be made invariant modulo that filtration by subtracting a constant from the local field. Positivity of the prior exponent forces that constant to be zero; without positivity one retains at least the zero-filtration threshold. These are the precise properties needed in the later cup-product calculation.

We record the elementary valuation estimate controlling these operators. For integers $n \ge 0$ and c, set $I(n,c) = [-n,n] \cup [c-n, c+n]$ (integer intervals). If $n \le pM$ and $|c|+n \le H$ with $H \ge 1$, then for $d \ge 0$,

$$\sum_{j \in I(n,c)} (v_p(j) + d) \le (4M+2)\lfloor \log_p H \rfloor + (4n+2)d. \tag{23.11}$$

Indeed at most $4n+2$ integers occur and, in each interval of $2n+1$ integers, at most $2\lfloor n/p \rfloor + 2$ can be multiples of p. More sharply, an interval of length $2n$ contains at most $\lfloor 2n/p \rfloor + 1 \le 2M+1$ multiples if $n \le pM$; hence the two intervals together contain at most $4M+2$. For any nonzero j of absolute value at most H, $v_p(j) \le \lfloor \log_p H \rfloor$, and zero contributes no valuation by convention. This proves the bound. It is `sum_twoWindows_le` in `New/Lift/CokernelKillers.lean`. For steps with $a = 0$, if p divides no nonzero integer within distance n of $0, \pm c$ then the valuation contribution vanishes; in that case all the steps in these intervals cost only δ_0 each. With the rough count $8n+4$ of their integers the total cost is at most $(8n+4)\delta_0$. The formal-source derived assertion `pointsOver_nonempty_of_noMultipleNear` therefore improves a point at the original divisor at precisely this cost. In an unramified chart the cost is zero. This free improvement is the first of several places where the proof needs a gap in the congruence classes of large weights, not just a gap in their ordinary archimedean sizes.

23.5 Cokernel estimates: divisibility by p and by the Eisenstein element

For an extension side (T, C_0, α) as above let $J_l(i;s)$ be the ideal of A_{inf} defined as the sum of (s) and the ideal of those multipliers which send the letter-i part of $A_{\mathrm{inf}} \otimes C_0$ into the image under π of the extension's exponent lattice with exponents st_l. Thus, for a scalar $u \in A_{\mathrm{inf}}$, the condition $u \in J_l(i;s)$ means that one can choose $v \in A_{\mathrm{inf}}$ such that for every $c \in C_0$ an exponent-st_l point lifts $(u - sv)e_i(1 \otimes c)$. We say that $p^\kappa \mathfrak{e}^s$ kills the cokernel at i if

$$p^\kappa \mathfrak{e}^s \in J_0(i;0).$$

These are `liftIdeal` and `KillsCokernel` in `New/Lift/CokernelKillers.lean`. This definition is a statement about actual multipliers, not just equality after inverting p. Coefficient retraction allows the corresponding lifts to be taken in the localized value-field module. In

particular if $p^\kappa \mathfrak{e}^s$ kills, after dividing by $\mathfrak{e}^s$ one finds a point over $p^\kappa c$ having both exponents and depth raised$_{-s}$. This last assertion is *formal-source input* `pointsOver_nonempty_of_killsCokernel` in `New/Lift/LongGap.lean`.

Proposition 23.6 (Cokernel calculus)**.** Formal-source input. *For the ideals just defined, the following are valid.*

(i) *The level and letter are transported by Frobenius: in particular $p^c z^{pa} \in J_1(i^p; r_0)$ if and only if $p^c z^a \in J_0(i; r_0)$, for a fixed modulus r_0. If all allowed jumps are $\le n$, multiplication by $\mathfrak{e}^n$ sends $J_0(i; r_0)$ into $J_1(i; r_0)$.*

(ii) *If $p^c z^v \in J_1(i^p; r_0)$, $pa \le v$ and $c + v_p(v!) \le \kappa + v_p((pa)!)$, then $p^\kappa z^a \in J_0(i; r_0)$.*

(iii) *If $p^c z^a \mathfrak{e}^n \in J_1(i; r_0)$, then $p^{c+\delta_0 n + v_p(n!)} z^{a+n} \in J_1(i; r_0)$.*

(iv) *If all jumps on the side are $\le n$, for each $\kappa \ge 0$ there exists $a \ge 0$ with $(p-1)a \le e_0(n+\kappa)$ and $z^a \in J_0(i; p^\kappa)$ for all i.*

(v) *In the unramified chart and for $p \ge 5$, if all jumps are $\le (p-3)(\kappa + s)$ then $p^\kappa \mathfrak{e}^s$ kills the cokernel at every letter.*

All assertions occur in `New/Lift/CokernelKillers.lean`, respectively `p_pow_mul_z_pow_mem_liftIdeal_iff` and `mul_E_pow_mem_liftIdeal_one`, `mem_liftIdeal_zero_of_mem_one`, `mem_liftIdeal_of_mul_E_pow`, `exists_z_pow_mem_liftIdeal`, and `killsCokernel_of_isUnramified`.

We recall how the factorials arise. The finite-difference calculation from Lemma 23.5, at level zero or one, acts on $J_l(i; r_0)$ by the following rule. If $f \in \mathfrak{S} \cap J_1(i; r_0)$ and there is an identity

$$z\tau\mathfrak{e}(s + \beta f) = \alpha(gf - f), \qquad g \in \Gamma',$$

with the appropriate integral scalar coefficients, then $s \in J_1(i; r_0)$. It follows by taking the difference of integral lifts in the defining multiplier ideal, using a retraction to preserve the membership conditions. This rule is *formal-source input* `mem_liftIdeal_of_derivation`. Applied to $f = p^c z^v$, the rule lowers v by one up to $p^{v_p(v)}$ in the coefficient. Inducting from v down to pa gives the additive cost $v_p(v!) - v_p((pa)!)$; transporting the remaining power back by Frobenius proves (ii). Applied to $f = p^c z^a \mathfrak{e}^n$, the rule reduces n by one and increases the z-exponent, at step cost $\delta_0 + v_p(n)$. Their sum is $\delta_0 n + v_p(n!)$, giving (iii). These summations are independent of any rank. In the unramified case the polynomial identity $\mathfrak{e} = z + p$ permits one to combine the multipliers of (ii) with a binomial expansion of $\mathfrak{e}^s$. The resulting containment modulo $p^{\kappa+s+1}$ is upgraded to an actual containment by the elementary Nakayama argument: if $p^c \in J_0(i; 0) + (p^{c+1})$ then $p^c(1 - pq) \in J_0(i; 0)$ for some $q \in A_{\mathrm{inf}}$, and $1 - pq$ is invertible. Thus $p^c \in J_0(i; 0)$ and adding the higher modulus does not enlarge $J_0(i; 0)$. The quantitative minimum on a in (iv) and binomial comparison in (v) are the specific remaining cokernel estimates; they cannot be replaced by an unquantified assertion that some power of p annihilates the cokernel.

23.6 Cup-product vanishing from the three sides

Each pairing in (23.4) introduces three crystalline extension sides: the side Y carrying c, the side T carrying c', and the side $(Y \otimes T)^*(1)$ carrying the third class q. On the last side the quotient lattice is O with extension determined by q, and the allowed jumps, including zero,

are the numbers $-u-v-1$ with u and v running over the jumps on the first two sides. The integral equivariant pairing of this third extension with $Y \otimes T$ is

$$((\xi, s), u \otimes v) \longmapsto (p^M \xi + s\psi)(u \otimes v),$$

where p^M is the pairing modulus. Its equivariance follows directly from (23.4), since the extra term $p^M q(g)$ is exactly the change of ψ. On reduction it agrees with the imperfectly dual pairing. This is why controlling the first two extensions is not enough: the period computation must also supply a lift of a constant in the third extension.

Over the tilt consider $\mathcal{M} = W(\mathbb{C}_p^\flat) \otimes \mathbb{Z}/p^M(1)$ and let $\mathcal{M}_+$ be the image of $W(\mathfrak{m}_{\mathbb{C}_p^\flat})$. Call a pairing of Witt-vector components safe with cost κ if p^κ times its value lies in $\mathcal{M}_+$. The submodule $\mathcal{M}_+$ is stable under Frobenius and Γ. Different letters are orthogonal: for example if two vectors have labels i, j, then multiplying their value by $\omega(i)$ or $\omega(j)$ gives the same result, since the pairing is O-bilinear before applying the functional and reducing. For $i \neq j$ their difference is a unit. This is also true of the sum of their pairings, as required by `isSafe_sum`.

Here is the algebra of the safe-pair criterion. We use the elementary Artin–Schreier–Witt property of $\mathcal{M}$ with its small part: $\varphi - 1$ is surjective; if $(\varphi-1)z \in \mathcal{M}_+$ then z is congruent modulo $\mathcal{M}_+$ to a unique element of the constant module $\mathbb{Z}/p^M(1)$; and the intersection of the constants with $\mathcal{M}_+$ is zero. This property, including the convergence for Galois differences needed to get continuous cochains, is a *formal-source input*, `TateTwist.smallPart_spec` in `Lib/PadicHodge/WachGamma/KisinModuleDefs.lean`.

Suppose x, w in the first extension lie over the same constant $p^{e_1}c_0$, and y in the second lies over $p^{e_2}c_0'$. Write X, W, Y for their respective first components in the Witt-vector tensor modules of the kernels. Write a_g, b_g for the crystalline cocycles on these kernels, scaled by p^{e_1} and p^{e_2}. Thus the first component of $(g-1)x$ equals $D_g = gX - X + a_g$, that of $(g-1)w$ equals $D_g^W = gW - W + a_g$, and that of $(g-1)y$ equals $T_g = gY - Y + b_g$. Put $C = \varphi Y - Y$. Denote the pairing multiplied by p^{e_3} by P. We assume the following four expressions are zero in $\mathcal{M}/\mathcal{M}_+$, for all $g, h \in \Gamma'$:

$$P(D_g, gT_h), \quad P(W - X, T_g), \quad P(D_g^W, gC), \quad P(\varphi X - W, g\varphi Y - \varphi Y + b_g). \tag{23.12}$$

These are precisely the four safe-pair hypotheses in `Pairing.ScaledCupBounds.of_isSafe`, `New/Lift/CupProductBound.lean`. The last second argument is the first component of $(g-1)\varphi y$. Since the scaled cocycle values are constants for Frobenius it is also φT_g. The goal is to see that $P(a_g, gb_h)$ is a coboundary.

First define $k_0(g) = P(a_g, gY) - P(X, T_g)$. The graded Leibniz rule for the group-cohomology differential, which can equally be checked by expansion in degree one, gives

$$\mathrm{d}k_0(g, h) = P(a_g, gb_h) - P(D_g, gT_h).$$

This equality can be read in the quotient because P is equivariant by construction. By the first safety condition the last term is zero. It remains to choose the cochain with values in constants, rather than just the large Witt-vector quotient. Choose Z with $(\varphi-1)Z = P(W, C)$ by the Artin–Schreier–Witt property and set $k(g) = k_0(g) + gZ - Z$. In the quotient, direct expansion gives

$$\begin{aligned}(\varphi-1)k(g) &= P(a_g, gC) - P(\varphi X, gC - C)\\ &\quad - P(\varphi X - X, T_g) + P(gW, gC) - P(W, C)\\ &= P(D_g^W, gC) - P(W - X, T_g) - P(\varphi X - W, \varphi T_g).\end{aligned}$$

This is zero by the remaining three conditions. Hence each $k(g)$ is congruent to a unique constant $\kappa(g)$. The differential of k equals that of k_0, and $P(a_g, gb_h)$ is a constant. Injectivity of the constants into the quotient implies that $\mathrm{d}\kappa(g,h) = P(a_g, gb_h)$ as an identity of constants. Continuity of κ is part of the preceding formal-source small-part estimate: the lift formula involves finite congruence data, and sufficiently small changes under the Galois action contribute to the small part, so the unique constant representative is locally constant. This proves vanishing after restriction to the chart group Γ', and then in the cohomology group on Γ by corestriction, since its index is prime to p. We may further increase the total exponent $e_1 + e_2 + e_3$. This gives the claimed scaled cup product, and is the mechanism behind the formal theorem `Pairing.ScaledCupBounds.of_isSafe`. Notice that the third condition involving w is not redundant: it provides precisely the correction to make the cobounding cochain Frobenius-invariant modulo the small part.

The remaining challenge is to produce safe points with sufficiently small exponents e_1, e_2, e_3. The main safety test is a strict filtration inequality. Suppose v, w in the two kernels lie in exponent lattices of exponents $o(a,u), o'(a,v)$. Suppose that in the third extension one has a point over p^κ at the predecessor letter $i[-l]$, with exponents $o_2(a,u'')$, antitone as a function of the jump. The required numerical inequality is

$$o(a,u) + o'(a,v) + \mathbf{1}_{a \geq l} o_2(a-l, -u-v-1) > 0 \tag{23.13}$$

for each a and all the jumps u, v at $i[-a]$. Under these hypotheses multiplying either input by $z\tau$ makes its pairing safe with cost κ. The corresponding safe-pair deduction and cochain calculation are treated in Chapter 25; the source declaration is `SidePairing.isSafe_of_exponents` in `New/Lift/LongGap.lean`. It uses the equivariant third pairing above and the rule that multiplication of filtered values adds exponents. The target line $O(1)$ has jump -1; applying the frame comparison for that line produces the factor τ, and the positive valuation of z puts the result in $W(\mathfrak{m}_{\mathbb{C}_p^\flat})$. The strict inequality is indispensable, as safe is a condition of vanishing modulo the maximal ideal of the tilt, not merely integrality. A Galois difference of a point above a constant projects to zero and is divisible by the auxiliary $z\tau$ in the required localized kernel; this is how the depth inequalities for x, w, y imply (23.12).

One uses cuts on Frobenius cycles of letters. Let $I_0 \subset I$ be a subset meeting each cycle, $l(i)$ the distance to the first strictly preceding cut, and choose one leading side at each cut. The point at every other letter is transported from a preceding cut by Frobenius; at a cut one may choose directly constructed points on the leading side. The complementary choice gives two systems of points such that the Frobenius of one system is the other. More concretely the integer shifts are $d_P(i) = 0$ if i is a cut at which the side leads, and $d_P(i) = l(i)$ otherwise. For the complementary systems the identity

$$d_{\mathrm{lag}}(i) = d_{\mathrm{lead}}(i[-1]) + 1$$

uses minimality of the preceding cut. Consequently summing $\varphi^{d_P(i)}$ of the chosen points at $i[-d_P(i)]$ yields the points in the previous safe-pair criterion. Orthogonality of different letter summands reduces all four tests to the separate inequalities (23.13), with appropriate depth and minimum exponents of differences of points at a cut. The precise assertion with the additional depth inequality for the leading side is *formal-source input* `SidePairing.scaledCupBounds_of_points_at_cuts` in `New/Lift/LongGap.lean`; it concludes that an exponent $e_1 + e_2 + \kappa$ is sufficient. The optimization of the cut selection matters, because a Frobenius walk of length l introduces new filtration inequalities at all preceding divisors $\mathfrak{e}_a$.

23.7 The window estimates: the three ways to obtain safe points

For a sign $\epsilon \in \{1, -1\}$ put

$$I(\eta, \epsilon, b, W_0) = [-\eta, \eta] \ \cup \ \epsilon[b, b + W_0],$$

and write $\mathrm{NM}(p, n, b)$ for the condition that no nonzero integer within distance n of $0, b, -b$ is divisible by p. These two definitions are `IsInJumpWindow` and `NoMultipleNear` in `New/Lift/CokernelKillers.lean`. The estimates use opposite signs for the far windows of paired sides. The third extension has many possible jump sums, but for two opposite windows the sum of any pair of large jumps is near zero if both are far, and near a far window if precisely one is far. This is the basic reason that the window structure is stable enough for (23.13) to be useful.

Sparse thick letters. First consider opposite windows $I(\eta, \mp\epsilon_i, \beta_i, V)$ at letter i, where $\beta_i = b$ for thick letters and $\beta_i = G$ for thin ones, and $G + V \le \eta$. Let $W_0 \ge V + 4\eta$, and suppose that either all the signs are constant or the chart is unramified and $p \ge 7$. For an integer $J \ge 0$ assume that there cannot be three thick letters $i, i^{p^s}, i^{p^{s+t}}$ with $1 \le s, t \le J$. The following full set of size hypotheses is sufficient:

$$8e_0 \le p^{J+1}, \qquad 2e_0 b \le p^{J+1} X,$$
$$11(p^J + 1)W_0 \le b, \qquad 4(J + 1 + \delta_0 + \lfloor \log_p(2b) \rfloor)W_0 \le X, \qquad \eta \ge 1.$$

Under them every cup product of crystalline cocycles of the paired sides is killed, in the sense of (23.12), by $p^{10^4 X(J+1)}$. This is *formal-source input* `SidePairing.scaledCupBounds_sparse_thick` in `New/Lift/SparseThickLetters.lean`. One must use these hypotheses on both sides, not just on the distinguished block.

The combinatorial operation in its proof is concrete. Define a cut at i if the next J letters $i^p, \dots, i^{p^J}$ are thin. On a Frobenius orbit every letter has a preceding cut at distance at most $2J + 1$. Indeed, for $J = 0$ all letters are cuts. For $J > 0$, suppose every one of the preceding $2J + 1$ positions fails to be a cut. Failure at distance $J + 1$ yields a thick letter at distance a with $1 \le a \le J$. Failure at distance $a + J + 1 \le 2J + 1$ yields a thick one at distance $a + b$, $1 \le b \le J$. Failure at distance a yields a third thick letter at forward distance c, $1 \le c \le J$, from the first. This is a forbidden triple with gaps b, c. The formal implementation of this selection is the first part of `scaledCupBounds_sparse_thick`; the cut definition also gives the short thin run used in the bound on $J_0(i; s)$ of each extension. Along the run, multiplying by a suitable power of $\mathfrak{e}$ and transporting under Frobenius changes the required exponent of z by a factor p per step. Proposition 23.6(ii),(iii) gives the extra factorial and different costs at each stage. The number X in the hypotheses bounds those costs. The starting powers of z are obtained from Proposition 23.6(iv). At cuts the resulting points can improve their depth in the two windows, paying the bound (23.11) and summing over at most a bounded multiple of $J + 1$ steps. With budget κ satisfying $12(J + 1)X < \kappa$, the source obtains the points over $p^{25\kappa}$ on both sides and a third point over p^{κ}. Thus 51κ pays for the safe-pair criterion. Taking κ of order $(J + 1)X$ and allowing ample slack gives the displayed exponent $10^4 X(J + 1)$. The general cut lemma specifying both the permissible sign pattern and the gap of the large weight is `SidePairing.scaledCupBounds_of_cut_windows`; it is a further formal-source input. This is one of the technically new results needed for the proof, not a standard property of p-adic representations cited without the quantitative bound.

Two sizes over an unramified chart. The second general regime assumes $p \geq 7$ and $e_0 = 1$, $\mathfrak{e} = z + p$. The windows are again $I(\eta, \pm\epsilon_i, \beta_i, V)$ with $\beta_i = b$ or G, and $W_0 \geq V + 4\eta$, $\eta \geq 1$. If

$$5W_0 \leq G \leq b, \qquad b < Kp^2, \quad 2G < Kp, \quad W_0 \leq K,$$

then the exponent

$$10^4 K(\lfloor \log_p(2b) \rfloor + 1) \tag{23.14}$$

annihilates all imperfect cup products. This is *formal-source input* `SidePairing.scaledCupBounds_two_sizes` in `New/Lift/SparseThickLetters.lean`. It is used when p is large and a valuation budget is available, regardless of the sparsity of far-apart letters. Here the cut choice uses one Frobenius step and at i an additional order

$$s_i = \begin{cases} \max(0, \lfloor b/p \rfloor - 7K), & i^p \text{ thick}, \\ 0, & \text{otherwise}. \end{cases}$$

It is a cut if $s_i < \beta_i$. Some cuts require $\mathfrak{e}^{s_i}$ in the cokernel bound, rather than only p^κ; Proposition 23.6(v) is the point that makes these useful. When traversing the single large gap the necessary deep point is constructed by extracting a coefficient in an Eisenstein expansion, avoiding the large cost of paying for every intermediate jump. This explains the role of the long-gap lemma below. For the precise verification of the cup-product depth inequalities in this regime one uses the same `scaledCupBounds_of_cut_windows` with $J = 1$ and $n = K$.

The long-gap operation. The auxiliary statement used for these cuts is worth formulating. In an unramified chart suppose $l > 0$, $3e_0 \leq p^l$, p^κ kills the cokernel at $i[-l]$, $\mathfrak{e}^s$ kills at i, all allowed jumps at $i[-j]$ are $\leq G_1$ for $0 \leq j < l$, and put

$$N = s + G_1 \sum_{j=0}^{l-1} p^j.$$

If $N < A$ and $e_0 A \leq p^l \delta$, there is a point in the extension over $p^{\kappa+4\delta}c$ with level-l standard exponents modified at $a = 0$ by adding $\min(A, -u - N)$ on the line of jump u, and with depth modified at $a = 0$ by adding A independent of u. This is *formal-source input* `PairingSide.pointsOver_nonempty_of_long_gap` in `New/Lift/LongGap.lean`. A brief account of the mechanism suffices to locate the nontrivial ingredients: by starting with a point over the pure Eisenstein annihilator, one multiplies by a monic polynomial in $\mathfrak{e}$ of degree N to arrive in a Frobenius-transported module. In this step each Frobenius predecessor contributes the corresponding power of p to the degree; explicitly the polynomials for the unramified chart arise from $\varphi^j(z + p) = (\mathfrak{e} - p)^{p^j} + p$. Applying the $p^{4\delta}$ estimate of Proposition 23.4 expresses the resulting element as a linear combination of nearly invariant points. Expand the coefficients of that combination in powers of $\mathfrak{e}$ using $\mathfrak{S}/\mathfrak{e} \simeq \Lambda$. Comparing the coefficient of degree N gives a new point, now with exponents controlled by N rather than by the full far weight. The strict inequality $A > N$ ensures that the original leading coefficient survives the comparison. The precise invariant coefficient extraction is expressed by `JumpLattice.coeff_mem_constantsModE` in that file. Thus the exponent 4δ in the frame theorem and the geometric degree N are both visible in the cost of the long-gap operation.

Few low jumps. In the last regime it would be wasteful to multiply every class by the exponent needed for a small exceptional part. The module of crystalline cocycles of the distinguished side is used as the trivial quotient C_0, with its identity cocycle map. In an

unramified chart suppose $p \geq 5$, $\eta \geq 1$, and all paired jumps lie at each letter in the opposite windows $I(\eta, \pm\epsilon_i, \beta_i, W_0)$. Let A_0 be a nonnegative integer such that for every letter

$$\beta_i + W_0 + \eta \leq (p-3)A_0, \qquad 6(A_0 + \eta + W_0) < \beta_i, \qquad \mathrm{NM}(p, 6(A_0 + \eta + W_0), \beta_i). \tag{23.15}$$

Then the following estimate holds.

Theorem 23.7 (Length estimate for few low jumps)**.** Formal-source input. *Let T have the indicated second jump window and let $\mathcal{Q}$ be the crystalline lattices with the opposite first jump sets. Suppose that at every coefficient embedding the sum of the multiplicities of jumps between $-\eta$ and η is at most μ. Then*

$$\Delta_p(T; \mathcal{Q}) \leq 11 f_0 \left(\mu(A_0 + 1) + (W_0 + 1)\,\mathrm{rk}_O\, T\right).$$

Moreover $\Delta_p(T; \mathcal{Q}) = 0$ if neither of the original jump sets meets $[-A_0, 0)$, or if at every embedding the second jump set is contained in $[-\eta, 0]$.

This is `KummerChart.few_low_jumps` in `New/Lift/ExponentTheorems.lean`. Its bound involves the width of the far window and the number of low jumps separately, rather than the rank times A_0; this refinement is essential in summing the loss.

Here is the internal length argument behind it. For each letter, consider the constants which lift to points with exponent raised$_0$ and prescribed depth. To increase depth on one graded line of multiplicity $m_T(i, u)$ at the Frobenius divisor $\mathfrak{e}_a$, the index can be controlled by counting relations modulo p among its leading coefficients. For an unramified chart, if the depth conditions at $\mathfrak{e}$ can be advanced to the next Frobenius threshold without a nonzero multiple of p intervening, then the leading coefficients of $f_0 m_T(i, u) + 1$ points satisfy a nontrivial integral relation with one coefficient equal to 1 modulo the deeper submodule. The leading-term calculation and normalization of a relation by its coefficient of least valuation are made precise by `exists_comb_of_jumpMultiplicity_le` in `New/Lift/LongGap.lean`. Indeed the leading terms on a Hodge line of multiplicity m may be represented in at most m copies of the finite unramified coefficient module of rank f_0 over $\mathbb{Z}_p$; the rank count supplies a relation. The theorem includes the important compatibility of the retraction with this leading-term comparison, so it is a statement about the actual integral exponent modules, not just dimensions of their rational Hodge pieces.

For finite free modules the passage from this relation property to a length bound is elementary. If $p^d L \subset L'$ and every $q+1$ elements have an integral combination, one coefficient a unit, belonging to L', then the elementary-divisor decomposition of L/L' shows that at most q cyclic elementary divisors can be nonzero, and all their exponents are at most d. The length is therefore at most dq over $\mathbb{Z}_p$. Applying this estimate to each depth step and using subadditivity of relative length bounds the cost of changing the exponents from $c(a, u)$ to $c'(a, u)$ by a sum of step costs weighted by the corresponding multiplicity. More concretely, let n_i and r_i bound f_0 times the low and the total jump counts at i. The prescribed module of deep constants at letter i has colength at most

$$(2A_0 + 1)n_i + 3(W_0 + 1)r_{i[-1]} + 6A_0 n_{i[-1]}. \tag{23.16}$$

The precise modules and sign-dependent depths are `deepConstants`, `relLength_deepConstants_le`, and `exists_relLength_deepConstants_le` in `New/Lift/FewLowJumps.lean`. If the sign of the second side at i is positive its depth increases by $2A_0 + \eta$ at the single Eisenstein element. For the negative sign the prescribed depth is $\beta_i + W_0 + 2A_0$ at $a = 0$ and $W_0 + 1 - u + 2A_0 \mathbf{1}_{|u| \leq \eta}$ at $a = 1$, with $-u$ elsewhere; this is the depth underlying (23.16). The verification of (23.16) from the individual steps is a formal-source input; the source further

checks that these depths give safety through the cut lemma. Taking the intersection over all letters can only cost the sum of their colengths (inject the quotient of the intersection into the direct sum of quotients). Since $i \mapsto i[-1]$ is a permutation, summing (23.16) gives a bound by

$$11 \sum_i ((A_0+1)n_i + (W_0+1)r_i).$$

For $n_i = f_0 \sum_{|u|\le\eta} m_T(i,u)$ and $r_i = f_0 \sum_u m_T(i,u)$ this is at most $11 f_0 [K_O : \mathbb{Q}_p](\mu(A_0+1) + (W_0+1)\,\mathrm{rk}_O\, T)$ over $\mathbb{Z}_p$. Dividing by $[K_O : \mathbb{Q}_p]$ as in (23.1) proves the displayed normalized length bound from the integral depth statements. The zero-loss part of Theorem 23.7 does not follow by making this right side zero: when the low negative jumps are absent, pure Eisenstein annihilators and zero-cost depth steps give safe points at exponent zero. The second alternative is treated separately by the no-multiple condition and one Frobenius transport of the low-jump side. These are the two cases in `SidePairing.scaledCupBounds_of_killsCokernel`, `New/Lift/LongGap.lean`.

23.8 The tame type bound, uniform in the ramification degree

At the selected primes g_0, g_1 the local representations may only be potentially crystalline. Here the ramified inertial type has a special form. Let r_0, N_0' be bounds and suppose that for every inertial element s the polynomial class $\mathcal{K}(s,\cdot)$ contains $X^M - 1$, where $M = \mathrm{lcm}_{1\le j\le r_0}(p^j - 1)$, and contains a monic polynomial in $\mathbb{Z}_p[X]$ of degree at most N_0'. This is precisely `IsTameType` in `New/Lift/ExponentTheorems.lean`. The dual condition is formed by using polynomials at s^{-1}. Suppose now that the local field is $\mathbb{Q}_p$, the paired jump sets lie in opposite windows $I(\eta, \pm\epsilon, b, W_0)$, and

$$\kappa_0 = r_0^2(N_0'+1)^2, \qquad b + 6(W_0+\eta) + \kappa_0 + 2 < p.$$

Then multiplication by

$$p^{3\kappa_0 + 150(\eta + W_0 + 1)} \tag{23.17}$$

sends $Z_f(T)$ into the orthogonal part for all lattices in the opposite polynomial class. This is *formal-source input* `tame_types_exponent`, in the same file.

For the uniformity of (23.17), it is especially useful to recall its proof mechanism. All the roots of $X^M - 1$ are prime-to-p roots of unity; one may use a tame chart of ramification index M, which can be extremely large. But $M \le \prod_{j\le r_0} p^j \le p^{r_0^2}$, and the monic small-degree annihilator allows one to avoid paying for all that ramification. A generator s of the tame quotient acts on the chart coordinate as $z \mapsto \xi z$ for a Teichmüller root of unity of order M. If Q is a monic annihilator on the kernel lattice, then $Q(X)(X-1)$ kills the extension of a trivial lattice by that kernel. Indeed $s - 1$ maps the extension into the kernel on crystalline periods, and the polynomial Q kills that kernel. The corresponding operator on framed sections is the polynomial of the tame generator; its image is divisible by z in the integral frame. This last assertion is the specific formal-source input `TameChart.exists_z_smul_eq_sum_tameGen` in `New/Lift/SparseThickLetters.lean`. It is proved by applying the relation on the period matrix and using the divisibility test in the frame theorem. To lower a power z^j in the multiplier ideal one now uses this polynomial evaluated at ξ^j: unless its reduction vanishes, it is a unit and no power of p is needed. Where it vanishes the finite-difference step pays at most $\log_p M \le r_0^2$. The reductions of ξ^j for $0 \le j < M$ are distinct, hence there are at most $\deg Q + 1$ exceptional steps. Thus the cost at a side is bounded by $r_0^2(\deg Q + 1)$, independently of M itself. The third extension uses a polynomial of degree bounded by the product of the two degrees (plus the contribution of the trivial extension). The chosen κ_0 is sufficient for all three sides. This exact bound for the cokernel is `TameChart.killsCokernel_of_isKilledBy`. Together

with the no-multiple condition, guaranteed by the size hypothesis $b + 6(W_0 + \eta) + \kappa_0 + 2 < p$, the depth and safe-pair calculation gives (23.17); the factor 150 is a generous allowance for the three sides. The rank-independent cost of the earlier depth operators and the degree rather than ramification bound for the tame operators are distinct features of the proof.

In the selected-prime application, the root condition itself gives the small monic annihilator. A root of order dividing $p^j - 1$ lies in an unramified extension of degree dividing j, hence has monic minimal polynomial over $\mathbb{Z}_p$ of degree $\le j$. The product of at most N such polynomials is a monic polynomial of degree $\le Nr_0$ vanishing on the indicated roots. This explains `isTameType_of_roots` in `New/Selmer/RowBalanced.lean`, which is applied to the polynomial class of the finitely many columns. It contrasts with a direct estimate based on the order of the whole finite group at g_0, g_1.

23.9 Case analysis for the primes and the explicit local loss

The form in which the three crystalline-window estimates are combined is the following. We give the constants, since an asymptotic assertion without them would not specify the interface to the global proof. Temporarily write b_* and G_* for the two window sizes, h_* for the half-gap parameter, and let $Q > 0$ be the prime cutoff. Put

$$\lambda_p(h_*, b_*) = v_p\left((2b_*)!\left(\prod_{\max(b_*-30h_*^3,0)<j\le b_*+30h_*^3} j\right)^{h_*^3}\right). \tag{23.18}$$

Assume the size conditions

$$99(Q^2 + (W_0+1)Q + Q + W_0 + m) \le h_*, \qquad 99h_*^3 \le G_*,$$
$$99(G_*h_*^3 + G_*) \le b_*, \quad 2\log(4b_*) < m, \quad \delta < \eta, \quad 99\big((\eta + W_0)Q + \eta + h_*\big) \le h_*^3,$$

with $e_0 \le h_*$. Let the opposing jump windows be $I(\eta, \mp\epsilon, b_* + \delta, W_0)$ at thick embeddings and $I(\eta, \mp\epsilon, G_* + \delta, W_0)$ otherwise. Suppose that either all negative jumps on both sides satisfy $2u \le -h_*$ and the multiplicity of jumps of the second lattice T in $[-\eta, \eta]$ is at most μ per embedding, or the second jump sets consist entirely of jumps in $[-\eta, 0]$. At letters above a given embedding the sign and the thick indicator must be compatible whenever the embedding is thick or $Q < p$. If $p > Q$ the chart must be unramified with $p \ge 7$; at $p \le Q$ one may instead use constant signs and the sparsity of thick letters (no three successive thick letters with the two positive gaps at most $\lfloor \log_p(h_*^3) \rfloor$). This is the set of hypotheses used by `crystalline_defect_le` in `New/Selmer/CrystallineBound.lean`, where T is crystalline.

Under these hypotheses the formal-source local estimate is

$$\Delta_p(T; \mathcal{Q}) \le d_F\, 10^5 \left(\mathrm{rk}_O\, T \left\lfloor \frac{2\lambda_p(h_*, b_*)}{Q} \right\rfloor m + \mu\lambda_p(h_*, b_*)\right). \tag{23.19}$$

The same statement uses an arbitrary upper bound y for the parenthesized expression including 10^5 in its source formulation. Here are the quantitative subcases in the proof. Write $v = \lambda_p(h_*, b_*)$ and $K = \lfloor 2v/Q \rfloor$. By the factorial term and Legendre's formula,

$$\lfloor 2b_*/p \rfloor \le v, \qquad 2b_* < p(v+1), \qquad KQ \le 2v < (K+1)Q.$$

Also $\lfloor \log_p(4b_*) \rfloor < m$, since $p \ge 2$ and $2\log(4b_*) < m$.

- For $p \le Q$ the resource inequality $2b_* < p(v+1)$ forces a large bound K; after enlarging the small window to absorb the thin weight, use the sparse-thick exponent with $J = \lfloor \log_p(h_*^3) \rfloor$. The source compares $2(b_* + \delta)$ with $h_*^2 K$ and takes the depth budget compatible with K. The rank bound of Lemma 23.2 then suffices, at cost bounded by the first term of (23.19).

- If $p > Q$ and $v \geq h_*^3$, use the two-sizes exponent (23.14). The floor inequality, $p > Q$, and the size assumptions ensure the requisite bounds on $b_* + \delta, G_* + \delta$ in terms of K. The logarithmic factor is at most m, so by Lemma 23.2 the normalized loss is at most $2 \cdot 10^4 d_F \operatorname{rk}_O TKm$, absorbed by the first term.
- If $p > Q$ and $v < h_*^3$, the product factor in (23.18) has no multiple of p in its interval: a single such integer would contribute h_*^3 to v. The factorial term also implies that p is large compared with G_*, since $2b_* < p(h_*^3 + 1)$. These facts yield the no-multiple condition for the two sizes in Theorem 23.7 at $A_0 = 2(v+1)$. That theorem either gives zero loss (in the pure-low second case or when the negative weights avoid $[-A_0, 0)$) or gives a contribution of order $d_F(\mu(v+1) + (W_0+1)\operatorname{rk}_O T)$. In the nonzero case the large negative gap entails $h_*/5 \leq v$, and the numerical assumptions imply $W_0 + 1 \leq K$. Thus this too is absorbed by (23.19).

The exact floor and scale comparisons, including the sparse-thick budget in the first subcase, are part of the formal-source input `crystalline_defect_le`. Notice the structural improvement: the second term has the number of low jumps, not the rank, multiplying the large factorial valuation. Without that refinement, the eventual bound at rank-one string blocks and at the main row block would not have the required common budget.

23.10 Insertion of the global weights and the sum over all blocks

Return to the admissible datum $\mathcal{D}$. For this application put $h_* = h_2 = r^8$, $Q = kt$, and set

$$b_1 = N_1 \left\lceil \frac{a}{xN_1} \right\rceil, \qquad b_2 = N_1 \left\lceil \frac{a}{x_1 N_1} \right\rceil,$$

so $w_j(\sigma) = (b_1 \text{ or } b_2) + 2jh_2 + h_2$ at the two kinds of places. The main row has rank r' with $1 \leq r' \leq r$, and its block string has length $t' \leq t$. Its jumps at a complex embedding σ satisfy the alternative

$$h_2 \leq 2|u| \leq 2h_2^2 \quad \text{or} \quad |\epsilon_\sigma u + w_j(\sigma)| \leq t,$$

for one $j \leq m$; the number of jumps of absolute value at most $2h_2^2$ is at most $r'm/k$ per embedding. Here ϵ_σ is the sign of the CM embedding, $+1$ for the chosen CM type, -1 for its conjugate. The source also requires the sum of negative jump counts at $\sigma, \bar{\sigma}$ to be at most r'. These are the hypotheses `IsRowType` in `New/Selmer/DeterminantBound.lean`, used for the actual main row representations. The one-dimensional extra blocks are cyclotomic powers with jumps $c \in [0, t')$. After taking Tate duals the new large windows are opposite to the original ones, by (23.2), while the rank-one blocks have all jumps small and nonpositive. The local substitution in (23.19) is $b_* = b_1$, $G_* = b_2$, $W_0 = 4t$, $\eta = h_2^2 + 1$, and $\delta = 2jh_2 + h_2 - 2t$. This is why the input to the crystalline theorem is a count of low jumps rather than an absolute bound on all weights.

The local embeddings above p have sign and place classes related by the Galois action on roots of unity. If $g'g^{-1}$ acts on the prime-to-p part of the cyclotomic generator by the p^n power, the ratio of the infinite-place classes is $\epsilon_\sigma \epsilon_{\sigma'} p^n$ modulo the relevant prime-to-p divisor of N_1. This is `placeClass_ratio`. Squaring removes the sign, and the three-term avoidance in $\mathcal{S}$ gives `not_thick_of_short_arc`: for $p \leq kt$, if two thick embeddings are separated by a Frobenius arc of length $n \leq 27s$ then a third at interior distance $0 < i < n$ cannot be thick (where s is the cyclotomic tower exponent). Furthermore at $p \mid N_1 d$ two thick embeddings in one local Galois orbit have the same CM sign, by `cmSign_eq_of_thick`.

These are specific arithmetic *formal-source inputs*, `New/Selmer/ThickSpacing.lean`, used to verify the small-prime hypotheses of (23.19). At large primes away from the conductor the local extension is unramified and the larger-prime hypotheses apply. These separate facts explain why it is not sufficient merely to take a large set of thick places, without the progression-freeness built into $\mathcal{D}$.

Define now the local cost used in the global proof:

$$\mathrm{loss}_p = 10^6 \frac{m}{k}(b_1 v_p(g) + \lambda_p(h_2, b_1)), \qquad g = g_0 g_1. \tag{23.20}$$

At the selected split level primes use (23.17), with the degree bound for the polynomials derived from the good-type eigenvalues at those primes; the source uses $r_0 = r$ and $N_0' = (r+t)r$. At other primes use (23.19). With the preceding scale and spacing verifications the result is the following exact statement.

Proposition 23.8 (One-place block estimate)**.** Formal-source input. *In the above setting, let $v \mid p$ and let T_i be the lattices of the row plus string blocks (one row, and one for each $c \in [0, t')$). Let $\mathcal{Q}_v$ be the class of local lattices inherited from the chosen column representations, with the same jump restrictions and crystalline condition, or the stated tame polynomial condition when $p \mid g$. For all sufficiently large coefficient rings, a block belonging to this class has*

$$\Delta_p(T_i^*(1); \mathcal{Q}_v) \le [E_v : \mathbb{Q}_p] \operatorname{rk}_O T_i \frac{\mathrm{loss}_p}{q_i}, \qquad q_i = \begin{cases} 1, & i \text{ the row}, \\ t, & i \text{ a cyclotomic string block}. \end{cases}$$

This is `block_defect_le_at_place` in `New/Selmer/OnePlaceDefect.lean`. The verification of the numerical parameters and the precise embeddings-to-letters dictionary are included in that declaration. It uses the very large gap between $p = g_0, g_1$ and the weights in the tame case. The row-type assumptions for actual automorphic Galois representations must be proved by the other part of the congruence argument; the local proposition does not give their existence. In particular, the gain for string blocks arises from their special single low weight, not from an assertion that all rank-one blocks have zero crystalline defect.

Let $\mathrm{bloss}_{i,p}(M_i)$ be the normalized relative length of the module of obeying cocycles on the i-th block modulo those also satisfying the crystalline relations simultaneously at all $v \mid p$. By Proposition 23.3 and subadditivity of relative length on intersections,

$$\begin{aligned} \sum_i \mathrm{bloss}_{i,p}(M_i) &\le \sum_{v \mid p} [E_v : \mathbb{Q}_p] \left(r' + \frac{t'}{t} \right) \mathrm{loss}_p \\ &\le 4 r' f \, \mathrm{loss}_p. \end{aligned}$$

The last step uses $r' \ge 1$, $t' \le t$, and $\sum_{v \mid p} [E_v : \mathbb{Q}_p] = [E : \mathbb{Q}] = 2f$. It is valid for arbitrary choices of the moduli M_i. This is exactly `sum_blockLoss_le` in `New/Selmer/RibetElement.lean`. It gives the $4f \, \mathrm{loss}_p$ term in the row budget per rank.

Finally the summation of this term over all rational primes with logarithmic weights is especially transparent. Write $J_* = (2b_1)! \left(\prod_{\max(b_1 - 30h_2^3, 0) < j \le b_1 + 30h_2^3} j \right)^{h_2^3}$. For a finite set of primes P_0,

$$\begin{aligned} \sum_{p \in P_0} 4f \, \mathrm{loss}_p \log p &\le 4 \cdot 10^6 f \frac{m}{k} (b_1 \log g + \log J_*) \\ &\le 4 \cdot 10^6 f \frac{m}{k} (b_1 \log g + 2b_1 \log(2b_1) + 60 h_2^6 \log(b_1 + 30h_2^3)). \end{aligned}$$

Here unique factorization gives the first inequality, and the number of factors in the consecutive product is at most $60h_2^3$. On the admissible asymptotic family $b_1 \sim b$, $\log g \le 22 \log b$, and

h_2^6 is polynomial in x whereas b grows faster than every fixed power of x. Thus the cost is $O(f(m/k)b\log b)$. Its ratio to $B_0 = fa(\log b)^3/(mk)$ is

$$O\left(\frac{m^2}{x(\log b)^2}\right) = O((\log b)^{-2})$$

using $a \geq xb$, $m^2 \leq 4x$. The source proves the requisite fixed-budget inequalities for all admissible data, including the other part of the row budget, in `exists_isFitFamily`, `New/Selmer/FitTwistsExist.lean`. The displayed calculation identifies why the arithmetic parameters make even the high-weight local integral comparison affordable. At this point a loss of order $\log J_*$ is acceptable, but the large powers of r entering the rank and level cannot be permitted to multiply the factorial part indiscriminately. This is the significance of the rank-uniform depth method, the separate tame-polynomial estimate, and the low-jump colength bound in the proof.

Chapter 24

Integral period frames and the algebra of the local loss

We now examine the steps behind the local estimate in Chapter 23. There are two rather different kinds of assertion in that estimate. The first concerns ordinary finite-dimensional local cohomology, the second a uniform integral comparison problem. Ordinary weak admissibility, even if applied to every congruence representation separately, does not by itself supply the second assertion. In these pages we prove the elementary and classical cohomological deductions, develop the integral algebra used by the comparison, and isolate the additional comparison inputs in their specific matrix and module forms. This separation is important: it makes the normalization, the dependence on ramification, and the few-low-jumps saving intelligible independently of the notation of the implementation. Statements still used as formal-source inputs are identified as such; the discussion is not an independent proof of every theorem of the period-ring development.

24.1 The rank calculation and duality over a finite coefficient ring

Let $F/\mathbb{Q}_p$ be finite of degree d_F, let G_F be its absolute Galois group, and let T be a free O-lattice, of rank r, as in Section 23.1. Put $V = T[1/p]$. We use the convention that the jump of $O(1)$ is -1. In particular, a line with negative jump contributes one to the dimension of $D_{\mathrm{dR}}(V)/\mathrm{Fil}^0$. The dimensions of this quotient, calculated at the different embeddings, are the quantities that should control the rank of the cocycle lattice. It is worth giving the reason for the extra r in Proposition 23.1.

Proposition 24.1 (Rank of the finite-condition cocycle lattice)**.** *Suppose that V is de Rham and that the finite condition is the usual crystalline finite condition. Let $t_V = D_{\mathrm{dR}}(V)/\mathrm{Fil}^0 D_{\mathrm{dR}}(V)$. Then*

$$\mathrm{rk}_O\, Z_f(T) = r + \dim_{K_O} t_V.$$

Consequently, if the coefficient field splits the embeddings of F, and J_σ is the jump multiset at σ, then

$$\mathrm{rk}_O\, Z_f(T) = r + \sum_{\sigma : F \hookrightarrow K_O} \#\{u \in J_\sigma : u < 0\}.$$

Proof. The classical fundamental exact sequence for the finite condition gives

$$0 \longrightarrow H^0(F,V) \longrightarrow D_{\mathrm{cris}}(V) \xrightarrow{(1-\varphi,\mathrm{can})} D_{\mathrm{cris}}(V) \oplus t_V \longrightarrow H^1_f(F,V) \longrightarrow 0.$$

Here dimensions are taken over K_O (the Frobenius on the first factor is linear over K_O). This sequence, which is the usual Bloch–Kato cohomological application of the fundamental sequence of period rings, does not require the positive-dimensional vector space $D_{\mathrm{cris}}(V)$ to have its maximal possible dimension. Taking the Euler characteristic gives

$$\dim H^1_f(F,V) = \dim H^0(F,V) + \dim t_V.$$

The coboundary map $V \to B^1(F,V)$ has kernel $H^0(F,V)$, so that

$$\dim Z^1_f(F,V) = \dim H^1_f(F,V) + \dim B^1(F,V) = \dim t_V + r.$$

The finiteness theorem for local Galois cohomology of lattices implies that $H^1(F,T)$ is finitely generated over O, and $B^1(F,T)$ is a quotient of T. Thus $Z_f(T)$ is finitely generated and torsion free. A continuous rational cocycle on the compact group G_F takes values in some $p^{-n}T$, so it becomes integral after multiplication by p^n; the finite condition is unchanged by multiplication by this scalar. Hence $Z_f(T)[1/p] = Z_f(F,V)$. The rank calculation follows. Finally $F \otimes_{\mathbb{Q}_p} K_O = \prod_\sigma K_O$, and the dimensions of the separate filtered quotients count precisely the negative jumps. □

This argument uses the classical Bloch–Kato sequence, local cohomology finiteness and the meaning of de Rham Hodge jumps. It explains the rank computation in `Lib/PadicHodge/BlochKato/BlochKato.lean` and `New/Selmer/CrystallineRank.lean`; it does not replace the latter files' precise comparisons with the code's period-ring presentation.

There is a useful elementary way to prove the index bound behind the reduction to a dual finite condition. We give it since it is easy to miss the factors arising from coboundaries if one works entirely with classes. Put $M = T/p^aT$, $M^D = T^*(1)/p^aT^*(1)$, and $C = \mathbb{Z}/p^a\mathbb{Z}(1)$. The perfect finite pairing with target C is obtained by composing the O-valued evaluation with a generator of $\mathrm{Hom}_{\mathbb{Z}_p}(O,\mathbb{Z}_p)$ as an O-module. Indeed this dual is the inverse different $\mathfrak{D}^{-1}_{K_O/\mathbb{Q}_p}$ under the trace pairing. If d generates the inverse different, then $x \mapsto \mathrm{Tr}(dx)$ induces the required perfect pairing also after reduction. This assertion does not require that the ordinary trace pairing on O itself be perfect.

Let $Z = Z^1(F,M)$, $Z' = Z^1(F,M^D)$, with coboundary subgroups B, B'. By finite local Tate duality their cup pairing factors to a perfect pairing of Z/B and Z'/B' (with values in $H^2(F,C) \simeq \mathbb{Z}/p^a\mathbb{Z}$). Classical Bloch–Kato orthogonality gives the orthogonality of the images A, A' of the two finite-condition lattices. One justification passes first to the rational cup pairing on finite-condition classes; to pass back integrally, $H^2(F,\mathbb{Z}_p(1)) \cong \mathbb{Z}_p$ is torsion free. Here the coefficient functional can be understood on the entire tensor product of lattices before reduction. In the source the cup is read in $H^2(F,\overline{\mathbb{Q}}_p^\times)$. The inclusion $\mu_{p^a} \hookrightarrow \overline{\mathbb{Q}}_p^\times$ induces an injection on H^2 by the Kummer exact sequence and $H^1(F,\overline{\mathbb{Q}}_p^\times) = 0$ (Hilbert 90). Thus no orthogonality information is weakened by reading its image there.

Lemma 24.2 (Index by cocycle annihilation)**.** *Let $X' \subset A'$, and let $S \subset Z$ be a subgroup with zero pairing against X'. Then*

$$[S : S \cap A] \le [A' : X'].$$

Proof. Write $\bar{X}'$ for the image of X' in Z'/B'. The exact order of the annihilator of X' in Z is

$$|\operatorname{Ann}(X')| = |B| \frac{|H^1(F,M)|}{|\bar{X}'|}.$$

In particular it is at most $|B||B'||H^1(F,M)|/|X'|$. The local Euler–Poincaré formula for finite p-primary modules gives

$$|H^1(F,M)| = |M|^{d_F} |H^0(F,M)||H^0(F,M^D)|.$$

The two orders of the coboundary groups are $|M|/|H^0(F,M)|$ and $|M^D|/|H^0(F,M^D)|$. Since $|M| = |M^D|$, their product with $|H^1(F,M)|$ equals $|M|^{2+d_F}$. The rank calculation and the complementary jumps u and $-1-u$ give

$$|A||A'| = |M|^{2+d_F}.$$

For the equality here we use saturation, so that reduction of the crystalline cocycle lattice has precisely the order computed from its rank. We deduce $|\operatorname{Ann}(X')| \le |A||A'|/|X'|$. The subgroup $S + A$ is contained in $\operatorname{Ann}(X')$, so

$$[S : S \cap A] = \frac{|S+A|}{|A|} \le \frac{|\operatorname{Ann}(X')|}{|A|} \le [A' : X'],$$

as claimed. □

A practical consequence of Lemma 24.2 is the following. Suppose one has a second lattice of lifts of some reductions in A', and can prove that its intersection with the actual finite-condition lattice of the dual has normalized colength at most D. Pairing the first side against this intersection is sufficient to show that the first side can have at most D units of normalized defect. Thus the integral pairings need not lift to pairings of dual rational representations at all. What is essential is the existence of one large orthogonal submodule. This is why the special imperfectly dual pairings in Section 23.1 are natural, and why an estimate of relative length, rather than a mere assertion of nondegeneracy, occurs in the argument. The classical inputs in this reasoning correspond to `Lib/ClassField/LocalDuality/CrystallineDuality.lean`; the later applications additionally specify the period-valued cocycles and the coefficient functionals.

24.2 The coordinates in a Kummer chart

We explain the origin of several scalars that enter the local comparison. A letter E used just in this subsection and the ensuing frame calculation denotes a scalar generating a period-ring ideal; it is not the global number field. Let F have residue degree f_0 and ramification index e_0, and put $\Lambda = W(\mathbb{F}_{p^{f_0}})$. Work with the completed algebraic closure C_p of F. Its tilt $C_p^\flat$ has valuation normalized by the value on compatible systems of roots of p. Write

$$A = W(\mathcal{O}_{C_p^\flat}), \qquad A_{\max} = \widehat{A[\xi/p]}^{\,p}, \qquad B^+ = A_{\max}[1/p],$$

where ξ is a generator of $\ker\theta$ for Fontaine's homomorphism $\theta : A \to \mathcal{O}_{C_p}$. Let ϵ be a compatible system of roots of unity, with first root equal to 1 and subsequent primitive p^n-th roots. Set

$$\mu = [\epsilon] - 1, \qquad \tau = [\epsilon^{1/p}] - 1, \qquad \xi = \frac{\mu}{\tau} = \sum_{j=0}^{p-1} [\epsilon^{1/p}]^j.$$

This is a choice of generator of $\ker\theta$; the relations $\varphi(\tau) = \mu$ and $\theta(\tau) = \zeta_p - 1$ follow directly. Here B^+ is the presentation of a positive period ring using the completed ring generated by ξ/p, rather than by divided powers of ξ. In customary notation it is a maximal period-ring presentation. We do not identify its integral model with the divided-power model of B^+_{cris}. The precise admissibility and kernel comparisons for the presentation in the development are part of the period-ring inputs. This distinction matters in statements involving p-divisibility.

Choose a uniformizer ϖ_F with $\varpi_F = -p$ if F is unramified. Select compatible p^n-th roots of ϖ_F in $\overline{F}$. Their system defines $z^\flat \in \mathcal{O}_{C_p^\flat}$; put $z = [z^\flat]$. The map sending the formal variable z to this element is an embedding

$$S = \Lambda[[z]] \longrightarrow A.$$

Here the infinite series are evaluated using the topology defined by (p, z). To see injectivity, reduce a nonzero series modulo p after dividing by its greatest common power of p. The first nonzero coefficient of the remaining series has valuation zero in the tilt, whereas higher terms have strictly larger valuation. The non-Archimedean triangle inequality shows its value is nonzero. Successive division by p, together with p-adic separatedness, proves injectivity in general. The restriction of θ to S is onto $\mathcal{O}_F$: every element of that ring can be written as a convergent series in ϖ_F with residue representatives in Λ. Choose a unit $a(z) \in S^\times$ lifting the unit $-p/\varpi_F^{e_0}$ via this restriction and set

$$E = z^{e_0} a(z) + p. \tag{24.1}$$

Then $\theta(E) = 0$, and modulo p, E has the same valuation as ξ. Since $\ker \theta = (\xi)$, it follows that E is a unit multiple of ξ in A. In the unramified case we can take $a(z) = 1$, so $E = z + p$. Define $v \in A^\times$ by $E\tau v = \mu$. The same ring $A_{\max}$ is obtained using E/p as generator before completion. Formula (24.1) implies $z^{e_0}/p \in A_{\max}$, a fact used repeatedly in the denominator estimates.

There is a useful supply of Galois elements acting on z. For each integer $m \geq 0$ and $k \in \mathbb{Z}_p$, an element of G_F sends the chosen p^m-th root of ϖ_F to its product with the chosen p^m-th root of unity to the k-th power, because $X^{p^m} - \varpi_F$ is Eisenstein and hence irreducible over F. Fix k and let m vary. The sets of elements giving the desired action at level m are nonempty compact subsets of G_F, nested by compatibility. Their intersection is nonempty. For any element g_k in this intersection,

$$g_k(z) = [\epsilon]^k z.$$

This applies to the ordinary G_F; when a chart is defined with additional allowed-tame-action conditions, the chart existence statement has to supply the corresponding subgroup. It is the exact chart of the source, including its subgroup conditions, that is used in the later formal-source frame theorem.

Put $E_j = \varphi^j(E)$ for $j \geq 0$. Frobenius acts on S by the usual Frobenius on Λ and $z \mapsto z^p$. Evaluation at φ^{-j} followed by θ sends z to the selected p^j-th root of ϖ_F and acts on Λ by an automorphism. It is thus onto the integers of the extension F_j generated by this root. Its kernel is generated by E_j. For example, this follows by Weierstrass division: E_j is a unit times an Eisenstein polynomial of degree $e_0 p^j$ over Λ. Its quotient ring is finite free of that rank over Λ and has the indicated uniformizer and residue field. Equivalently this evaluates the quotient ring as $\mathcal{O}_{F_j}$. In particular E_j generates a prime of S. If $j \neq k$, evaluating E_k by $\theta\varphi^{-j}$ gives two terms in (24.1) with unequal valuations, since the valuation of the first term is p^{k-j} while the value of p is one. Their sum does not vanish. Therefore the primes (E_j) are distinct. In the de Rham filtration, $\varphi^{-j}(E_j)$ is a generator of the first filtration step, while $\varphi^{-j}(E_k)$ for $k \neq j$ is a unit. In particular estimates at different j can track different prime factors in the coefficient ring S.

Two more elementary nonvanishing observations are useful. The evaluation of z in each F_j is nonzero; and $\theta\varphi^{-j}(\tau) = \zeta_{p^{j+1}} - 1 \neq 0$. Thus the powers of z and the powers of τ in an inequality cannot simply be replaced by powers of a single E_j. A loss expressed in these scalars is first a statement about integral comparisons; only later, along with the factorial bounds and truncation modulo p^κ, does it become a bound in p-adic length. The chart existence and these division calculations are found in `Lib/PadicHodge/WachKisin/KummerChartExists.lean`, `Lib/PadicHodge/WachKisin/ChartScalars.lean` and `Lib/PadicHodge/WachKisin/DeRhamValues.lean` (notably `exists_kummer_chart`, `Chart.weierstrass_division`, `frobE_dvd_iff`).

24.3 A coefficient retraction and why it is used

The period-ring calculations often need to descend a relation over B^+ to relations over S. It is not legitimate simply to regard B^+ as a power-series ring over S, nor to assume that a basis of its residue field is a finite basis of the period ring. An adic coordinate argument provides the appropriate substitute. We give the algebraic details, as they also clarify why the integral comparison need not have a loss proportional to the number of entries of a matrix at every stage.

Lemma 24.3 (Adic coordinates)**.** *Let $S = \Lambda[[z]]$, $\mathfrak{m} = (p, z)$, and M an S-module. Assume p acts injectively on M and z acts injectively on M/pM. Choose lifts $b_i \in M$ of a basis of $M/\mathfrak{m}M$ over $S/\mathfrak{m}$. There are S-linear coordinate maps $\lambda_i : M \to S$ with $\lambda_i(b_j) = \delta_{ij}$ having the following approximation property: for each $x \in M$ and each n, a finite sum $\sum_{i\in J} \lambda_i(x)b_i$ approximates x modulo $\mathfrak{m}^n M$, and all but finitely many coordinates are in $\mathfrak{m}^n$. If M is $\mathfrak{m}$-adically separated and has finite rank over S, it is finite free.*

Proof. Write $P = \bigoplus_i Se_i$ and $\beta : P \to M$ for $e_i \mapsto b_i$. First, for each $n \geq 1$,

$$\beta^{-1}(pM + z^n M) = pP + z^n P.$$

Modulo pM, if a finite linear combination belongs to $z^n(M/pM)$, its reduction modulo $z(M/pM)$ vanishes. By the independence of the chosen residue basis each coefficient reduces to a multiple of z. Divide the coefficients by z and use its injectivity on M/pM to reduce n by one. This proves the assertion by induction. Next we claim that $\beta^{-1}((p^m, z^n)M) = (p^m, z^n)P$. The case $m = 1$ was just proved. If $\beta(x) = p^m y + z^n w$, the case $m = 1$ permits us to write $x = px_1 + z^n x_2$. Subtracting the second term gives

$$p(\beta(x_1) - p^{m-1}y) = z^n(w - \beta(x_2)).$$

By reduction modulo p, the expression in parentheses on the right is divisible by p. Cancel p (injective on M) and apply induction. In particular

$$\beta^{-1}(\mathfrak{m}^{2n}M) \subset \beta^{-1}((p^n, z^n)M) = (p^n, z^n)P \subset \mathfrak{m}^n P.$$

For each $x \in M$ choose a finite combination $a_1 \in P$ with $\beta(a_1) \equiv x \bmod \mathfrak{m}M$. Given an approximation modulo $\mathfrak{m}^k M$, express its error as a finite sum of the monomials $p^i z^{k-i}$ times elements of M, approximate these elements modulo $\mathfrak{m}M$, and add the correction to obtain an approximation modulo $\mathfrak{m}^{k+1}M$. The displayed inclusion shows that the resulting finite-support coefficient vectors are Cauchy in the $\mathfrak{m}$-adic topology of P and are independent of the approximations in its completion. Taking the coordinate limits in the complete ring S defines the maps λ_i. Uniqueness of the limit also proves S-linearity. For a fixed n, the coordinate vector differs modulo $\mathfrak{m}^n$ from a finite-support approximate vector, which gives the asserted approximation property. To justify the first approximation property with the actual limits, note that a finite-support approximation modulo $\mathfrak{m}^{2n}M$ has coefficients equal to the limits modulo $\mathfrak{m}^n$; replacing them by the limits changes the approximation only by $\mathfrak{m}^n M$.

Finally the coordinate maps prove the linear independence of the entire family (b_i) over S, so if M has rank at most r, at most r of the b_i exist. For any x, the element $x - \sum_i \lambda_i(x)b_i$ belongs to every $\mathfrak{m}^n M$ by the approximation property. By separatedness it is zero. Thus the b_i are a finite basis. □

For $M = A$, the hypotheses of Lemma 24.3 hold. The module is p-torsion free, and modulo p it is the tilt's valuation ring, in which z is a nonzero element of positive valuation. To

check separatedness, if x lies in every $(p,z)^nA$, reduction modulo p shows it belongs to every $z^n(A/pA)$, hence $x = px_1$. For each n the inclusion $\mathfrak{m}^{2n}A \subset (p^n, z^n)A$ and injectivity of z^n modulo p show that $x_1 \in (p^{n-1}, z^n)A$. Therefore x_1 too lies in every power of $\mathfrak{m}$. Iteration shows that x is divisible by every power of p, hence vanishes. Choose the residue basis to contain the class of 1, and denote this coordinate by λ_0. It gives a retraction $A \to S$ of the scalar inclusion. It also preserves all divisibilities by elements of S: $\lambda_0(sx) = s\lambda_0(x)$.

We next explain its extension to the positive period ring. On the dense subring $A[E/p]$ the prescription

$$\widehat{\lambda}\left(\sum_{j=0}^{n} a_j (E/p)^j\right) = \sum_{j=0}^{n} \lambda(a_j)(E/p)^j$$

is well-defined, since any equality becomes an equality over A on multiplication by p^n. It preserves $A[E/p]$ and is p-adically continuous there. It consequently extends on completion to $A_{\max}$ and, by inverting p, to B^+. (The same argument works for any S-linear map $A \to A$ rather than the particular coordinates.) For manipulating A itself inside $A_{\max}$, the precise divisibility criterion used in the source is

$$a \in A, \quad p^m \mid a \text{ in } A_{\max} \quad \Longleftrightarrow \quad a \in (p, E)^m A. \tag{24.2}$$

The forward direction of this integral criterion is significant: it is not a criterion for divisibility by p in B^+, where p is invertible. The elementary reverse direction follows because both generators of the ideal are divisible by p in $A_{\max}$. The converse, called `Amax.mem_span_pow_of_p_pow_dvd` in `Lib/PadicHodge/PeriodRings/AmaxStructure.lean`, is an integral period-ring structure input. The adic-coordinate argument above follows the lemmas `Lift.exists_adicBasis` and `free_finite_of_rank_le` in that file. In particular the argument concerns coordinates of arbitrary A-elements, not just the values of an element in the scalar subring.

Let S_B denote the bicommutant scalar ring used in the chart. In concrete terms its elements are those scalars $x \in B^+$ such that multiplication by x commutes with each extended S-linear endomorphism of the type just constructed. This is a subring containing $S[1/p]$, and if an element is a unit of B^+ its inverse also has the commuting property. Its utility is the identity

$$\widehat{\lambda}(xy) = x\widehat{\lambda}(y) \qquad (x \in S_B,\ y \in B^+). \tag{24.3}$$

Thus matrix equations whose coefficients belong to S_B can be transformed by arbitrary extended coordinate maps. The coordinates themselves give elements in the bicommutant: evaluate them first on the dense set $A[E/p]$, on which the result is in $S[E/p]$, and then use continuity and inversion of p. The stronger properties used in the frame theorem also concern Frobenius and the action of the subgroup fixing both the scalar chart and its theta values. One needs the extension of (24.3) to fields of fixed and de Rham values at different inverse Frobenius evaluations. These properties are not consequences merely of S-linearity; they are the specialized chart-scalar inputs in `Lib/PadicHodge/WachKisin/ChartScalars.lean` and `Lib/PadicHodge/WachKisin/DeRhamValues.lean`. They allow both a reduction to coefficients in S and comparison between different E_j. Without them, a lattice over the very large ring A would not yield the required finite-rank module over S.

24.4 From crystalline periods to an integral frame

We distinguish the elementary manipulations of a period matrix from the integral existence assertion about such matrices. Let T first be a $\mathbb{Z}_p$-lattice of rank d. Over the positive period ring localized at μ, choose a matrix M whose columns are a basis of invariant crystalline

periods, relative to an integral basis of T, in the crystalline case. The twisting convention determines whether this is read in T or in the space of linear forms on T; here we fix the convention that the columns are periods and Galois action on them is the diagonal action. Assume all jumps lie in $[-a, b]$ with $a, b \geq 0$. Write $h = a + b$.

Over B^+_{dR}, a filtration-adapted matrix factorization has the following shape:

$$MY = W_0 \operatorname{diag}(t^{k_1}, \dots, t^{k_d}), \qquad -a \leq k_i \leq b.$$

Here t is a generator of its first filtration step, W_0 is invertible over B^+_{dR}, and Y is a change of crystalline basis compatible with the filtration at the chosen embedding. The statement at other embeddings and inverse Frobenius evaluations gives these inequalities at each of the scalars E_j from Section 24.2. A bound on pole orders is plainly preserved by any invertible matrices over a discrete valuation ring. If W_0 and the power of t are separated as above, multiplying M by t^a and multiplying M^{-1} by t^b puts both in matrices over the positive de Rham ring at that evaluation. Returning from these pointwise pole bounds to the positive period ring itself is a separate assertion, of the type proved using a fundamental-kernel/divisibility calculation in the period-ring development. It is not simply an intersection of finitely many de Rham rings inside an ordinary number field.

In the notation of the chart, the normalized matrix in the source is

$$U = (E\tau)^a M, \qquad UW = (E\tau)^h I_d,$$

where $U, W \in \mathrm{Mat}_d(B^+)$. It satisfies

$$\varphi U = v^a U F_1, \tag{24.4}$$

where F_1 is invertible over B^+ with entries in the scalar ring S_B. Indeed, wherever the crystalline invariant comparison applies, $\varphi M = MC$ for a matrix C of fixed-period coefficients. Since $\varphi(E\tau) = \varphi(E)\mu = (E\tau)v\varphi(E)$, the formula has $F_1 = (\varphi E)^a C$. The field in which C takes its values, and its corresponding membership in S_B, are part of the invariant-ring comparison. The Galois transformation law for U is inherited from that of M after accounting for the scalar ratio $g(E\tau)/(E\tau)$. The normalized period matrix obeys one more important condition: extended coordinate endomorphisms preserve the constraints $Uy \in E^h(B^+)^d$. These are the requirements of `IsPeriodMatrix`, in `Lib/PadicHodge/WachKisin/CrystallinePeriodMatrix.lean`; the step from a crystalline matrix to this package is `IsChartOf.exists_isPeriodMatrix`. We use this package as a formal-source input when the integral pole and fixed-scalar comparisons have not been independently established.

Here is the module to which adic algebra is applied:

$$P = \{x \in A^d : \ \exists y \in S_B^d, \ Uy = E^h x\}. \tag{24.5}$$

The reason for requiring $x \in A^d$ is to obtain an integral Galois transformation matrix after selecting an integral basis. This is stronger than requiring the relation over B^+ alone. The two regularity conditions needed by Lemma 24.3 are natural: if px belongs to P for $x \in A^d$, one divides its witnessing y by p in S_B to see that $x \in P$. If $zx \in pP$ for $x \in P$, reduction in A^d/pA^d and injectivity of z show that $x = px'$ for $x' \in A^d$; the first observation gives $x' \in P$. Thus P is p-saturated inside A^d, and z is injective on P/pP. Since P is a submodule of the separated module A^d, it is separated. To conclude finite freeness from Lemma 24.3, one still needs to bound its rank over S; note that the rank of A over S is certainly not one. The period-matrix coordinate preservation is used precisely to show that any more than d independent S-directions would violate the period relation together with the extended coordinate identities. This rank bound is part of `IsPeriodMatrix.exists_basis` in

`Lib/PadicHodge/WachKisin/PeriodMatrixBasis.lean`. Invoking it gives a free S-module of rank at most d inside A^d.

In passing from this module to the final frame, it is important to obtain saturation not just with respect to p, but with respect to scalars preserving the entire system of period equations. The following schematic linear algebra records what is needed. Suppose Q is a square basis matrix for P, and $UZ = E^h Q$ with Z over S_B. If for a row $l \in A^{1\times d}$ and an element $q \in A$ not divisible by p we have $lQ = qr$ with $r \in S^{1\times d}$, one needs divisibility $q \mid lx$ for all appropriately allowed solutions $Uy = E^h x$, even beyond $y \in S_B^d$. Such a statement is not a result of row-reduction over B^+: it concerns divisibility in A. The comparison uses a Frobenius contraction argument (`exists_Ainf_eq_of_contraction`) and the extended-coordinate saturation lemma in `Lib/PadicHodge/PeriodRings/BdR.lean` and `Lib/PadicHodge/WachKisin/PeriodMatrixBasis.lean`. Also the assertion that Q is square and that p does not divide $\det Q$ is not an automatic consequence of finite freeness of P. That assertion uses a minimal-support and fixed-coefficient argument in `IsPeriodMatrix.exists_frame` in `Lib/PadicHodge/WachKisin/PeriodMatrixFrame.lean`. We isolate both as formal-source inputs.

Once these statements hold, several matrix consequences are elementary. In particular the adjugate of Q, applied row by row, has $(\operatorname{adj} Q)Q = (\det Q)I_d$. The stronger divisibility property yields $x \in QA^d$ for allowed solutions x. Taking $Uy = (E\tau)^h e_i$ using the i-th column of W produces $x = \tau^h e_i$, and hence a matrix L over A with

$$QL = \tau^h I_d. \tag{24.6}$$

The comparison furthermore obtains $U = QY$ with $Y \in \mathrm{Mat}_d(B^+)$. If the coordinate matrix is Z as above, then cancellation of Q yields $YZ = E^h I_d$. Thus the determinant-clearing factor of the matrix Z in every rational comparison is exactly E^h, not E^{dh}. The distinction between the identity and the much weaker assertion on determinants is crucial. When only the determinant bound is used, applying an inverse matrix through the adjugate can cost d times the pole bound on each independent direction. Instead the right inverse after multiplication by E^h gives the pole bound per direction. Relation (24.6) and the analogous period identity describe how the frame lies between two scalar multiples of the integral lattice. A reader should not infer from this that a Wach module or a Kisin module over an arbitrary coefficient ring has been proved to be categorically equivalent to every representation in the argument. What is used here is a particular simultaneous frame bound, in a fixed chart, with its Galois and Frobenius relations and its uniform scalar denominator.

For clarity we record one consequence of a frame when estimating divisibility of coordinates. Let $Z'Z = E^h I_d$, with entries of both matrices in a comparison ring R. If Zx is known to be divisible by q for some $q \in R$ acting injectively, then $E^h x = Z'(Zx)$ is also divisible by q. If q and E are relatively prime in the sense that the image of E acts injectively on R/qR, this implies $x \in qR^d$. More generally the conclusion records a possible loss of E^h, not a loss of a determinant. The same elementary calculation works after Frobenius, with E_j in place of E. When the divisibilities are tested by evaluation of period-valued linear forms, the additional ingredient is the chart-coordinate comparison which identifies such evaluations with Zx. In the frame theorem this evaluation has, in symmetric weight range $[-a, a]$, the form of a factor E_j^a times a period-valued linear form evaluated on $\varphi^j(n)x$, where n denotes the frame. Substitution of n using the normalized matrix U proves the assertion over B^+; the divisibility conclusions back in A or S use the integral saturation. The source result `frame_theorem` packages these assertions in `Lib/PadicHodge/WachKisin/FrameTheorem.lean`. They, rather than rational filtered-module theory alone, are the starting point for the integral approximation below.

To spell out the precise matrix identity in symmetric weight range, use a for the absolute jump bound and $h = 2a$. Let n be the frame with columns $\tau^{-a}Qe_j$ in the ambient localized module. Relation (24.6) gives

$$\tau^a n = Q, \qquad nL = \tau^a I_d.$$

In particular $\tau^a nA^d \subset A^d$ and $\tau^a A^d \subset nA^d$. The Frobenius relation delivered by the integral frame computation is $\varphi Q = v^a Q B_1$ for a matrix B_1 over S. Here $E\tau v = \mu = \varphi\tau$. Therefore

$$\varphi n = \mu^{-a} v^a Q B_1 = E^{-a} n B_1.$$

This proves the stated finite-height inclusion for the frame span once it is identified with the prescribed Kisin lattice. The inclusion is not multiplied by the rank and holds simultaneously at every letter.

Now substitute $U = (E\tau)^a M$ in $UZ = E^{2a}Q$, obtaining $MZ = E^a n$. Apply Frobenius m times, then the comparison to B_{dR} and an arbitrary crystalline linear form λ; write ψ_m for this embedding of the period coefficients after Frobenius evaluation, and let $\mathrm{val}(\lambda)$ be the row of its values at the integral basis of the original lattice. For $x \in A^d$ the identity becomes

$$E_m^a \lambda(\varphi^m(n)x) = \mathrm{val}(\lambda)\, \psi_m(M)\, \varphi^m(Z)x, \tag{24.7}$$

where the scalars on each side are mapped to the common evaluation ring. Since M is invertible and its columns are invariant, the rows of its inverse yield crystalline forms whose values on the invariant basis give the coordinate rows. For these forms the right side becomes the separate coordinates of $\varphi^m(Z)x$. Suppose all of their values on $\varphi^m(n)x$ lie in a given B^+-submodule of the target ring. Multiplying by E_m^a preserves the submodule. Multiplying the rows by the entries of $\varphi^m(Z')$ and adding, with $Z'Z = E^{2a}I$, shows that every $E_m^{2a}x_j$ belongs to that submodule. For an ideal-principal test modulo a scalar d_0 this is the origin of the divisibility claim in the frame theorem: if division by E_m is allowed on the target quotient and membership in d_0 tests divisibility back in the integral coordinates, one obtains $x \in d_0 A^d$. The doubled exponent $2a$ originates in the two-sided jump bound, not in a matrix determinant. Identifying the span of the frame with the intrinsic module defined by all Hodge inequalities uses the simultaneous evaluations, the criterion for τ-divisibility from the zero theta values at positive Frobenius powers, and the coefficient-retraction scalar comparison. These identifications in `FrameTheorem.lean` remain integral comparison inputs; (24.7) is their explicit linear algebra.

24.5 The radius estimate and its fourth power

We can now explain the quantitative part of the integral radius lemma. The strength of the uniform period comparison is not merely that a suitable denominator exists for each T separately. Its denominator can be chosen as a scalar depending only on the chart and a range of weights; one avoids a choice depending on the rank or on a basis of T. The first half of the argument is a formal-power-series construction and the second half is a linear algebra calculation with truncated series. We develop the latter in detail.

Choose an integer ℓ_0 such that $p^{\ell_0} \geq 3e_0$, and put

$$X = z^{p^{\ell_0}}/p^2 \in A_{\max}.$$

Since $z^{e_0} \in pA_{\max}$, we have $z^{p^{\ell_0}} \in p^3 A_{\max}$ and $X \in pA_{\max}$. Thus formal series in $S[[X]]$ can be evaluated in $A_{\max}$ by p-adic completeness. The radius lemma applies to a matrix Z over

S_B with $Z'Z = E^h I$. Its transformation for a chart element g with $g(z) = [\epsilon]z$ is $gZ = ZG$, with G integral over A. It establishes a factorization

$$\varphi^{\ell_0} Z = E_{\ell_0}^{b}\, C\, V, \tag{24.8}$$

where C is fixed by the fixer of Λ and V comes from an invertible matrix over $S[[X]]$ evaluated at the above argument. In particular V, V^{-1} are integral over $A_{\max}$. The integer b is a parameter in the radius lemma. To see the mechanism of this factorization, one first multiplies Z by an appropriate fixed matrix B to achieve $BZ = I + zW$. The finite-difference identity in the chart promotes $gZ = ZG$ to the divisibility of $G - I$ by $z\tau$. A recursive approximation of BZ then expands its coordinates in powers of z. At the n-th step the division of the difference of z^{n+1} by $z\tau E$ produces the factor $p^{1+v_p(n+1)}$, by the formula for $[\epsilon]^{n+1} - 1$. The accumulated cost is bounded by

$$\sum_{j=1}^{n} (1 + v_p(j)) = n + v_p(n!) \le 2n.$$

The residual equation and the coefficient retraction of Section 24.3 select the coefficients in S. After applying φ^{ℓ_0} the budget p^{-2n} is absorbed by $z^{np^{\ell_0}}$, precisely the n-th power of X. The resulting series has constant coefficient I, hence is invertible as a formal matrix. The normalization by B, the divisibility at every recursive step, and the fixed property of C are the first half of `radius_lemma` in `Lib/PadicHodge/Comparison/CrystallineLimit.lean`, used here as an integral comparison input. The displayed valuation calculation explains why the matrix is evaluable in the positive integral period ring; it explains in particular the appearance of $3e_0$, as opposed to taking p^{ℓ_0} arbitrarily large depending on Z.

Here is the next part of the argument, including the reason for the fourth power in the bound. Let $D \ge 0$ and $A_0 \ge 0$ be integers satisfying $e_0 A_0 \le p^{\ell_0} D$, and let $\overline{R} = B^+/E^{A_0}B^+$. A series $H(X) = \sum h_n X^n$ with coefficients in S has

$$p^{2D} H(X) = H_0(z) + z^{p^{\ell_0} D} H_1(X), \qquad H_0 \in S, \quad H_1 \in S[[X]]. \tag{24.9}$$

Indeed take the first D terms, multiplying each by p^{2D}, and rewrite the tail by shifting its index by D. The first term is integral because for $n < D$ the coefficient of $z^{p^{\ell_0} n}$ is $p^{2(D-n)} h_n$. The scalar $z^{p^{\ell_0} D}$ is divisible by p^{A_0} in $A_{\max}$. We use the precise truncation of $p^{A_0}\widehat{\lambda}_0(y)$ modulo E^{A_0} for $y \in A_{\max}$: there exists $s \in S$ with

$$p^{A_0}\widehat{\lambda}_0(y) - s \in E^{A_0} B^+.$$

This is the special E case of `exists_pow_mul_extendedCoord_eq_add`, obtained from the integral decomposition $p^{A_0} y = a + E^{A_0} y'$ with $a \in A$ in `Amax.exists_p_pow_mul_eq_add`, followed by $\widehat{\lambda}_0$. It is the exact ideal-theoretic approximation required here; truncation merely modulo some high power of p in B^+ would be meaningless. Since $\widehat{\lambda}_0$ fixes the convergent evaluations of $S[[X]]$ (the partial sums are fixed and converge p-adically after the prescribed scaling), applying it to (24.9) gives a polynomial approximation in the sense that $p^{2D} H(X)$ has the same image in $\overline{R}$ as some element of S.

Apply this entry by entry to the two formal units giving V and V^{-1}. Obtain $N, M \in \mathrm{Mat}_d(S)$ satisfying, after mapping to $\overline{R}$,

$$\overline{N} = p^{2D}\overline{V}, \qquad \overline{M} = p^{2D}\overline{V}^{-1}.$$

Consider the subset of S^d consisting of all y for which Vy is a constant vector of natural numbers modulo E^{A_0}. The S-span of these y contains every column of M, since $\overline{V}\,\overline{M} = p^{2D} I$.

It contains every column of $p^{4D}I - MN$, since these columns are mapped to zero by $\overline{V}$. For every $x \in S^d$, the identity

$$p^{4D}x = (p^{4D}I - MN)x + M(Nx) \tag{24.10}$$

then expresses $p^{4D}x$ in this span. This proves the second half of the radius lemma, with its original bound $p^{4D}S^d$. Notice the use of two approximations, each costing $2D$: the cost is not a determinant of an approximating matrix and is independent of d. This is an instance of a general pattern in the local proof: instead of solving an entire system of equations by its determinant, one approximates both a matrix and the matrix that provides a prescribed multiple of its inverse, and subtracts the residual identity.

At $A_0 = D = 0$, the integral radius factorization supplies the uniform positive comparison for all crystalline lattices with jumps in a fixed interval $[0, h]$. The multiplier is the image of $\varphi^{\ell_0}((E\tau)^h)$ in $A_{\max}$, independent of the dimension of T. Indeed in the normalized equations $UZ = E^hQ$ and $QL = \tau^hI$, choose $b = 0$ in (24.8). Apply Frobenius and replace the image of Z by CV; the product $\varphi^{\ell_0}(E^hQ)V^{-1}$ is integral and equals the invariant period matrix times the fixed matrix C after applying Frobenius. The invariance holds because Frobenius commutes with the Galois action and C is fixed by the group in question. Multiplication by $V\varphi^{\ell_0}(L)$ then gives precisely $\varphi^{\ell_0}((E\tau)^h)I$. Inverting this multiplier over the crystalline period ring is possible since $E\tau$ is associated to μ and the Frobenius translates of E needed here are units over the relevant localized period ring. This is the matrix version of `uniform_period_bound` in `Lib/PadicHodge/Comparison/CrystallineLimit.lean`. Its mathematical role differs from the usual statement that a sequence of bounded-weight crystalline representations has a crystalline limit: one also knows an explicit scalar uniform in the integral lattice and need not take a different period denominator on every individual line. The source cites the Dwork-trick arguments of Kisin and the integral comparison in the work of Bhatt and Scholze as conceptual antecedents to this construction; one needs the exact radius and scalar bounds for the subsequent length calculations.

24.6 Exponent lattices, Frobenius letters and the small part

We explain the role of the two variables of an exponent lattice by using the evaluations of the chart. Let $I = \mathbb{F}_{p^{f_0}}^{\times}$ be the set of letters. Write $i[-j]$ for the inverse Frobenius translate of a letter by j. At a fixed letter, the relevant jump spaces are those of the crystalline forms for each evaluation. An exponent system $o(j, u)$ assigns a lower bound on filtered valuation to the evaluations of those forms with jump u at letter $i[-j]$. The standard exponents of level l are

$$\mathrm{st}_l(j, u) = \begin{cases} 0 & j < l, \\ -u & j \geq l. \end{cases}$$

The membership conditions also restrict the point to the appropriate level of the localized frame and require that the evaluations take values in the chart value fields. Consequently they are stronger than the displayed inequalities over B_{dR} alone. If the lower bounds increase, the corresponding exponent module decreases. Intersecting two systems of bounds gives their pointwise maximum. If all exponents are shifted by m at $j = j_0$, multiplication by $E_{j_0}^m$ implements the shift in the rational module, because $\varphi^{-j}(E_{j_0})$ has valuation one precisely when $j = j_0$ and zero otherwise. Likewise applying Frobenius l times transports the letter $i[-l]$ to i and displaces the old filtration bounds by l places. This is why applying Frobenius to a lift over one letter can make the inequalities easier at old evaluations but introduces new inequalities at the first l places. Exact membership statements, including the required

value field and local-frame conditions, are in `New/Lift/ExponentLattices.lean`; the evaluation-subspace and strictness assertions needed in subsequent length estimates occur in `New/Lift/ExponentTheorems.lean`. They are not automatic properties of arbitrary sequences of filtrations.

We recall the algebra controlling the coefficient decomposition. If $q_0 = p^{f_0} - 1$, the operator multiplying by a primitive q_0-th root ζ from the coefficient ring admits the integral projectors

$$e_i = q_0^{-1} \sum_{j=0}^{q_0-1} (\omega(i)^{-1}\zeta)^j.$$

The denominator is a unit and all differences of distinct Teichmüller roots $\omega(i)$ are units. Hence $e_i e_{i'} = 0$ for $i \neq i'$, $e_i^2 = e_i$, and their sum is one: all three follow by evaluating the polynomial in the simple-root decomposition of $X^{q_0} - 1$, which is separable even modulo p. Frobenius sends e_i to e_{ip}, since it raises the Teichmüller factor, but not the coefficient-lattice factor, to the p-th power. If a bilinear pairing is balanced for multiplication by ζ on its two arguments, the evaluations on unlike letters are orthogonal. Indeed $\omega(i)P(x,y) = \omega(i')P(x,y)$ for the two arguments in their respective letter components; subtracting gives zero because the difference is invertible. These calculations hold integrally modulo p^m for every m. They explain why it suffices to test exponent inequalities one letter at a time, and why summing lengths over letters does not multiply them by the full number of letter pairs.

The cup-product argument uses a further fact about the small part, whose core is ordinary Artin–Schreier–Witt algebra. We prove the corresponding basic statement explicitly. Let $L = C_p^\flat$ be the algebraically closed tilt field, let $\mathfrak{m}_L$ be the ideal of elements of strictly positive valuation in its valuation ring, and let $W_m(L)$ denote the truncated Witt vectors. Put $Q = W_m(\mathfrak{m}_L)$, by which we mean the vectors all of whose Witt coordinates belong to $\mathfrak{m}_L$. This is an ideal of $W_m(\mathcal{O}_L)$ and an additive subgroup of $W_m(L)$ (not an ideal of the field's whole Witt ring). Frobenius on the latter acts on Witt coordinates by the p-th power. On the Tate-twisted tensor the auxiliary Frobenius acts on the Witt part, while the Galois group acts diagonally. Thus the fixed constants of that auxiliary Frobenius form the finite Tate line used in Section 23.6.

Lemma 24.4 (Artin–Schreier–Witt and small vectors)**.** *The map $\varphi - 1 : W_m(L) \to W_m(L)$ is onto, with kernel $W_m(\mathbb{F}_p)$. Its restriction to Q is a bijection. Thus if $(\varphi - 1)x \in Q$, there is a unique $c \in W_m(\mathbb{F}_p)$ with $x - c \in Q$.*

Proof. For $m = 1$, the equation $x^p - x = a$ has a root in the algebraically closed field L and its kernel is $\mathbb{F}_p$. For higher m, use the exact additive sequences given by truncation to length $m - 1$ and its kernel (identified with L by the last coordinate). Frobenius commutes with truncation and acts on the last coordinate by p-th power. Inductively solve the equation after truncation, lift the solution to length m, and correct its error in the last coordinate using the length-one equation. The Frobenius-fixed Witt vectors are precisely those all of whose coordinates lie in $\mathbb{F}_p$.

For $a \in Q$, the sum $y = -\sum_{n=0}^{\infty} \varphi^n(a)$ converges in the valuation topology of the truncated Witt vectors. Indeed every coordinate of $\varphi^n(a)$ tends to zero, all with valuation tending to infinity. Witt addition in the finite number of coordinates is continuous, and for each $\varepsilon > 0$ the subgroup of vectors whose coordinates have valuation at least ε is closed, by the integral polynomial formulas for addition and negation. Thus the tails of the sum tend to zero; completeness gives a limit. The valuations of all coordinates in the limit are still bounded below by a positive number (the minimum of the starting coordinate valuations that are finite), so $y \in Q$. Subtracting adjacent terms from the convergent sum shows $(\varphi - 1)y = a$. Uniqueness follows since $Q \cap W_m(\mathbb{F}_p) = 0$. Finally, if $(\varphi - 1)x \in Q$, solve that equation in Q

and subtract the solution; the difference is in the constant kernel just computed. The same zero-intersection assertion gives the uniqueness of its constant representative. □

The small part is open in the valuation topology on the finite Witt coordinates. If a Galois-dependent expression in them varies continuously and is required to have a representative among the constants modulo Q, that representative is locally constant: a sufficiently small change in the expression belongs to Q, and two distinct constants do not differ by an element of Q. In a tensor expression with a finite reduction of a Galois lattice, the action on that finite reduction is locally constant, and one reduces continuity to the finitely many Witt coordinates. This gives the continuity needed to use the Artin–Schreier–Witt correction in Section 23.6. The precise formal counterpart includes the passage to the diagonal tensor with the finite Tate module in `TateTwist.smallPart_spec`, `Lib/PadicHodge/WachGamma/KisinModuleDefs.lean`. The elementary proof here concerns the truncated Witt algebra and its topology; it should not be confused with an assertion that every integral period-ring primitive or invariant is automatically continuous in the topology of the full period ring.

24.7 Exact cokernel multipliers via finite differences

We expand the arithmetic in Proposition 23.6. In this subsection $J_l(i; r_0)$ denotes exactly its multiplier ideal: multiplication by a scalar in this ideal, up to the modulus r_0, permits the lift of every constant of letter i to standard exponents of level l. The letter subscript is essential in a Frobenius step. In writing A_{inf} for its coefficient ring we again identify it with A above.

There is an elementary Nakayama argument that removes a large modulus but must be used with the precise power. If $p^c \in J_l(i; r_0)$ with $r_0 = p^{c+1}$, write $p^c = x + p^{c+1}y$ for $x \in J_l(i; 0)$ and $y \in A_{\mathrm{inf}}$. Then $p^c(1 - py) = x$, and $1 - py$ is a unit of the p-adically complete ring. Thus $p^c \in J_l(i; 0)$ and $J_l(i; p^{c+1}) = J_l(i; 0)$, since p^{c+1} already belongs to it. Importantly this argument does not work if the modulus on the right is only p^c. The strict difference of one in the powers of p accounts for the $+1$ in the proof of the unramified cokernel estimate.

We recall more explicitly the finite difference behind part (ii). For $v > 0$, the chart calculation gives $\alpha, \beta \in A_{\mathrm{inf}}$ such that

$$\alpha([\epsilon]^v - 1) = E\tau(p^{v_p(v)} + \beta z). \tag{24.11}$$

This is the precise scalar-lifting statement `exists_tau_mul_E_eq_mul_epsTilt_pow_sub` in `Lib/PadicHodge/WachKisin/PeriodMatrixBasis.lean`. To understand the valuation in it, factor $[\epsilon]^v - 1$ first by $[\epsilon] - 1 = E\tau v_0$ with v_0 a unit. At first order in $[\epsilon] - 1$ the other factor has constant term v. The lifting-the-exponent argument of the source uses the relation between ξ and (p, z) to remove the prime-to-p factor of v and control the residual term by z in the needed quotient. The exact existence of the coefficients in (24.11), rather than a valuation identity just in a field, is used below.

If $f = p^c z^v$ and $g(z) = [\epsilon]z$, multiplication of (24.11) by $p^c z^v$ gives

$$\alpha(gf - f) = z\tau E(p^{c+v_p(v)} z^{v-1} + \beta f).$$

Apply the finite-difference property of the ideal (`mem_liftIdeal_of_derivation`) to this identity. Thus $p^{c+v_p(v)} z^{v-1} \in J_1(i; r_0)$ whenever $p^c z^v$ belongs. This property uses the $z\tau$-divisibility of the lifted Galois differences and the coefficient retraction, and is stronger than the assertion that an ordinary ideal is invariant under Galois. Induction lowers v to pa for any $pa \le v$, with cumulative extra p-power $\sum_{pa<j\le v} v_p(j) = v_p(v!) - v_p((pa)!)$. Applying the inverse Frobenius step transports the letter i^p to i and the power z^{pa} to z^a. Multiplying by

any remaining power of p in the ideal proves the stated form of Proposition 23.6(ii). For part (iii), the analogous depth-step calculation for E^n replaces one E by one z, with the bound $\delta_0 + v_p(n)$ (here the chart's specified different bound occurs). Iterating adds precisely $\delta_0 n + v_p(n!)$. The leading coefficient needed for this step comes from `exists_depth_step`, used inside `mem_liftIdeal_of_mul_E_pow`. Thus (ii) and (iii) reduce to these two specialized scalar and ideal-derivation properties; no extra implicit constants are permitted in the factorial sums.

We can now prove the two remaining arithmetic deductions of that proposition. We state the proof because it yields the strict coefficient of a in z^a and the special coefficient $p-3$ of the unramified weight bound.

Proposition 24.5 (A minimal exponent bound)**.** *Assume all allowed jumps are at most n, and put $r_0 = p^\kappa$. Using the frame inclusions and the Frobenius and finite-difference properties of the multiplier ideals, there exists an integer $a \geq 0$ with*

$$z^a \in \bigcap_i J_0(i; r_0), \qquad (p-1)a \leq e_0(n+\kappa).$$

Proof. The integral frame inclusion (with its scalar τ^h for a bounded weight range of the fixed extension) yields a power τ^h in the exact standard-exponent multiplier ideal of every letter. For each fixed κ, this implies that some power of z is in the intersection of those ideals modulo p^κ. Indeed in A_{inf}/p the class of τ has positive valuation and is nonzero. A sufficiently high power of z is divisible by τ^h there, so a power of z belongs to $(p, \tau^h)A_{\mathrm{inf}}$. Raising this containment to the power κ puts a higher power of z in $(p^\kappa, \tau^h)A_{\mathrm{inf}}$. Let a be the least exponent for which all letters belong modulo p^κ. This also covers $\kappa = 0$, where the ideal already contains the modulus 1 and $a = 0$.

By the standard-exponent shift in Proposition 23.6(i), $E^n z^a$ lies in every level-one ideal $J_1(i; r_0)$. From (24.1), $z^{e_0} \in (p, E)$ in A_{inf}. Hence

$$z^{e_0(n+\kappa)+a} \in z^a (p, E)^{n+\kappa} \subset (p^\kappa, E^n z^a) \subset J_1(i; r_0),$$

where in the middle containment each monomial either has E-exponent at least n or p-exponent at least κ. Set $N = e_0(n+\kappa) + a$ and $q = \lfloor N/p \rfloor$. The factorials $(pq)!$ and $N!$ have the same p-valuation, since none of the integers $pq+1, \ldots, N$ is divisible by p. Applying part (ii) of Proposition 23.6 with zero cost sends $z^N \in J_1(i^p; r_0)$ to $z^q \in J_0(i; r_0)$. This holds for all i, because Frobenius permutes the letters. Minimality implies $a \leq q \leq N/p$, or $(p-1)a \leq e_0(n+\kappa)$. □

The proof uses an arbitrarily large initial height h only to show that the set of possible exponents is nonempty. Its numerical conclusion is independent of h; it is the minimization and the Frobenius operation, not an estimate for that initial height, which reduce the final upper bound to n. This illustrates the value of applying the arithmetic reasoning to exact ideals rather than assigning an ad hoc approximation to each individual constant.

Proposition 24.6 (Unramified cokernel bound)**.** *Suppose $e_0 = 1$, $E = z + p$, $p \geq 5$. If all allowed jumps are at most $(p-3)(s+t)$ for nonnegative integers s, t, then $p^s E^t \in J_0(i; 0)$ for every letter i.*

Proof. Put $L = s + t$ and work at modulus $r_0 = p^{L+1}$. Proposition 24.5 gives a nonnegative integer b such that z^b belongs to every $J_0(i; r_0)$ and

$$(p-1)b \leq (p-3)L + L + 1 = (p-2)L + 1.$$

Legendre's valuation formula implies

$$(p-1)v_p(b!) \le b, \qquad (p-2)(b+v_p(b!)) \le (p-1)b \le (p-2)L+1.$$

For the middle inequality, use $(p-2)v_p(b!) \le b$. Since $p-2>1$ and the summands are integral, $b+v_p(b!) \le L$.

We claim that for every $0 \le m \le L$,

$$p^{L-m}z^m \in J_0(i;r_0). \tag{24.12}$$

If $m \ge b$, ideal closure proves the claim. If $m<b$, transport z^b by Frobenius to $z^{pb} \in J_1(i^p;r_0)$. The difference of factorial valuations needed to lower pb to pm is

$$v_p((pb)!) - v_p((pm)!) = b - m + v_p(b!) - v_p(m!) \le L-m,$$

using $v_p((pq)!) = q + v_p(q!)$. Part (ii) of Proposition 23.6 now proves (24.12). For $m=0$ this puts p^L in the ideal modulo p^{L+1}, so the Nakayama step at the start of this subsection shows $J_0(i;r_0) = J_0(i;0)$. Finally expand

$$p^s E^t = p^s(z+p)^t = \sum_{m=0}^{t} \binom{t}{m} p^{L-m}z^m,$$

each term of which belongs to the exact ideal by (24.12). This proves the assertion. □

The last two arguments follow respectively from `exists_z_pow_mem_liftIdeal` and `killsCokernel_of_isUnramified`, `New/Lift/CokernelKillers.lean`. In particular the proof does not use an unspecified "large enough power" of p at the last step. If t is positive, the conclusion can be interpreted as allowing lifts over p^s at the expense of a deeper E-exponent at the first place, by dividing by E^t in the localized module. This is the exact trade-off between the cost of a class and the exponent safety test of Section 23.6. The conversion of multiplier ideal membership to points with the specified depth, and the improvement of depth at no cost under NM, are the more specialized lift/depth inputs in `New/Lift/LongGap.lean`.

24.8 A rank-sensitive bound for raising filtered conditions

We now give further details of the length mechanism for Theorem 23.7. Let L denote a finite free $\mathbb{Z}_p$-module of constants, and $D \subset L$ the submodule of constants with admissible lifts satisfying specified depth conditions. Such submodules need not be O-linear before the final intersection, because the coefficient retraction and intermediate linear-combination lemma are over $\mathbb{Z}_p$. It is therefore appropriate to do the following count over $\mathbb{Z}_p$, converting the resulting lengths by (23.1) only at the end.

Lemma 24.7 (Length from relations)**.** *Suppose $p^cL \subset D$, and for every list of $q+1$ elements of L there is a linear combination with coefficients in $\mathbb{Z}_p$, at least one equal to 1, whose value belongs to D. Then* $\operatorname{length}_{\mathbb{Z}_p}(L/D) \le cq$.

Proof. Decompose the finite quotient over the DVR as a direct sum of nonzero cyclic modules $\mathbb{Z}_p/p^{a_j}\mathbb{Z}_p$ with $a_j \ge 1$. Each $a_j \le c$. If there were $q+1$ of them, taking lifts of their generators would give $q+1$ classes independent in $(L/D)/p(L/D)$. A relation of the stated type would contradict their independence. Thus there are at most q summands and their length is at most cq. □

In the integral comparison, a single filtration step at index a and jump u has $q = f_0 m(i[-a], u)$, where $m(i[-a], u)$ is the corresponding jump multiplicity. We explain algebraically why this dimension, not the rank of the full extension lattice, governs the count. Fix a letter i and an exponent system o, strictly decreasing from the lower jumps to the selected jump u at that index. When the single exponent $o(a, u)$ is increased by one, only the leading coefficients in the jump-u quotient have to be cancelled: the forms of lower types already satisfy the stricter bound. They can be read as a leading coefficient vector of dimension at most $m = m(i[-a], u)$ over the chart's value field. For a list x_j of $f_0 m + 1$ points satisfying the original exponents, choose a maximal independent sublist of their leading coefficient vectors for which every remaining vector is an integral combination of the selected ones. Such a choice is possible by the following elementary argument. If the leading coefficient vectors lie in a valued field's vector space of dimension $r \le m$, fix coordinates in their linear span. Of all the r-tuples choose one whose nonzero determinant has maximal absolute value. Every coordinate of the expansion of any other vector in this basis is a ratio of determinants, by Cramer's rule; replacing any one column cannot increase that absolute value. Hence each ratio lies in the valuation ring. If all vectors are zero the selected list is empty. This is the non-Archimedean basis selection used by the source under the name `exists_auerbach`.

The chart comparison identifies and lifts these integral coefficients in such a way as to write, in the localized module,

$$x_j - \sum_{k=1}^{r} s_{kj} x_{b_k} \in M', \qquad s_{kj} \in S,$$

up to a summand EM; here M represents the original exponent module and M' the module with the one extra step. More precisely the source establishes that the difference, after coefficient retraction, lies in M', and that the discrepancy of the coefficient lift from S is accounted for by the summand EM. This is the value-field and filtration part of `exists_comb_of_jumpMultiplicity_le` in `New/Lift/LongGap.lean`. It is an integral assertion, involving not just a change of leading-coefficient basis over a field, and we retain it as a specialized input. The remaining linear algebra is as follows. Since the chart is unramified, reduction modulo $E = z + p$ gives $S/(E) \simeq \Lambda$ and every s_{kj} can be written uniquely modulo E as $w_{kj} + E y_{kj}$, with $w_{kj} \in \Lambda$. For instance the residue is the value of the convergent series at $-p$. The map

$$\mathbb{Z}_p^{f_0 m+1} \longrightarrow \Lambda^r, \qquad (c_j)_j \longmapsto (\textstyle\sum_j c_j w_{kj})_k$$

has a nonzero kernel over $\mathbb{Q}_p$, since $[\Lambda : \mathbb{Z}_p] = f_0$ and $r \le m$. Clear denominators and divide the relation by the coefficient of least valuation. This gives a relation with all coefficients integral and one coefficient equal to 1. Substituting in the points gives a combination in $M' + EM$, as desired. Notice the factor f_0, which comes from the rank of the unramified chart's scalar residue, not from the number of letters.

There is a further important distinction: the EM error is not automatically zero on lifting constants. One needs the depth operator and the auxiliary no-multiples condition to remove it. In the unramified chart the elementary scalar congruence

$$E^{p^a} - E_a \in pS \tag{24.13}$$

follows from $E = z + p$ and $E_a = z^{p^a} + p$, since both reduce modulo p to z^{p^a}. If no multiple of p lies strictly between the current depth at E and that depth plus p^a, one can use successively the zero-cost operators at E to replace an EM error by an $E^{p^a} M$ error modulo deeper conditions. Relation (24.13) allows this to contribute in $E_a M + pM$. In the proof of `relLength_deepConstants_le`, the precise correction uses `exists_reaches_deep` to change the lifts while keeping the prescribed constants and uses the kernel exponent lattices within

the extension. This is the remaining step needed to pass from a relation in $M' + EM$ to a relation on the actual constants in D'. Thus one may not replace the actual constants by the full extension lattice when applying Lemma 24.7. The arithmetic after this correction is exact: the exponent step costs $c_p(a, n)$, the multiplicity bound gives q, and its loss is at most $qc_p(a, n)$.

We unfold the sums of such steps into the actual two-window bound. For this purpose assume the conclusions of the preceding integral relation/depth comparisons (precisely `relLength_deepConstants_le` and the initial depth construction in `exists_relLength_deepConstants_le` in `New/Lift/FewLowJumps.lean`). Write $A_0 = s$ for the budget in (23.15), η for the width of the low window, W for W_0, and b for β_i at the present letter. The steps compare depth systems c, c' differing by increases at a only, with any decreases at the other divisors allowed. Suppose the new depth at a is antitone as a function of the jump, the old depth is at least $-u$, and no nonzero multiple of p arises strictly between $c(0, v)$ and $c(0, v) + p^a$ for the old system. If for each jump u at $i[-a]$,

$$\sum_{n \in [c(a,u),\, c'(a,u))} c_p(a, n) \le A_1 + B_1 \mathbf{1}_{|u| \le \eta},$$

then the colength paid for the depth change is at most $A_1 r_{i[-a]} + B_1 n_{i[-a]}$. Here n_j and r_j bound respectively the low and total counts, multiplied by f_0, as in (23.16). To see the sum, raise one coordinate at a time, always using the smallest jump requiring a further increase so that the strictness in the lower-jump conditions is preserved. Each step is bounded by the preceding application of Lemma 24.7 after the depth correction. Multiplying the separate bounds by the multiplicities and summing yields the expression $A_1 r_j + B_1 n_j$.

The initial construction of constants of depth $2s + \eta - u$ at $a = 0$, remaining depth $-u$ elsewhere, costs at most sn_i. The construction uses the exact unramified cokernel for $p^v E^{s-v}$, $0 \le v \le s$, and zero-cost no-multiple depth operators, then successively imposes the low negative jump conditions for $-\eta \le u < 0$. At the u-step multiplication by p^s gives the annihilator and the relation bound is the corresponding jump count. In particular the initial depth construction's cost is s times this small count, and is already sufficient for the positive sign. The correctness of those lifts in the localized extension is the initial-depth part of the formal-source theorem just cited, not a consequence of the free-module rank argument alone.

For the negative sign the final required depth at $a = 0$ is $b + W + 2s$. Choose the integer q with

$$b + W + 2s \le qp < b + W + 2s + p.$$

First strengthen the bound at $a = 0$ to the constant qp. The sum of the possible nonnegative step costs on a low jump is at most $v_p((qp)!)$. The window condition gives $qp < (p - 1)s + p$, so

$$v_p((qp)!) \le \frac{qp}{p-1} < s + \frac{p}{p-1} < s + 2,$$

and this integer is at most $s + 1$. On a far negative jump, the starting depth $2s + \eta - u$ lies near b. Every step up to $b + W + 2s$ lies within the interval of radius $6(s + \eta + W)$ about b and hence has zero cost by NM. Every remaining step below qp also has zero cost, since qp is the first multiple of p at or above $b + W + 2s$. If the starting depth exceeds qp, there are simply no steps to pay. Thus this operation costs at most $(s + 1)n_i$. The choice to end at a multiple of p is deliberate: no multiple lies strictly between qp and $qp + p$. This supplies the auxiliary no-multiples condition at $a = 1$. If one had first chosen the exact desired depth $b + W + 2s$, the auxiliary condition would in general fail.

Strengthen the bound at $a = 1$ from $-u$ to

$$W + 1 - u + 2s\mathbf{1}_{|u| \le \eta}$$

at letter $i[-1]$. Each step lies within one of the no-multiples windows about $0, \beta_{i[-1]}, -\beta_{i[-1]}$. The number of steps is $W+1$ on a far jump and $W+1+2s$ on a low jump. Since the chart is unramified, every such step has cost $c_p(1,n) \le 3$ (with $v_p(0)$ set to zero). We can now weaken the bound at $a=0$ from qp to $b+W+2s$ with no additional cost. The total bound is

$$sn_i + (s+1)n_i + 3(W+1)r_{i[-1]} + 6sn_{i[-1]},$$

which proves the arithmetic part of (23.16), conditional on the exact relation and operator statements. Summing over all letters gives

$$(8s+1)\sum_i n_i + 3(W+1)\sum_i r_i \le 11\sum_i \big((s+1)n_i + (W+1)r_i\big).$$

The Frobenius letter permutation is essential in the equality after the summation. The intersection over the letters of the preimages of these deep-constant modules has at most that colength, by the injection into the direct sum of their separate quotients. After the length normalization this yields the precise coefficient $11f_0$ of the few-low-jumps theorem, once the points are verified to meet the safety inequalities. That last verification, including the two zero-loss alternatives in Theorem 23.7, remains the specific formal-source input `exists_scaledCupBounds_few_low_jumps`, together with `scaledCupBounds_of_killsCokernel` in `New/Lift/LongGap.lean`. The argument explains analytically why a narrow far window, even at a very large absolute weight, costs $O(r(W+1))$; the cost proportional to s need only be charged to the low multiplicities. A crude bound for the whole rank, or replacing the low-count term by an annihilator of every constant, would lose this advantage.

Chapter 25

Integral crystalline comparison, depth operators and the Frobenius walk

The use of crystalline representations in the lower bound is unusually quantitative. It asks for a comparison of finite congruence classes with torsion coefficients which is nearly insensitive to the rank and to large isolated Hodge–Tate weights. Neither the equality of the dimensions of two rational period spaces nor a bound linear in the largest Hodge–Tate weight has this property. In this chapter we give the underlying comparison and lifting arguments in substantially more detail. We use the same conventions as Chapters 23 and 24. The present account separates the ring-theoretic descent, the construction of a frame, the elementary moving operators, the Taylor coefficient at a long gap, and the cup-product argument. In particular it identifies which assertions require comparison theorems for the period rings and which follow from elementary algebra once those theorems have been proved.

We give references to both layouts of the formal development where useful. The compact source places the local lifting results in `New/Lift/`; the expanded readable source separates them under `New/Congruences/Lifting/`. The entries called `withOrders` in the expanded source are the exponent lattices called `withExponents` in the compact source. Likewise `IsScaledCupCoboundary` and `ScaledCupBounds` designate the corresponding conclusion about a cup product. We reconstruct that conclusion as a statement about continuous cohomology, not as an assertion about pointwise vanishing of a particular cocycle.

25.1 The period rings and an integral contraction

Fix a prime p. Put $C_p = \widehat{\overline{\mathbb{Q}}_p}$, $A = W(\mathcal{O}^{\flat}_{C_p})$, and write φ for Witt Frobenius. Choose a compatible system ϵ of primitive p^n-th roots of unity; let $\mu = [\epsilon] - 1$, $\tau = [\epsilon^{1/p}] - 1$, and $\xi = \sum_{j=0}^{p-1} [epsilon^{1/p}]^j$. Thus

$$\varphi(\tau) = \mu = \xi\tau, \qquad \xi - p = \tau w \quad (w \in A).$$

The equality on the right follows by expanding the finite geometric sum in powers of τ. We use the rings

$$A_{\max} = \widehat{A[\xi/p]}^{\,p}, \quad R = A_{\max}[1/p], \quad B = R[1/\mu].$$

The last ring is the version of the crystalline period ring used in the formal source. We use B^{+}_{dR} with fraction field B_{dR}, and put $\mathrm{Fil}^{j} = t^{j} B^{+}_{\mathrm{dR}}$ for any integer j, where t generates the maximal ideal. The cyclotomic convention is $\varphi(\tau) = \mu$, not $\varphi(\mu) = \tau$.

Here are the elementary period-ring properties needed in the contraction. We state explicitly the structural facts used:

(i) A and R are domains, $A \hookrightarrow A_{\max} \hookrightarrow R$, A is (p,ξ)-adically complete and separated, $A_{\max}$ is p-adically complete and separated, and φ acts on them, agreeing on the overlaps.

(ii) Fontaine's map $\theta : A \longrightarrow \mathcal{O}_{C_p}$ is surjective, with kernel (ξ); its extension to R has kernel (ξ). In particular if $a \in A$ becomes divisible by ξ^m in R it already is divisible by ξ^m in A.

(iii) Write $X = \xi/p \in A_{\max}$. If $Xb \in p^n A_{\max}$ for $b \in A_{\max}$, then $b \in p^n A_{\max}$. For $q = \tau$ and $q = \xi$, every $x \in A_{\max}$ and $n \geq 0$ satisfy $p^n x \in A + q^n A_{\max}$. Furthermore every $x \in A_{\max}$ can be approximated modulo $p^k A_{\max}$ by an element whose appropriate p-power multiple belongs to A: for each k there are m, a, w with $a \in A$, $w \in A_{\max}$, $p^m x = a + p^{m+k} w$.

The second implication in (ii) requires no separate divisibility theorem beyond the kernel calculation: if $a = \xi b$ in R, then $\theta(a) = 0$, so $a = \xi a'$ in A and cancellation in R gives $a' = b$; iterate. We include an explanation of (iii). At the polynomial stage use the binomial theorem. Since $\xi = p + \tau w$, $p^n P(\xi/p)$ has an integral polynomial remainder modulo τ^n: expand $P(1+\tau w/p)$; the terms of degree $j < n$ contain $p^{n-j}\tau^j w^j$. For $q = \xi$ use the expansion directly in powers of ξ/p. Passing from polynomials to the completed ring requires an integral closure statement. The approximation in the last sentence of (iii) follows by density of the polynomial subring in $A_{\max}$ and clearing the denominators of one polynomial. The first property of (iii) is the regularity of the image of X in $A_{\max}/p$. In the expanded source these facts are parts of `Lib/PadicHodge/PeriodRings/CrystallineRings.lean`, particularly `Amax.exists_p_pow_mul_eq_add`, `p_pow_dvd_of_dvd_aeval_mul`; they use the presentation as a completion of the ring with $pX = \xi$. We retain the integral structural properties in (i)–(iii) as period-ring inputs. Their consequence for μ has the cost $2n$ instead of n: first write $p^n x = a + \tau^n w$, then $p^n w = a' + \xi^n w'$, so $p^{2n} x = p^n a + \tau^n a' + \mu^n w'$. In particular, using the last approximation of (iii) with sufficiently high k gives

$$p^m y = a + p^{m+k}\mu^n w' \qquad (y \in R,\ a \in A,\ w' \in A_{\max})$$

with some m depending on y, k, n. Indeed first choose $p^d y \in A_{\max}$, approximate it by a polynomial up to $p^{d+k+2n} A_{\max}$, and clear polynomial denominators by p^{m-d}. Replace the normalized error w by its two-stage decomposition just given. This is `BPlus.exists_Ainf_approx`, used in Corollary 25.7. The subsequent contraction argument itself will not require a bound on a matrix determinant.

Lemma 25.1 (Contraction of powers of p)**.** *Under the preceding structural properties,*

$$A \cap p^m A_{\max} = (p,\xi)^m A \qquad (m \geq 0).$$

Proof. The containment from right to left follows from $\xi = pX$. For the other direction suppose $a \in A \cap p^{m+1} A_{\max}$. Since θ extends to $A_{\max} \longrightarrow \mathcal{O}_{C_p}$, $\theta(a)$ is divisible by p^{m+1} there. Surjectivity of θ gives $c \in A$ with $a - p^{m+1}c = \xi b$ for $b \in A$. In $A_{\max}$ cancellation of p gives $Xb \in p^m A_{\max}$, so $b \in p^m A_{\max}$ by (iii). Induction gives $b \in (p,\xi)^m$. Hence both summands $\xi b, p^{m+1}c$ belong to $(p,\xi)^{m+1}$. The assertion starts at $m = 0$. □

Theorem 25.2 (Uniform Frobenius contraction)**.** *Let $V \subset A_{\max}$ be a subset, let $u \in A$ and $n \geq 0$. Suppose that for every $v \in V$ there exists $w \in V$ with $\xi^n v = u\varphi(w)$ in $A_{\max}$. Then $p^{n+1} V \subset A$.*

Proof. For each $k \geq 0$ and each $v \in V$ we construct $a_k \in A$ and $z_k \in A_{\max}$ with

$$p^{n+1}v = a_k + \tau^{n+1}p^k z_k.$$

For $k = 0$ this is (iii) applied to v. Given v, choose a predecessor $w \in V$ as in the hypothesis and choose the k-th approximation $a_k(w), z_k(w)$. Apply $u\varphi$ to it:

$$\begin{aligned} u\varphi(a_k(w)) &= p^{n+1}\xi^n v - u(\xi\tau)^{n+1}p^k\varphi(z_k(w)) \\ &= \xi^n\big(p^{n+1}v - \tau^{n+1}p^{k+1}uX\varphi(z_k(w))\big). \end{aligned}$$

The left side belongs to A, the final expression to $\xi^n R$. Property (ii) permits division by ξ^n in A. This proves the next approximation with $z_{k+1} = uX\varphi(z_k(w))$. For $j \geq k$ the two approximations satisfy $a_j - a_k \in p^k A_{\max}$. By Lemma 25.1 they are congruent modulo $(p, \xi)^k A$. Thus they converge to an element $a \in A$. Their images converge in $A_{\max}$ to $p^{n+1}v$, since $(p, \xi)^k A \subset p^k A_{\max}$. Separatedness gives $a = p^{n+1}v$. □

This is the argument of `exists_Ainf_eq_of_contraction`. It is uniform in V and in the possibly non-invertible scalar u. It is an important device because a preliminary bound after passage to $A_{\max}$ may be quite ineffective in A. The contraction removes that difficulty with the extra cost p^{n+1}, even though one has no uniform bound on the number of elements in V. The factor ξ gained at the induction step is what turns a fixed denominator into an adically convergent series; cancellation of ξ^n in the original integral ring is essential.

25.2 The period-basis argument

We formulate an algebraic result underpinning `IsPeriodMatrix.exists_basis` and `IsPeriodMatrix.exists_frame`. Use the chart $S = \Lambda[[z]] \subset A$, E associated to ξ, where Λ is finite unramified over $\mathbb{Z}_p$. Denote the ring of scalar period coordinates by $S_B \subset R$. Recall that S_B is defined in the source by a bicommutant condition: multiplication by its elements commutes with the extensions of continuous S-linear endomorphisms of A preserving the integral period ring. We need the following properties of this construction.

(a) S_B contains S and p^{-1}. A chosen topological S-basis (b_j) of A, with coordinate functions $\lambda_j : A \to S$, has extended coordinate functions $\widehat{\lambda}_j : R \to S_B$. These extend the original maps, are S_B-linear for multiplication, and finite sums $\sum_{j \in J} b_j \widehat{\lambda}_j(y)$ approximate $y \in R$ modulo arbitrarily large powers of p times $A_{\max}$.

(b) The extensions of S-linear endomorphisms of A preserve $A_{\max}$; they commute with Frobenius conjugation in the sense that the extension of $\varphi^{-1} f \varphi$ satisfies $\varphi \widehat{f'} = \widehat{f}\varphi$. Here φ is bijective on A, and f' is S-linear for the Frobenius-stable chart.

(c) The principal ideal (q) for $q \in A$ with $p \nmid q$ is closed in the (p, ξ)-adic topology. Also $\bigcap_j z^j A = 0$.

(d) If d divides μ in A, any fixed matrix M over R has a bound k such that

$$My = d^h w, \quad y \in A_{\max}^r,\ w \in R^r \quad \Rightarrow \quad p^k w \in A_{\max}^r.$$

The k may depend on M, d, h but is independent of y, w. Finally $\varphi(\tau)^h \mid \varphi(b)$ implies $\tau^h \mid b$ for $b \in R$.

These properties, with the integral comparison of kernels, are the structural inputs of the period-basis proof, not assertions about every infinite S-module. For (d), the nonzero divisor d divides μ and the latter has a controlled reduction in $A_{\max}$; a denominator selecting M first suffices, and division by d^h costs a fixed extra p-power in $A_{\max}$. The source proves the required statement in `exists_p_pow_mul_mem_amaxImage`. For (c), closure of (q) can also be seen by the same successive reduction modulo p used to prove completeness of A in its two-generator topology: reduction of q is nonzero in the valuation domain $\mathcal{O}^\flat_{C_p}$ whose nonzero principal ideals are closed for the topology of any element of positive valuation. Lifting the compatible quotients successively modulo p^n requires the integral closure of the image of a nonzero principal ideal at each finite Witt length. We use the source's closure lemma for this last integral step. These are the only asserted inputs to the next theorem. They are supplied by `Lib/PadicHodge/Kisin/PeriodMatrixBasis.lean` and the chart and period-ring constructions it imports.

Let $U, W \in M_d(R)$ obey $UW = (E\tau)^h I$ and suppose U is a period matrix in the following sense:

$$\varphi(U) = vUF, \qquad v \in A^\times, \quad F \in \mathrm{GL}_d(S_B),$$

and whenever $Uy = E^h x$ with $x \in A^d$, then $U\widehat{f}(y) \in E^h R^d$ for all coordinate endomorphisms $f : A \to A$ considered in (b). We assume $\widehat{f}$ is S_B-linear and that $\varphi(E) \in S_B^\times$. The period-matrix construction of a crystalline representation with weights in $[-a, b]$ supplies $h = a + b$ and the unit factor v (see the more explicit calculation in Section 24.4). In the following, we treat the existence and the stability of the period matrix as inputs; its passage to an actual integral frame will be proved.

Lemma 25.3 (Integral bound for all coordinate transforms)**.** *There exists an integer N such that, for every solution $Uy = E^h x$ with $y \in R^d$, $x \in A^d$, and every f as above, if $U\widehat{f}(y) = E^h w$, then $p^N w \in A^d$.*

Proof. The relation $UW = (E\tau)^h I$ implies injectivity of U and also $WU = (E\tau)^h I$. First $W(E^h x) = (E\tau)^h y$, so (d) supplies one fixed p^a putting all such y into $A^d_{\max}$. The extensions preserve this integral ring. Applying (d) to $U(p^a \widehat{f}(y))$ puts every w into $A^d_{\max}$ after multiplication by the same p^{a+b}, with some fixed b.

For each solution define $x' = \varphi^{-1}(vE^h x) \in A^d$. We claim there is a solution $Uy' = E^h x'$ with $y' \in R^d$. In $\mathrm{Frac}(R)^d$ its unique candidate is $y' = \tau^{-h} W x'$. Applying Frobenius gives

$$F\varphi(y') = \varphi(E)^h y,$$

by cancellation of vU. The right side lies in R^d and F is invertible over R. Thus $\varphi(\tau)^h$ divides $\varphi(Wx')$, and the inverse Frobenius property in (d) gives $y' \in R^d$. Let $f' = \varphi^{-1} f \varphi$ and define w' by $U\widehat{f'}(y') = E^h w'$. The conjugation property of extensions and their linearity over the entries of F and over $\varphi(E)$ give

$$F\varphi(\widehat{f'}(y')) = \varphi(E)^h \widehat{f}(y).$$

Apply φ to the defining equation of w', use $\varphi(U) = vUF$, and cancel $\varphi(E)^h$ to get $\varphi(w') = vE^h w$. Write $\xi = Ee_1$ for a unit e_1 of A. Each coordinate of $p^{a+b}w$ is now in a set $V \subset A_{\max}$ satisfying

$$\xi^h p^{a+b} w = e_1^h v^{-1} \varphi(p^{a+b} w').$$

Theorem 25.2 applies with $n = h$, $u = e_1^h v^{-1}$. Take $N = a + b + h + 1$. □

Theorem 25.4 (A basis and saturated solution space)**.** *Under these assumptions there are $Q \in M_d(A)$ and $Z \in M_d(S_B)$ with $\det Q \neq 0$ and $UZ = E^h Q$, such that*

$$P = \{x \in A^d : \ Uy = E^h x \text{ for some } y \in S_B^d\} = QS^d.$$

For every solution $Uy = E^h x$ with $y \in R^d, x \in A^d$, every row $l \in A^{1\times d}$, and every $q \in A$ not divisible by p, the divisibility $q \mid lQ$ entry by entry implies $q \mid lx$.

Proof. As explained in Section 24.4, P is p-saturated and z is injective on P/pP. We verify the rank estimate before using the finite freeness lemma. Let $x_1, \dots, x_r$ be S-linearly independent elements of P, with corresponding solutions $y_1, \dots, y_r \in S_B^d$. Suppose $\sum_j a_j x_j = 0$ for $a_j \in A$. Since U is injective, this is equivalent to $\sum_j a_j y_j = 0$. Apply the extended coordinate $\widehat{\lambda}_i$; by the commutant property,

$$\sum_j \lambda_i(a_j) y_j = 0.$$

Multiply by U and cancel E^h to obtain $\sum_j \lambda_i(a_j) x_j = 0$. All the coefficients vanish by S-linear independence, for all i. The original coordinates separate A, hence $a_j = 0$ for all j. Thus the x_j are independent even over A and $r \le d$. The separated regular-module lemma, applied to $P \subset A^d$ with the regular sequence (p, z) and this rank bound, gives a finite free S-module (Lemma 24.3). Pad a basis by zero columns to size d, forming Q, and put the corresponding y in Z.

Now let x solve the wider equation $Uy = E^h x$, with $y \in R^d$. For each i, by Lemma 25.3, $p^N E^{-h} U \widehat{\lambda}_i(y)$ lies in A^d. Its solution $p^N \widehat{\lambda}_i(y)$ lies in S_B^d, so the vector can be written Qs_i with $s_i \in S^d$. Choose k from (d) for the matrix U and divisor E. Given m, approximate y by $\sum b_i \widehat{\lambda}_i(y)$ modulo $p^{m+k} A^d_{\max}$, using a common finite set J for its finitely many entries. The integral error

$$t = p^N x - Q \sum_{i \in J} b_i s_i$$

satisfies $p^{-(m+k)} E^h t \in U A^d_{\max}$. Property (d) implies $t \in p^m A^d_{\max}$ and Lemma 25.1 implies $t \in (p, \xi)^m A^d$. If $lQ = qg$, then

$$l(p^N x) - qg \sum_{i \in J} b_i s_i = lt \in (p, \xi)^m A.$$

Since this holds for every m, closure of (q) gives $q \mid p^N lx$. The element p is prime in A, since $A/p = \mathcal{O}^\flat_{C_p}$ is a domain. When $p \nmid q$, cancellation of p^N in this divisibility is justified by repeatedly observing that $p \mid qa$ implies $p \mid a$. Thus $q \mid lx$.

If $\det Q = 0$, choose a nonzero row l over A that annihilates Q, clearing denominators of a kernel row over its fraction field. For every $x_0 \in A^d$ the vector $\tau^h x_0$ is a solution with witness Wx_0. Hence $q \mid \tau^h l x_0$ for every $q = z^m$, since $p \nmid z$. Separatedness for (z) implies $\tau^h l x_0 = 0$. Take x_0 successively to be the standard basis vectors; cancellation of τ gives $l = 0$, a contradiction. In particular no padding was in fact needed: the original basis of P has rank d. □

The point of this proof is not the existence of a rational basis of periods; that was already given by the crystalline hypothesis. Its novelty is that the columns simultaneously span an integral solution space over S and control divisibility of all solutions with coefficients in the larger positive period ring. The argument is a convenient place to separate the structural period-ring lemmas from the finite and adic linear algebra. No bound on the exponent N of Lemma 25.3 is needed for the final divisibility result, since it is removed by the primality of p. In contrast, replacing the integral row-divisibility statement by a rational basis theorem would not give any such cancellation.

25.3 Reduction of the determinant and the integral frame

The matrix Q in the preceding theorem is nonsingular, but its determinant need not be a unit. The weaker property required below is $p \nmid \det Q$. It is stronger than nonsingularity and cannot be deduced from the finite freeness of P alone. We first give the fixed-coefficient argument of `PeriodMatrixFrame.lean`.

Lemma 25.5 (Minimal-support vector)**.** *Let V be a valuation domain and let $M \in M_d(V)$ have zero determinant. There exists $a \in V^d$ in its kernel, with $a_\nu = 1$ for some ν, such that for every unital ring endomorphism g of V,*

$$Mg(a) = 0 \quad \Longrightarrow \quad g(a) = a.$$

Proof. There is a nonzero vector in the kernel by solving over the fraction field and clearing denominators. Choose one with minimal support. In a valuation domain the principal ideals of its nonzero coordinates are totally ordered. Choose a coordinate dividing every other coordinate and divide it out, obtaining $a_\nu = 1$ without changing the support or the kernel equation. If $g(a) - a$ were a nonzero kernel vector, its support would be contained in the support of a but would not contain ν, a contradiction. □

Write $V = A/p = \mathcal{O}^\flat_{C_p}$ and let a bar denote reduction modulo p. We need two precise chart properties:

(i) For $j \geq 0$, let $\pi_j : V \to \mathcal{O}_{C_p}/p$ be the j-th projection of the inverse limit along Frobenius. If $a \in V$ is fixed by every Galois element fixing S, then for each m and all sufficiently large n there is $s \in S$ with
$$\pi_m(a) = \pi_{m+n}(\bar{s}).$$

(ii) If $x \in A$ satisfies $\pi_j(\bar{x}) = 0$, every original coordinate $\lambda_i(x) \in S$ has $\pi_j(\overline{\lambda_i(x)}) = 0$ as well.

The first is the fixed-coefficient approximation in the chosen chart, `Chart.eventually_exists_coeff_eq_of_fixed`; it is an assertion about the arithmetic fixer of the Kummer tower, not just an arbitrary action on a ring. The second is the chart's divisibility-preservation result `exists_frobE_dvd_sub_of_dvd`, applied with multiplier p. We retain (i) and (ii) as chart comparison inputs. They are needed only at the indicated place and at finite Witt length, not as a statement that all fixed elements of the tilt lie in the residue of S.

Proposition 25.6 (Frobenius and p-primitivity)**.** *Suppose $Q \in M_d(A)$ is nonsingular and p-saturated over S, meaning that $Qt \in pA^d$ for $t \in S^d$ implies $t \in pS^d$. Suppose there is $B_1 \in M_d(S)$ with $p \nmid \det B_1$ and a relation*

$$QB_1 = U_1\varphi(Q), \qquad U_1 \in M_d(A).$$

Assume moreover that for each Galois element g fixing S, there is a matrix $U_g \in M_d(A)$ with $Q = U_g g(Q)$. With the preceding two chart properties, $p \nmid \det Q$.

Proof. Consider the rows of Q after applying the coordinate maps λ_i entry by entry. These are rows over S and separate S^d/pS^d. In fact, if every coordinate of Qt is divisible by p in S, then Qt is divisible by p in A: approximate by the topological basis and reduce modulo p, using separatedness of V for the topology of $\bar{z}$. Saturation gives $t \in pS^d$.

The separating rows have rank d over the fraction field of S/pS, which embeds into the fraction field of V. Indeed otherwise a nonzero kernel vector over the former field can be

scaled into $(S/pS)^d$; lifting it contradicts the separation just proved. Thus a choice of d such rows gives a matrix $N \in M_d(S)$ with $\overline{\det N} \neq 0$. Each row of Nt is a specified coordinate of one entry of Qt.

Suppose for contradiction that $\det \bar{Q} = 0$. Apply Lemma 25.5 to $\bar{Q}$ over V, obtaining a with $a_\nu = 1$. Each g fixing S carries a into the kernel of $g(\bar{Q})$ and hence of $\bar{Q}$ by $Q = U_g g(Q)$. Minimality gives $g(a) = a$. Define

$$C_n = B_1 \varphi(B_1) \cdots \varphi^{n-1}(B_1), \qquad D = \overline{\det B_1}\, \overline{\det N}.$$

Induction using the Frobenius relation shows that

$$\bar{Q} \bar{C}_n a^{p^n} = 0 \qquad (n \geq 0),$$

where the power is componentwise. Furthermore

$$D \det \bar{C}_n \mid D^{p^n} \quad \text{in } V.$$

To see this, $\det \bar{C}_n = (\overline{\det B_1})^{1+p+\cdots+p^{n-1}}$ and $1 + (p^n - 1)/(p-1) \leq p^n$. The assertion for $n = 0$ is immediate.

Fix any component m in the inverse limit description of V. By (i), choose a common n large enough and $s \in S^d$ with $\pi_{m+n}(\bar{s}) = \pi_m(a) = \pi_{m+n}(a^{p^n})$. In particular $\pi_{m+n}(\bar{s}_\nu) = 1$. Then $\pi_{m+n}(\bar{Q}\bar{C}_n \bar{s}) = 0$. Property (ii) and the selection of the rows of N give $\pi_{m+n}(\bar{N}\bar{C}_n \bar{s}) = 0$. This is a matrix equation over $\mathcal{O}_{C_p}/p$, which need not be a domain. Nevertheless, multiplication by the adjugate shows that $\pi_{m+n}(\det \bar{N} \det \bar{C}_n) = 0$, because one coordinate of $\bar{s}$ has image 1. Consequently $\pi_{m+n}(D \det \bar{C}_n) = 0$. The displayed divisibility gives $\pi_{m+n}(D^{p^n}) = 0$, i.e. $\pi_m(D) = 0$. Every component is zero, so $D = 0$. This contradicts $\det \bar{N} \neq 0$, $\det \bar{B}_1 \neq 0$ in the domain V. □

This proves `not_dvd_det_of_frobenius` from the precise two chart properties. It also explains why the inverse-limit components cannot be replaced by a single residue-field reduction. The matrix N need not be invertible modulo the maximal ideal of S, only nonsingular modulo p. Taking a component of D far enough back in the tilt, multiplied by Frobenius iterates of B_1, is the way the proof obtains a contradiction without assuming either of these determinants is a unit.

We apply this result to Theorem 25.4. First, Q is p-saturated over S: if $Qt = px$ with $x \in A^d$, then $p^{-1}Zt$ solves $Uy = E^h x$ over S_B. Hence $x = Qs$ for $s \in S^d$ and nonsingularity gives $t = ps$. Next,

$$U \frac{F\varphi(Z)}{\varphi(E)^h} = v^{-1}\varphi(Q)$$

follows from $UZ = E^h Q$. Multiplying by E^h yields a period solution over S_B whose corresponding vector $v^{-1}E^h\varphi(Q)$ is integral; to obtain $v^{-1}\varphi(Q)$ itself, keep the undivided equation $UF\varphi(Z) = v^{-1}\varphi(E)^h\varphi(Q)$ and use $E^h\varphi(E)^{-h}$ in the witness. Thus $v^{-1}\varphi(Q) = QB_1$ for B_1 over S. Taking determinants in $\varphi(Q) = vQB_1$ shows $p \nmid \det B_1$: if p divided $\det B_1$, then $p^k \mid \det Q$ would imply $p^{k+1} \mid \varphi(\det Q)$ and hence $p^{k+1} \mid \det Q$, contradicting nonsingularity and p-separatedness.

The Galois relation used in the proposition derives from the compatibility of a crystalline period matrix with the representation and the scalar normalization. More precisely the source gives $\rho_g g(U) = U$ with ρ_g integral when g fixes the chart constants and $\theta(z)$, after the scalar normalization has been included (`IsPeriodMatrix.exists_gal`). For g fixing S, this hypothesis holds and $g(E) = E$. Furthermore g is an S-linear endomorphism of R preserving $A_{\max}$. As the entries of Z belong to the bicommutant, evaluating the commutation

$g \circ m_z = m_z \circ g$ at 1 for each such entry z gives $g(Z) = Z$. Applying g to $UZ = E^h Q$ and multiplying on the left by ρ_g yields $UZ = E^h \rho_g g(Q)$. Canceling E^h gives $Q = \rho_g g(Q)$. This is `IsPeriodMatrix.exists_eq_mul_smul_of_forall_smul_eq`. With the period-matrix compatibility input, Proposition 25.6 applies.

Corollary 25.7 (Uniform scalar inverses)**.** *Under the above hypotheses, every integral solution of $Uy = E^h x$ with $y \in R^d$ lies in QA^d. There are matrices $L \in M_d(A)$ and $Y \in M_d(R)$ with*

$$QL = \tau^h I_d, \qquad U = QY, \qquad YZ = E^h I_d.$$

Proof. Apply the divisibility statement of Theorem 25.4 to every row of the adjugate of Q with $q = \det Q$; this is permitted by $p \nmid \det Q$. Each component of $\operatorname{adj}(Q)x$ is divisible by $\det Q$ in A. This yields $x = Qx'$ with $x' \in A^d$. In particular the solutions $\tau^h e_j$ with witnesses We_j give L.

To obtain Y, use the approximation modulo μ^h of each entry of U. Choose a common exponent m and matrices $K \in M_d(A)$, $H \in M_d(A_{\max})$ with

$$p^m U = K + p^m \mu^h H.$$

Write $\mu = E\tau d_0$ for a unit $d_0 \in A$. Since $UW = (E\tau)^h I_d$,

$$K = Up^m (I_d - d_0^h WH).$$

Therefore every column of K has a witness for the equation $Uy = E^h x$ in R^d after multiplication of the parenthesized matrix by E^h. We obtain $K = QK_0$ for an integral matrix K_0. Now

$$U = Q\left(p^{-m} K_0 + (Ed_0)^h LH\right).$$

This proves $U = QY$; substituting in $UZ = E^h Q$ and canceling Q yields $YZ = E^h I_d$. □

Thus the part of the frame theorem giving the rank-free scalar denominators reduces to the structural chart comparison, the exact period equations, and the preceding proofs. This removes the abstract basis and determinant assertions previously treated as inputs in Section 24.4. The remaining transformation properties for all Galois elements in the chosen chart, the simultaneous Hodge–Tate pole estimates, and the identification of this span with the intrinsic Kisin module are still necessary comparison statements. In the source the complete `frame_theorem` combines these statements with the integral basis and the calculation above. In particular we may not replace Q by an arbitrary rational basis of a filtered φ-module.

25.4 Scalar depth steps at each Frobenius divisor

Return to the jump lattices of Chapter 23. A Kummer chart has $S = \Lambda[[z]]$ and a subgroup G' of prime-to-p index, acting trivially on Λ. Let $E = z^{e_0} A_0(z) + p$ with A_0 a unit of $\Lambda[[z]]$, let $E_a = \varphi^a E$, and put $\Delta = z\tau$. Write $\theta_a = \theta\varphi^{-a}$. The valuation v_p on C_p is normalized by $v_p(p) = 1$, and for an integer $n = 0$ we put $v_p(0) = 0$ only in the cost function. In particular the convention does not extend v_p as a valuation of the field.

The different bound of the source is

$$d_0 = \begin{cases} 0 & e_0 = 1, \\ v_p(e_0) + 1 & e_0 > 1. \end{cases}$$

We recall the exact chart properties used in the calculation. For every nonnegative m there is $g \in G'$ with $g(z) = [\epsilon]^m z$; every $g \in G'$ fixes $\theta(z)$ and Λ. The ideal (E_a) equals the kernel of

θ_a; Δ is divisible by none of the E_a; $\Delta \mid \varphi^a\Delta$. Write G for the unit satisfying $F(z) = EG$, where $F(X)$ is the monic Eisenstein polynomial of the actual uniformizer $\theta(z)$ over $\Lambda[1/p]$. Then

$$\theta(F'(z)) \mid p^{d_0}, \qquad \theta(G) \in \mathcal{O}_{C_p}^{\times}.$$

For an element $g \in G'$ fixing the base constants, $gE = \omega E$ with $\omega \in A^{\times}$ and $\Delta \mid \omega - 1$. The last relation is the scalar near-identity part of the chart, not a general statement about all Galois actions on A; it is established in `zTau_dvd_sub_one_of_mul_eq_smul_E`. We retain these existence and divisibility properties as Kummer chart inputs. In the unramified chart $E = z + p$ and $F(z) = E$, so the polynomial and derivative assertions are immediate. In a general chart the inequality $v_p(F'(\theta(z))) \le d_0$ follows from the different bound: the leading derivative term has valuation $v_p(e_0) + (e_0 - 1)/e_0$ and the nonzero derivative summands of distinct degrees have distinct fractional parts of their valuations, so the valuation of the sum is the minimum and is no larger than the leading valuation. This proves the claimed divisibility for this polynomial; the unit quotient and the action assertions refer to the precise embedding in A.

Lemma 25.8 (First-order scalar estimate)**.** *Choose g with $g(z) = [\epsilon]^{p^a} z$ and write $g(E) = \omega E$. Then*

$$\theta(\omega - 1) \mid p^{d_0+a}\theta(\Delta) \quad \text{in } \mathcal{O}_{C_p}.$$

Proof. Write $g(z) - z = Eb$. This is possible since g fixes $\theta(z)$. Factor

$$[\epsilon]^{p^a} - 1 = ([\epsilon] - 1) \sum_{j=0}^{p^a-1} [\epsilon]^j = \xi\tau \sum_{j=0}^{p^a-1} [\epsilon]^j.$$

Because E and ξ are associated and $\theta([\epsilon]) = 1$, this gives $\theta(b) = p^a\theta(\Delta)u$ for a unit u of $\mathcal{O}_{C_p}$. Apply the Taylor identity to F:

$$F(gz) = F(z) + F'(z)Eb + E^2b^2H \qquad (H \in A).$$

Since g fixes the coefficients of F, the left side is $g(EG) = \omega E g(G)$. Divide by E and apply θ, using $\theta(gG) = \theta(G)$. We obtain

$$\theta(\omega - 1)\theta(G) = \theta(F'(z))\theta(b).$$

The derivative divides p^{d_0}, and $\theta(G)$ is a unit. Substituting the expression for $\theta(b)$ proves the result. □

This proof is the scalar calculation `theta_sub_one_dvd_pow_mul_theta_zTau` in the readable `ExponentLattices.lean`. At higher places, we use the same calculation after applying inverse Frobenius, because E_a generates that kernel. Keeping the chart polynomial in the argument (rather than replacing E by an arbitrarily associated element) is why a bound in terms of the different remains possible.

Lemma 25.9 (Lifting the exponent at p)**.** *Suppose $\omega \in A^{\times}$, $\Delta \mid \omega - 1$, and $\theta(\omega - 1) \neq 0$. Then, for every nonzero integer n,*

$$v_p(\theta(\omega^n - 1)) = v_p(\theta(\omega - 1)) + v_p(n).$$

Proof. Put $x = \theta(\omega) - 1$. Since $\theta(\tau) = \zeta_p - 1$ and $\theta(z)$ has positive valuation,

$$v_p(x) \ge v_p(\theta(\Delta)) = v_p(\theta(z)) + \frac{1}{p-1} > \frac{1}{p-1}.$$

If $m > 0$ is prime to p, the binomial expansion of $(1+x)^m - 1$ has leading term mx of valuation $v_p(x)$, strictly less than those of all higher terms. For exponent p, the linear term px has valuation $1 + v_p(x)$. The terms $\binom{p}{j}x^j$ for $1 < j < p$ have strictly greater valuation; the last term x^p does also, since $(p-1)v_p(x) > 1$. Thus raising to p increases the valuation by exactly one. Iterate and use $n = p^r m$. Replacing positive n by $-n$ changes $(1+x)^n - 1$ by a unit times its negative and preserves the valuation. □

The strict inequality in this proof matters: the classical formula can fail for a p-th root of unity where $v_p(x) = 1/(p-1)$. Multiplication of the cyclotomic element by the chart coordinate z is what supplies strictness. In the application of Lemma 25.8, $\theta(\omega - 1)$ cannot be zero: the formula expresses it as a nonzero derivative times the nonzero scalar $\theta(b)$ divided by a unit. In particular ω does not have finite order in the residue field of E.

Proposition 25.10 (The scalar depth identity)**.** *For each $a \geq 0$ and nonzero $n \in \mathbb{Z}$, put*

$$c(a,n) = v_p(n) + d_0 + 3a.$$

There are $g \in G'$, a unit $\omega_a \in A^\times$ and $X, V \in A$ satisfying

$$g(E_a) = \omega_a E_a, \quad g(s) - s \in E_a A \ (s \in S), \quad \Delta X = \omega_a^n - 1, \quad VX \equiv p^{c(a,n)} \pmod{E_a A}. \tag{25.1}$$

Proof. Use the element g of Lemma 25.8. It commutes with Frobenius, so $\omega_a = \varphi^a \omega$. Since $g(z) - z = z([\epsilon]^{p^a} - 1)$ and $\theta_a([\epsilon]^{p^a}) = 1$, $g(z) - z \in E_a A$. The same holds for each $s \in S$ by its power-series expression and the exact kernel condition; the continuity and power-series evaluation here are those of the chart. The element $\omega^n - 1$ is divisible by $\omega - 1$, also for negative n (because ω is a unit). Thus $\varphi^a(\omega^n - 1)$ is divisible by $\varphi^a\Delta$, hence by Δ; take X to be the quotient.

It remains to invert X modulo E_a at the indicated cost. Since $\theta_a(\Delta) \neq 0$,

$$v_p(\theta_a X) = v_p(\theta(\omega^n - 1)) - v_p(\theta_a \Delta).$$

By Lemmas 25.8 and 25.9, this is at most

$$v_p(n) + d_0 + a + v_p(\theta\Delta) - v_p(\theta_a\Delta).$$

For $a = 0$ the last difference is zero. For $a > 0$, $\theta(\Delta) = \theta(z)(\zeta_p - 1)$ divides p^2 because $v_p(\theta(z)) = 1/e_0 \leq 1$ and $1/(p-1) \leq 1$; the other valuation is nonnegative. Therefore the difference is at most $2a$ for all a. This proves $\theta_a X \mid p^{c(a,n)}$ in $\mathcal{O}_{C_p}$. The map θ_a is surjective, so lift the quotient $p^{c(a,n)}/\theta_a X$ to $V \in A$. Then the congruence holds by its kernel identity. □

In the compact source this is `exists_depth_step`, and in the readable source it is `exists_isDepthStep`. The proof also identifies the source of all three contributions to the integer $c(a,n)$: the power of p dividing n, the field different, and the shift of the divisor together with the comparison of the two evaluations of Δ. No factor involving the rank occurs.

25.5 Depth zero and finite averaging

The preceding proof does not apply when $n = 0$. The unit eigenvalue for degree zero yields no nonzero factor to cancel. The correct replacement is an averaging operator; its cost is $c(a,0) = d_0 + 3a$ as above.

Lemma 25.11 (Averaging modulo E_a)**.** *There are $p^{c(a,0)}$ elements g_j of G' such that the sum $\sum_j g_j(s)$ is fixed modulo E_a by all $g \in G'$ and $s \in S$.*

Proof. Choose $g_j(z) = [\epsilon]^j z$, $0 \le j < p^{c(a,0)}$. The scalar $\zeta = \theta_a([\epsilon])$ is a primitive p^a-th root of unity, with order one if $a = 0$. At θ_a each $s \in S$ can be expanded as a polynomial in $\theta_a(z)$ with coefficients in the unramified constants, of degree less than $e_0 p^a$. This follows from the chart's monic division by E_a: the evaluated scalar $\theta_a(z)$ is a root of the polynomial obtained from the Eisenstein polynomial of $\theta(z)$ by replacing X by X^{p^a} and twisting coefficients by inverse Frobenius. Division leaves a polynomial of at most that degree; the integral division statement in the chart is `exists_dvd_sub_sum_z_pow`.

For a monomial $s = wz^m$ with $w \in \Lambda$, its sum after evaluation is

$$\theta_a(w)\theta_a(z)^m \sum_{j=0}^{p^{c(a,0)}-1} (\zeta^m)^j.$$

Since $c(a,0) \ge a$, the geometric sum is zero unless $p^a \mid m$. In that case the monomial is fixed by all G': $\theta_a(z)^{p^a} = \theta(z)$, and G' fixes this element and the constants. The sum of the surviving terms is therefore fixed. This is equivalent to the congruence modulo E_a. □

Notice that it is a sum, not an average divided by $p^{c(a,0)}$; the latter would not be integral. The source lemma `exists_frobE_dvd_smul_sum_sub_sum` is precisely this statement.

25.6 From scalar depths to crystalline lattice depths

We make explicit the notation used to apply these scalar identities uniformly to a lattice. A jump lattice T is a finite free $\mathbb{Z}_p$-module with an integral Galois action (and, in the applications, compatible actions of a finite free coefficient ring). Over $W(C_p^\flat)$ we have its ambient module $\mathcal{T}$. The chart identifies a Kisin lattice with its localized space $\mathcal{N}(T)$, over $S[E_a^{-1} : a \ge 0]$. Let $\mathcal{N}_A(T)$ denote its A-span, also with the indicated divisors inverted in the localized coordinates. The fixed chart coefficient ring is denoted by Λ, and the letters i are the factors of the coefficient-ring decomposition, permuted by Frobenius. For a letter i, $i[-a]$ means a inverse Frobenius steps. We write $U_i(T)$ for the finite jump set containing the actual Hodge–Tate jumps. These sets are allowed to include further integers; in particular, 0 occurs in each extension of a trivial lattice.

Every crystalline form of type u at letter $i[-a]$ defines a value $\mathcal{V}_a(\lambda, x)$ in B_{dR} on $x \in \mathcal{N}_A(T)$ at letter i. This is the evaluation of λ after φ^{-a}. For a system $o(a,u)$ of integers, write $\mathcal{M}_i^T(o)$ for the submodule consisting of the x satisfying

$$\mathcal{V}_a(\lambda, x) \in \mathrm{Fil}^{o(a,u)} \quad (a \ge 0,\ u \in U_{i[-a]}(T),\ \lambda \text{ of type } u). \tag{25.2}$$

This is `JumpLattice.withOrders`. We use the following comparison inputs, whose precise strength is important.

(a) Values commute with the G'-action and with multiplication by evaluated scalars $\varphi^{-a}(r)$ ($r \in A$). The modules $\mathcal{M}_i^T(o)$ are G'-stable, and multiplication by E_b^m adds m to the order at $a = b$, for any integer m, without affecting the other order conditions.

(b) There is a coefficient retraction $r_0 : \mathcal{N}_A(T) \to \mathcal{N}(T)$, preserving all the conditions (25.2), commuting with projections onto letters and with equivariant maps between the lattices, and preserving congruences of values to constants as follows. If $\mathcal{V}_a(\lambda, x)$ is congruent to a scalar in the fixed field modulo Fil^n, then the same congruence holds for $r_0(x)$.

(c) For $x \in \mathcal{N}(T)$ and $g \in G'$, the vector $\Delta^{-1}(gx - x)$ belongs to $\mathcal{N}_A(T)$ and, if $gx - x \in \mathcal{M}_i^T(o)$, then the quotient belongs to $\mathcal{M}_i^T(o)$ as well.

(d) Let K' be the field fixed by the chart subgroup. If $y \in B_{\mathrm{dR}}$ is fixed modulo Fil^n by G', then y is congruent modulo Fil^n to an element of K'. Moreover the value of each form on a localized Kisin vector lies in the specified value field for that divisor; whenever g acts trivially on S modulo E_a, it acts trivially on the integral elements of this value field modulo Fil^1. The same holds for the sum in Lemma 25.11: the sum of conjugates of such integral values is invariant modulo Fil^1.

Part (a) follows from the evaluations at the distinct prime divisors and equivariance of forms. Parts (b), (c), (d) are the integral frame and fixed-scalar comparisons: the source anchors are `coordRetract_mem_withOrders`, `valueAt_coordRetract_sub_mem_hodgeFil`, `zTauUnit_inv_smul_sub_mem_withOrders`, `exists_const_sub_mem_hodgeFil_of_smul_sub_mem`, `smul_sub_mem_hodgeFil_of_valueField` in `ExponentLattices.lean` and `HodgeExpansion.lean`. The assertion on K' extends to negative n: it is a statement about fixed representatives in a filtered quotient, stronger than saying $B_{\mathrm{dR}}^{G'} = K'$. For example, a nonzero K'-scalar has filtration order zero, regardless of its p-adic valuation. All remaining derivations are carried out under these comparison statements. The scalar inverse calculations earlier in this chapter explain how uniform frame denominators used in (b),(c) are obtained; deriving the actual comparison of every form in (d) still requires the filtered period-ring theory.

Define the depth condition

$$\mathcal{I}_i^T(o) = \{x : \ gx - x \in \mathcal{M}_i^T(o) \text{ for all } g \in G'\},$$

where x is in the same letter part of the ambient localized span. By (a) and (d), this is exactly the condition that every $\mathcal{V}_a(\lambda, x)$ be congruent to a fixed-field scalar modulo $\mathrm{Fil}^{o(a,u)}$. The reverse implication uses $g(y - c) - (y - c) \in \mathrm{Fil}^{o(a,u)}$ if c is fixed.

Lemma 25.12 (Improvement of positive orders)**.** *Let $x \in \mathcal{M}_i^T(o) \cap \mathcal{I}_i^T(h)$. Then its order system can be improved to*

$$\mathrm{imp}(o, h)(a, u) = \begin{cases} h(a, u) & o(a, u) > 0, \\ \min(0, h(a, u)) & o(a, u) \le 0. \end{cases}$$

Proof. Fix a, u, λ and write $y = \mathcal{V}_a(\lambda, x)$, $o = o(a, u)$, $h = h(a, u)$. We have $y \in \mathrm{Fil}^o$ and $y - c \in \mathrm{Fil}^h$ for some $c \in K'$. Thus $c \in \mathrm{Fil}^{\min(o,h)}$. If both o and h are positive, $c = 0$, because a nonzero constant is a unit of B_{dR}^+. If $h \le 0$, every constant lies in Fil^h and $y \in \mathrm{Fil}^h$. If $o \le 0 < h$, the congruence only implies $y \in \mathrm{Fil}^0$. These are exactly the displayed cases. □

The source writes this as `mem_withOrders_improvedOrders`. In any use one can keep the stronger original order if the displayed imp happens to be smaller, because both order systems hold. The ability to turn a small positive order into an arbitrarily large depth bound is the decisive simplification in the long-gap argument.

Theorem 25.13 (A rank-free depth operator)**.** *Let $a \ge 0$, $n \in \mathbb{Z}$ and $c = c(a, n)$. There is an additive operator on $\mathcal{N}(T)$ of the shape*

$$B_{a,n}(x) = p^c x + \sum_j r_0(v_j \Delta^{-1}(g_j x - x)), \qquad v_j \in A, quad g_j \in G',$$

with the following properties.

(i) For any form λ, if $\mathcal{V}_a(\lambda, gx - x) \in \mathrm{Fil}^n$ for all $g \in G'$, the same assertion with $B_{a,n}x$ holds in Fil^{n+1}.

(ii) If x has depth o at letter i, then $B_{a,n}x - p^c x$ has orders o. Every depth condition already satisfied by x is still satisfied by $B_{a,n}x$.

(iii) If x maps to a fixed constant in a trivial quotient lattice and the retraction commutes with that quotient, then $B_{a,n}x$ maps to the constant multiplied by p^c.

Proof. For $n \neq 0$, choose g, ω_a, X, V as in Proposition 25.10. Let π stand for the de Rham image of E_a under φ^{-a}; it generates the maximal ideal. After subtracting the fixed scalar from the value of x, write $y = \mathcal{V}_a(\lambda, x) - s \in \mathrm{Fil}^n$. At this point y belongs to the value field from (d). Thus $\pi^{-n}y$ is integral in that field and g fixes it modulo Fil^1, since g acts trivially on S modulo E_a. The equation $g(\pi) = \widetilde{\omega}_a \pi$ (with a tilde denoting evaluation after inverse Frobenius) gives

$$g(y) - \widetilde{\omega}_a^n y \in \mathrm{Fil}^{n+1}.$$

This is also valid for negative n: multiplication by $g(\pi)^n$ shifts the first filtration step by n. Set $w = \Delta^{-1}(gx - x)$. Since $\widetilde{\Delta}$ is a unit in B_{dR}^+, the equation $\Delta X = \omega_a^n - 1$ implies

$$\mathcal{V}_a(\lambda, w) - \widetilde{X} y \in \mathrm{Fil}^{n+1}.$$

Use $B_{a,n} = p^c \mathrm{id} - r_0(V\Delta^{-1}(g-1))$. Before retraction its value differs from $p^c s$ by

$$p^c y - \widetilde{V}\mathcal{V}_a(\lambda, w) = -\widetilde{V}(\mathcal{V}_a(\lambda, w) - \widetilde{X} y) - (\widetilde{V}\widetilde{X} - p^c)y,$$

which belongs to Fil^{n+1} by (25.1). The retraction does not change this congruence by (b). This proves (i) via the equivalence between depth and congruence to a constant.

For $n = 0$, use the $m = p^{c(a,0)}$ elements g_j of Lemma 25.11, choosing $v_j = \Delta$ for each j. Before retraction the operator is

$$p^c x + \sum_{j=1}^{m} (g_j x - x) = \sum_{j=1}^{m} g_j x.$$

The value of x is integral, since it is congruent to a fixed scalar modulo Fil^0. By the value-field comparison in (d) and the scalar averaging lemma, the new value is invariant modulo Fil^1. It is therefore congruent to a fixed scalar, and the retraction preserves that congruence.

For (ii), each difference $g_j x - x$ has orders o. Division by Δ and retraction preserve these orders by (b),(c), as does multiplication by the integral scalar v_j. Hence the operator differs from $p^c x$ by an element of $\mathcal{M}_i^T(o)$. Since that submodule is G'-stable, the operator preserves the depth condition. Finally the image of each difference $g_j x - x$ in the trivial quotient is zero when x lies over a fixed constant. Retraction commutes with the projection, giving (iii). □

The theorem is the combination of `exists_valueAt_sub_mem_hodgeFil_succ` and `exists_isOperatorOfCost_mem_hodgeFil_succ`. It proves the one-step moving assertion in Chapter 23, conditional only on the stated chart and form comparisons. In particular the same operator works for every form and vector; one does not construct a separate operator for each cocycle using its individual coefficients. This uniformity is indispensable when the sum of low jump multiplicities is much smaller than the rank of the ambient extension.

25.7 From a single operator to the cost of a window

The notation of a side is as follows. Given a crystalline lattice Y, a trivial parameter lattice C, and a family of crystalline cocycles $\alpha_c : G' \to Y$ linear in $c \in C$, the extension lattice $\mathcal{E} = Y \oplus C$ has action

$$g(y, c) = (gy + \alpha_c(g), c).$$

Write $q : \mathcal{E} \to C$ for the equivariant projection, and c_i for the constant vector $1 \otimes c$ projected to letter i in the corresponding ambient trivial lattice. The set of points over $p^r c_i$ with orders o and depth d consists of the elements $x \in \mathcal{N}(\mathcal{E})$ at letter i such that

$$qx = p^r c_i, \qquad x \in \mathcal{M}_i^{\mathcal{E}}(o) \cap \mathcal{I}_i^{\mathcal{E}}(d).$$

We call a scalar $p^r E^s$ a cokernel killer at i if it carries the letter component of the trivial integral lattice into the image of the standard order module for $\mathcal{E}$ at that letter. The standard system at level zero is $\sigma_0(a, u) = -u$ for all a. As in Section 24.7, existence of the required scalar kills is an assertion about the image of a module, not a pointwise choice of lifts depending on c with no uniformity. Applying the coefficient retraction gives actual localized Kisin points, so multiplication by E^{-s} gives a point over $p^r c_i$ with orders and depth

$$\rho_{-s}(a, u) = -u - s\mathbf{1}_{a=0}.$$

Indeed E^{-s} affects the order only at $a = 0$, and each order module is G'-stable. This proves the point-existence implication of the cokernel definition (`pointsOver_nonempty_of_kills Cokernel`) from the coefficient-preserving form of that definition.

Proposition 25.14 (A common operator for finitely many steps)**.** *Let $J \subset \mathbb{N} \times \mathbb{Z}$ be finite and write $C(J) = \sum_{(a,n)\in J} c(a, n)$. For every side there is an additive operator B_J on its localized Kisin module such that, if x lies over a fixed constant and is of depth d at i, then*

$$B_J x - p^{C(J)} x \in \mathcal{M}_i^{\mathcal{E}}(d) \cap \ker q.$$

If J contains all (a, n) for which $d(a, u) \le n < d'(a, u)$ for some $u \in U_{i[-a]}$, then $B_J x$ has depth d'.

Proof. Order the pairs in J by nondecreasing second coordinate n, with any order among pairs with equal n. Compose the operators of Theorem 25.13 in that order. Each step preserves existing depth and the condition of being over a constant. A correction of cost c differs from multiplication by p^c by a vector of the old orders and zero projection. If B_1 and B_2 are two such operators, their correction relative to multiplication by $p^{c_1+c_2}$ can be written as

$$(B_2 B_1 x - p^{c_2} B_1 x) + p^{c_2}(B_1 x - p^{c_1} x),$$

so it has the old orders and zero projection. This proves the first assertion by induction.

Fix a, u and a form of type u. If $d'(a, u) > d(a, u)$, each integer step from the latter to the former occurs in the prescribed ordering. By property (i) of Theorem 25.13, traversing a step n upgrades its depth from at least n to at least $n + 1$. All intervening steps at other divisors or for other jumps preserve the depth. Thus after the last step the required depth holds for this form. This is true for all forms and all jumps, including those at which no step is required. □

This proves `exists_isOperatorOfCost_and_mem_deepElements`. In particular for a point over $p^r c_i$ with orders o and depth d, the operator provides a point over $p^{r+C(J)} c_i$ of the same orders and improved depth d'. To verify that the old orders remain, use the first assertion of

Proposition 25.14 with initial depth o; membership of x in $\mathcal{M}_i(o)$ already implies this depth by stability. Then Lemma 25.12 improves the orders to $\mathrm{imp}(o, d')$. Multiplication by further powers of p does not change any filtration order. This proves `pointsOver_nonempty_of_costs`. If many different jump intervals overlap, one pays only for their union J; no rank factor or multiplicity count is inserted here.

Here is the elementary uniform bound on the sum of costs of two integer windows. Define

$$J(n, b) = [-n, n] \cap \mathbb{Z} \ \cup \ [b - n, b + n] \cap \mathbb{Z}.$$

Define $v_p(0) = 0$ in sums and let $H \geq 1$ bound the absolute values of members of $J(n, b)$. The number of nonzero integers divisible by p^j in an interval of $2n + 1$ consecutive integers is at most $\lfloor 2n/p^j \rfloor + 1$. Counting the contribution of each p^j gives

$$\sum_{m \in J(n,b)} v_p(m) \leq \frac{4n}{p - 1} + 2\lfloor \log_p H \rfloor. \tag{25.3}$$

The use of the floor is justified because terms for $j > \log_p H$ vanish. The inequality remains valid if the windows overlap, by counting the union only once on the left and bounding by both individual sums. Consequently

$$\sum_{m \in J(n,b)} c(a, m) \leq (4n + 2)(d_0 + 3a) + \frac{4n}{p - 1} + 2\lfloor \log_p H \rfloor.$$

The readable source uses the variant `sum_twoWindows_le` obtained from the same count (in `SidePoints.lean`). For $p \nmid m$, the first term of the cost is zero; thus on the no-multiple windows of the few-low-jumps argument the bound simplifies to $c(a, m) \leq 3a$ for an unramified chart. The exceptional degree $m = 0$ obeys the same bound, by the averaging rather than the lifting-the-exponent construction of Section 25.5. It is not legitimate to apply the latter construction to $m = 0$ by writing $v_p(0) = 0$.

25.8 The monic Taylor coefficient at a long gap

A gap in the jumps is used in a different way from a no-multiple window. Fix an unramified chart $E = z + p$ and let i be a letter. For $l \geq 0$ define the standard order system at level l,

$$\sigma_l(a, u) = \begin{cases} 0 & a < l, \\ -u & a \geq l. \end{cases}$$

Let $M_l^T(i)$ denote the localized Kisin part of $\mathcal{M}_i^T(\sigma_l)$; it is an S-module. For $A \geq 0$, call $x \in M_l^T(i)$ a constant modulo E^A if it also belongs to

$$\mathcal{I}_i^T(\sigma_l + A\mathbf{1}_{a=0}).$$

Call this $\mathbb{Z}_p$-submodule $D_l^T(i, A)$. This is not the ordinary constant term of a formal power series of lattice vectors: the definition refers to near-invariance under G'. It is a useful auxiliary space because of the improvement of positive orders in Lemma 25.12.

For $l > 0$ and a function o of the jumps at i, let $\sigma_l[o]$ be the system $\sigma_l(a, u) + o(u)\mathbf{1}_{a=0}$. The following coefficient extraction lemma will be applied twice: once to an extension and once to its trivial quotient.

Lemma 25.15 (Taylor coefficient and a jump gap)**.** *Let $l > 0$. Suppose $\xi_0, \dots, \xi_N \in D_l^T(i, A)$ and $y \in M_l^T(i)$, and*

$$\sum_{n=0}^{N} E^n \xi_n + E^{N+1} y \in \mathcal{M}_i^T(\sigma_l[o]), \qquad o(u) \le A.$$

Then $\xi_N \in \mathcal{M}_i^T(\sigma_l[o_N])$, where

$$o_N(u) = \begin{cases} A & N < o(u), \\ 0 & N \ge o(u). \end{cases}$$

Proof. First consider $N = 0$. Since Ey has order at least one at $a = 0$ and the other standard orders, ξ_0 has order at least $\min(1, o)$ at $a = 0$. Here $\sigma_l(0, u) = 0$ because $l > 0$. If $o(u) > 0$ the order is positive and the depth A upgrades it to A, by Lemma 25.12; otherwise ξ_0 already has order zero there. This gives o_0.

Now write the sum as $\xi_0 + Ey_1$, with $y_1 = \sum_{n=0}^{N-1} E^n \xi_{n+1} + E^N y$ in $M_l^T(i)$. By the first case ξ_0 has order A at all jumps where $o > 0$; at the remaining jumps order zero suffices for $o \le 0$. Thus it belongs to $\mathcal{M}_i^T(\sigma_l[o])$ because $o \le A$. Subtract it and divide by E to get $y_1 \in \mathcal{M}_i^T(\sigma_l[o - 1])$. The inductive hypothesis gives the result for ξ_N, since $o(u) - 1 > N - 1$ exactly when $o(u) > N$. □

This is the proof of `mem_withOrders_of_add_E_smul_mem` and `mem_withOrders_of_expansion_mem`. It relies on $l > 0$; if $l = 0$ the standard orders at E are not uniformly zero and the same positive-order improvement cannot be made in this form.

We next produce the required expansions. The integral radius estimate (Section 24.5) says that for $3 \le p^l$ and $A \le p^l \delta$ one has

$$p^{4\delta} M_l^T(i) \subset \operatorname{span}_S D_l^T(i, A). \tag{25.4}$$

Precisely this follows from `frame_theorem.pow_smul_mem_span_of_mem_valueFieldPoints` after projecting to a letter. It uses the frame identification with the value-field points and the first part of the radius lemma, which still rely on the integral comparison; the elementary matrix truncation yielding the fourth power was proved in Section 24.5. We make (25.4) an explicit hypothesis for the long gap, so no assertion about arbitrary vectors of arbitrary period modules is implicit here.

Since the chart is unramified, $E = z + p$ and $S/(E) \simeq \Lambda$, each $s \in S$ can be expanded modulo every E^M in a sum with coefficients in Λ. Repeatedly choose the residue in Λ and divide the remainder by E to get

$$s = \sum_{n<M} w_n E^n + E^M s_M, \qquad w_n \in \Lambda,\ s_M \in S.$$

The space $D_l^T(i, A)$ is stable under Λ, since G' fixes the constants and multiplication by $w \in \Lambda$ preserves the orders. From (25.4) write $p^{4\delta} x = \sum_j s_j b_j$ with finitely many $b_j \in D_l^T(i, A)$. Expand the s_j and collect the coefficients:

$$p^{4\delta} x = \sum_{n<M} E^n \xi_n + E^M y_M, qquad \xi_n \in D_l^T(i, A), \quad y_M \in M_l^T(i). \tag{25.5}$$

The coefficients $\xi_n = \sum_j w_{jn} b_j$ are chosen once, independent of M. This is `exists_taylor_expansion_of_mem_stdOrdersPart`. It requires no convergence assertion about an infinite series in B_{dR}; we only extract a finite coefficient of this family of identities.

Theorem 25.16 (Long-gap lifting lemma)**.** *Let S be a side with extension $\mathcal{E}$ and parameter lattice C. Assume (25.4) for its extension and trivial quotient, and the chart comparison assertions of Section 25.6. Let i be a letter and suppose*

(i) *E^s kills the cokernel at i, and p^κ kills the cokernel at $i[-l]$;*

(ii) *$l > 0$, $3 \le p^l$ and $A \le p^l\delta$;*

(iii) *all jumps of $\mathcal{E}$ at $i[-a]$, $0 \le a < l$, are at most g;*

(iv) *$N = s + g(1 + p + \cdots + p^{l-1}) < A$.*

Then for each $c \in C$ there is a localized Kisin point over $p^{\kappa+4\delta}c_i$, with orders

$$\sigma_l(a,u) + A\mathbf{1}_{a=0,\,u<-N}$$

and depth $\sigma_l(a,u) + A\mathbf{1}_{a=0}$.

Proof. The cokernel at i gives a point y over c_i with orders $\rho_{-s}(a,u) = -u - s\mathbf{1}_{a=0}$. In the unramified chart

$$E_a = z^{p^a} + p = (E-p)^{p^a} + p.$$

Hence the polynomial

$$f(X) = X^s \prod_{a<l} \left((X-p)^{p^a} + p\right)^g \tag{25.6}$$

is monic in $\mathbb{Z}[X]$ of degree N. Multiplying y by $f(E)$ first cancels the pole of order s at E and then adds g at each of the divisors E_a with $a < l$. Thus $v = f(E)y$ has orders

$$-u + g\mathbf{1}_{a<l}.$$

Since the jumps at these places are at most g, $v \in M_l^{\mathcal{E}}(i)$, and at $a = 0$ one also has order at least $\min(A, g-u)$.

Apply (25.5) with $M = N+1$ to this v. Lemma 25.15 with $o(u) = \min(A, g-u)$ shows that the coefficient ξ_N has orders at least σ_l shifted by A at every jump $u < -N$, because both $A > N$ and $g - u > N$. It has uniform depth $\sigma_l + A\mathbf{1}_{a=0}$ by membership in $D_l^{\mathcal{E}}(i, A)$.

Apply the trivial-quotient map q to the Taylor identity. By the functoriality of order modules and constants modulo E^A, all the coefficient images satisfy the corresponding conditions in the trivial lattice. The projected sum is $p^{4\delta}f(E)c_i$. Subtract the coefficients of that polynomial, putting

$$\eta_n = q\xi_n - p^{4\delta}f_n c_i, qquad f(X) = \sum f_n X^n.$$

These η_n are constants modulo E^A in the trivial lattice (the projected c_i is genuinely fixed). Their sum $\sum_{n=0}^{N} E^n\eta_n + E^{N+1}qy_{N+1}$ is zero, hence of arbitrarily high orders. Applying Lemma 25.15 with $o(0) = A > N$ gives

$$w = q\xi_N - p^{4\delta}c_i \in \mathcal{M}_i^C(\sigma_l + A\mathbf{1}_{a=0}).$$

Monicity of f is used exactly in obtaining coefficient $f_N = 1$. Merely bounding the degree of f without ensuring its leading coefficient is a unit would not give the required projection.

The error $E^{-A}w$ lies in $M_l^C(i)$. By the Frobenius transport of the standard order module, p^κ killing the cokernel at $i[-l]$ implies that $p^\kappa E^{-A}w$ lifts at i to $v_0 \in M_l^{\mathcal{E}}(i)$. This is just transport of the killer at level zero to level l under φ^l, and the coefficient retraction from integral solutions to localized Kisin points. Then $p^\kappa\xi_N - E^A v_0$ lies over $p^{\kappa+4\delta}c_i$. The correction term has orders $\sigma_l + A\mathbf{1}_{a=0}$ at all jumps, and by stability is deep for the same system. It cannot reduce the orders on the negative jumps selected by the coefficient and cannot reduce its depth. This is the required point. □

This supplies the long-gap lifting theorem `pointsOver_nonempty_of_long_gap` in the compact and readable `LongGap.lean`. The proof exhibits a rather striking mechanism: a monic product of Frobenius translates of one distinguished scalar has a specified degree, and a Taylor coefficient at that degree simultaneously has a known projection to the trivial lattice and a very large order on sufficiently negative jumps. The large order does not have to be purchased one step at a time. The cost is $p^{\kappa+4\delta}$, where δ need only be about A/p^l. In particular the geometric factor in the upper bound $N = s + g(1 + p + \cdots + p^{l-1})$ is exactly the polynomial degree, not an estimate on the rank of the representation.

25.9 The triple pairing and the safety rule

Let T, T' be two sides with kernels Y, Y', parameter lattices C, C', and cocycles $\alpha_c, \beta_{c'}$. Let $B_0 : Y \times Y' \to \mathbb{Z}/p^m(1)$ be the bilinear congruence pairing, equivariant modulo p^m. It comes from an integral form whose failure of equivariance is divisible by p^m. The third side is the extension of the trivial lattice by the Tate dual of $Y \otimes Y'$ associated to that failure of equivariance (and all three sides have their specified integral crystalline data). Its jump set contains 0 and the integers $-u - u' - 1$ with $u \in U_i(Y)$ and $u' \in U_i(Y')$. The -1 is the weight of the Tate twist in the jump convention of these notes. Let Z denote the kernel of the third side, and n a point in its extension over a fixed constant. This side comes with the equivariant extension form

$$\mathrm{ev} : \mathcal{E}_Z \times (Y \otimes Y') \longrightarrow \mathcal{O}(1),$$

whose reduction modulo p^m, followed by the fixed coefficient map to $\mathbb{Z}_p$, has the property that evaluation at a point n over p^κ times the constant 1_i is $p^\kappa B_0$ on the letter i. It is worth verifying that the algebraic extension implementing this form is natural. Write F for an integral bilinear form lifting the congruence pairing (before applying the coefficient map) and write $D_g = (gF - F)/p^m$ for its equivariance defect in the dual. Then D is a cocycle, because

$$D_{gh} = D_g + gD_h.$$

The dual extension has action $(f, c) \mapsto (gf + cD_g, c)$; the functional

$$(f, c) \longmapsto cF - p^m f$$

is equivariant and reduces to cF modulo p^m. Changing the sign convention for D changes the sign of the correction. This verifies the algebraic model of the third side; crystallinity of that extension follows from the specified defect cocycle, which must itself satisfy the finite condition. The source uses precisely this model in `PairingSide.tateDualSide`, `extensionForm` and related declarations in `CupProductBound.lean` and `SidePoints.lean`.

The integral property of pairings that is required is the following. Suppose $b : T_1 \otimes T_2 \to T_3$ is an equivariant bilinear map of lattice data. At a given inverse Frobenius place, a form of jump v on T_3 pulls back into sums of products of forms of jumps u, u' on the two factors with $u + u' \le v$, with coefficients in the positive de Rham ring. Consequently if x_1, x_2 have orders o_1, o_2 at a letter, their pairing has order system o_3 whenever

$$o_3(a, v) \le o_1(a, u) + o_2(a, u') \quad (u + u' \le v) \tag{25.7}$$

at all the corresponding jumps. For this assertion one chooses a filtration-adapted basis of dual forms over the fixed field and expands a form of the stated type in the dual tensor basis. The filtration on the tensor product is the sum filtration, so only the indicated $u + u' \le v$ can contribute; scalar coefficients given by filtered linearity have no negative filtration orders. Evaluation preserves the sums and products, proving the inequality term by term from (25.2).

Membership of the result in the ambient localized Kisin span under b uses the tensor-compatible identification of the integral frames. This is the use of `wittPairing_mem_kisinAinf` and `wittPairing_mem_withOrders` in `HodgeExpansion.lean`; we take that tensor-compatible comparison as an input and use the elementary filtered calculation above.

For the Tate line $\mathcal{O}(1)$ itself, an element of standard orders $\sigma_0(a,-1)=1$ at all places lies, after multiplication by τ, in the A-span of the integral Tate lattice. This is the rank-one case of the frame comparison, `tau_pow_smul_mem_ainfLattice_of_eq_chi_pow_smul`. Its consequence in the mod-p^m target is the smallness property: multiplying once more by z, whose reduction modulo p has strictly positive valuation, makes all Witt coordinates small. Write

$$W_m(C_p^\flat)=W(C_p^\flat)/p^m, qquad \mathcal{K}_m=\operatorname{im}\left(W(\mathfrak{m}_{C_p^\flat})\to W_m(C_p^\flat)\right).$$

Here $W(\mathfrak{m}_{C_p^\flat})$ means the kernel of reduction of integral Witt vectors to the residue Witt ring, an additive subgroup. We use $\mathcal{K}_m(1)$ for the same subgroup in the cyclotomic twist. It is invariant under Galois and Frobenius. Multiplying a vector with integral Witt coordinates by an element of $W(\mathfrak{m}_{C_p^\flat})$ gives a small vector in this sense.

Theorem 25.17 (Three orders imply a safe pair)**.** *Let w,y lie over zero in the extensions of the first and second sides, at the same letter i, with orders o,o' respectively. Identify them in the kernel lattices. Let n be a point of the third side over $p^\kappa 1_{i[-l]}$, with orders v_0 antitone in its second argument. Put*

$$\widetilde{v}_0(a,u)=\begin{cases}0 & a<l,\\ v_0(a-l,u) & a\ge l.\end{cases}$$

Suppose that for all a and all jumps u,u' of the first and second sides at $i[-a]$,

$$\widetilde{v}_0(a,-u-u'-1)+o(a,u)+o'(a,u')>0. \tag{25.8}$$

Then

$$p^\kappa B_0(\Delta w,y)\in\mathcal{K}_m(1),\qquad p^\kappa B_0(w,\Delta y)\in\mathcal{K}_m(1).$$

Proof. Frobenius transports n to $\varphi^l n$ at letter i, with the indicated shifted orders. (The Frobenius shift of orders and letter parts is one of the frame identities of the chart.) Define $\zeta=w\otimes y$ in the tensor kernel. If $s\ge u+u'$, then $-s-1\le -u-u'-1$ and antitonicity gives

$$\widetilde{v}_0(a,-s-1)\ge\widetilde{v}_0(a,-u-u'-1).$$

The strict integer inequality (25.8) and (25.7) imply that ζ has orders at least $1-\widetilde{v}_0(a,-s-1)$ at the jump s. Pair ζ with $\varphi^l n$ under the extension form. For jumps s of the tensor and t of the dual with $s+t\le -1$, antitonicity gives $\widetilde{v}_0(a,t)\ge\widetilde{v}_0(a,-s-1)$, so the sum of the two order functions is at least 1. The Tate-line value therefore has standard orders at every place. Multiplication by τ makes it integral by the rank-one comparison, then multiplication by z makes its mod-p^m image small. By the reduction property of the extension form, that image is $z\tau p^\kappa B_0(w,y)$, as desired. Bilinearity permits either factor to absorb $\Delta=z\tau$. □

This is the proof of `isSafe_of_orders` in readable `SafeMoves.lean`. In applications w is a difference of points with the same trivial coordinate. Such differences lie in the kernel and have at least the pointwise minimum of their two order systems. To move the first component by w while a point y rests in the second extension, it suffices that for every $g\in G'$, $p^\kappa B_0(w,(gy-y)_1)$ be small, where $(\,)_1$ denotes the kernel coordinate. If the rest point has depth d', then $\Delta^{-1}(gy-y)$ has orders d' by the integral comparison of Section 25.6; it lies over zero. We apply Theorem 25.17 with $o'=d'$. For movement of the second component

while x rests, the required expression is $p^{\kappa}B_0((gx-x)_1, gw)$. The difference of x gives its depth and gw has the same orders as w by stability.

When the third side has the cokernel killer $p^{\kappa}E^t$ at $i[-l]$, we may choose $v_0 = \rho_{-t}$. In that case the exact numerical rule sufficient for a move is

$$R(l,t;a,s) = \begin{cases} 1 & a < l, \\ t-s & a = l, \\ -s & a > l. \end{cases} \tag{25.9}$$

The sum of the moving order and resting depth at jumps u, u' must be at least $R(l,t;a,u+u')$. Indeed the shifted third order at $-u-u'-1$ is zero below l, is $u+u'+1-t$ at l, and is $u+u'+1$ above it. In each case its sum with the right side of (25.9) is exactly 1. This proves the rule in `isSafeMoveFst_of_killsCokernel` and `isSafeMoveSnd_of_killsCokernel`. Notice the change in condition below l: the sum need only be positive there, not equal to a large weight-dependent expression. This is the reason for carrying previously chosen points forward along Frobenius before making a move. It cannot be seen from a pointwise estimate at a single de Rham place.

25.10 From safe pairs to an actual cup coboundary

One must justify that safety in the small Witt subgroup has a cohomological consequence over the original finite Galois module. We prove the pairing lemma itself in an elementary form; no rational duality alone supplies it.

First, work with $\mathcal{H} = W_m(C_p^{\flat})(1)$ and its subgroup $\mathcal{K}_m(1)$. Frobenius φ acts on Witt coordinates and is trivial on the finite Tate factor. It commutes with the Galois action.

Lemma 25.18 (Artin–Schreier–Witt modulo the small subgroup)**.** *The map $\varphi - 1 : \mathcal{H} \to \mathcal{H}$ is surjective. An element $x \in \mathcal{H}$ with $(\varphi-1)x \in \mathcal{K}_m(1)$ is uniquely congruent modulo $\mathcal{K}_m(1)$ to a constant in $\mathbb{Z}/p^m(1)$.*

Proof. Fix a generator of the Tate twist to work with ordinary Witt coordinates. In length one, $\varphi - 1$ is $x \mapsto x^p - x$, surjective because $C_p^{\flat}$ is algebraically closed. Inductively, solve for the first coordinate of a desired image, subtract its Teichmüller lift's image and divide the remainder by p to reduce by one length. The division here is carried out in the Witt ring over the perfect field; its ambiguity modulo p^m is exactly removed by reducing the quotient modulo p^{m-1}.

For the assertion modulo small, first consider length one. If $x^p - x$ is of strictly positive valuation, then x is integral: if $v(x) < 0$ its valuation would be $pv(x) < 0$. The residue of an integral such x satisfies $\bar{x}^p = \bar{x}$, hence lies in $\mathbb{F}_p$. Subtract a Teichmüller constant lifting this residue. The first coordinate is now small. Subtract its Teichmüller lift (which is itself small) and divide by p; the result of applying $\varphi - 1$ remains small modulo the shorter length, because the kernel of the residue map is saturated for p in this sense. This is immediate from the sequence of Witt coordinates in the perfect residue field: the vanishing of all coordinates of the residue after multiplication by p implies the vanishing of the first $m-1$ coordinates before multiplication. Induction then makes the entire vector congruent to a sum of Teichmüller constants times powers of p, i.e. an element of $\mathbb{Z}/p^m$. Uniqueness follows because reduction of a nonzero constant mod p^m in the residue Witt ring is nonzero. □

In addition, for every fixed $x \in \mathcal{H}$, $gx - x$ is small for g sufficiently near 1 in the Galois topology. Indeed the Galois action on the finite number of Witt coordinates is continuous

for the valuation topology on $C_p^\flat$, and addition and subtraction of Witt vectors are given by polynomials in the coordinates. At the identity all the coordinates of the difference vanish, so they lie in the maximal ideal for a neighborhood of the identity. The cyclotomic action is trivial modulo p^m sufficiently near 1. The same argument applies to fixed vectors in a tensor product with a finite free integral representation, after reducing it modulo p^m. This proves the near-identity smallness used in the source; it does not assert that the Galois action on all of $W_m(C_p^\flat)$ is discrete.

Theorem 25.19 (The pairing lemma via a four-move path)**.** *Let x, w be points of the first extension over the global (all-letter) constant $p^{e_1}c$, and x', w' points of the second over $p^{e_2}c'$. Let $e_3 \geq 0$. Assume that the four moves from $(\varphi x, \varphi x')$ through $(w, \varphi x')$, (w, w') and (x, w') to (x, x') are safe at cost p^{e_3}, and that for every $h \in G'$ the move of the second component by $hx' - x'$ while x rests is safe. Then, for every $N \geq e_1 + e_2 + e_3$, $p^N[\alpha_c \smile_{B_0} \beta_{c'}] = 0$ in continuous cohomology of the original local Galois group with coefficients in the finite Tate module.*

Proof. Write u, u_0, v, v_0 for the kernel coordinates of x, w, x', w' respectively, and put $U = \varphi(u)$, $V = \varphi(v)$. The cocycles

$$a_g = p^{e_1}\alpha_c(g), \qquad b_g = p^{e_2}\beta_{c'}(g)$$

give the first coordinates of the extension actions. In particular set

$$A_g = gu - u + a_g, qquad B_g = gv - v + b_g,$$

which are the first coordinates of $gx - x$ and $gx' - x'$. They are themselves cocycles. Define the bilinear scalar-valued pairing $\langle r, s\rangle = p^{e_3}B_0(r, s)$ extended to the Witt modules; it is Galois equivariant and Frobenius compatible. The rest of the proof uses only these identities and the fact that the target is $\mathcal{H}$.

By Lemma 25.18, choose $y \in \mathcal{H}$ satisfying

$$(\varphi - 1)y = \langle u_0, V - v_0\rangle + \langle u, v_0 - v\rangle. \tag{25.10}$$

Define a function (for now only algebraically) by

$$f(g) = \langle a_g, gv\rangle - \langle u, B_g\rangle + gy - y. \tag{25.11}$$

This is exactly the source's cochain after simplifying the four terms using equivariance and the extension action formula. We verify its two properties.

First, since φ commutes with g and is trivial on the integral cocycles, the difference of applying Frobenius to (25.11) is

$$\begin{aligned}(\varphi - 1)f(g) = {}& \langle gu_0 - u_0 + a_g,\ g(V - v_0)\rangle \\ & - \langle U - u_0,\ gV - V + b_g\rangle \\ & - \langle u_0 - u,\ gv_0 - v_0 + b_g\rangle \\ & + \langle gu - u + a_g,\ g(v_0 - v)\rangle.\end{aligned}$$

For completeness, insert (25.10) in $(g - 1)(\varphi - 1)y$. The remaining terms in $(\varphi - 1)f(g)$ are

$$\langle a_g, g(V - v)\rangle - \langle U, gV - V + b_g\rangle + \langle u, gv - v + b_g\rangle.$$

Adding the two terms of $(g - 1)(\varphi - 1)y$ and collecting by bilinearity yields the displayed four-move identity. Each summand is small by the assumed safety of the corresponding

move, so $(\varphi-1)f(g)$ is small. Lemma 25.18 gives a unique constant $k(g)\in\mathbb{Z}/p^m(1)$ with $f(g)-k(g)\in\mathcal{K}_m(1)$.

Second, use the inhomogeneous coboundary convention $\partial f(g,h)=gf(h)-f(gh)+f(g)$. The cocycle relations for a_g and B_g give directly

$$\partial\langle a_g, gv\rangle(g,h)=-\langle a_g, g(hv-v)\rangle,$$
$$\partial(-\langle u, B_g\rangle)(g,h)=-\langle gu-u, gB_h\rangle.$$

The last term of (25.11) is a coboundary and has zero differential. Using $hv-v=B_h-b_h$, we obtain the exact identity

$$\partial f(g,h)=\langle a_g, gb_h\rangle-\langle A_g, gB_h\rangle.$$

The last pairing is small by the assumption on moves by $hx'-x'$ while x rests. Replacing f by k changes ∂f by a small element since that subgroup is G'-invariant. Both ∂k and $\langle a_g, gb_h\rangle$ are constants, so their small difference is zero by uniqueness of constants modulo the small subgroup. Thus on G'

$$\partial k(g,h)=p^{e_1+e_2+e_3}B_0(\alpha_c(g), g\beta_{c'}(h)).$$

We still need continuity of k. Near the identity, a_g and b_g vanish modulo p^m by continuity of the integral cocycles. Also $gv-v$ and $gy-y$ are small, by the preceding near-identity argument. The fixed pairing with u of the difference $gv-v$ is small near the identity even if u has nonintegral coordinates: finitely many coordinates of a continuous bilinear expression vanish at the identity. Consequently (25.11) is small near the identity and $k(g)=0$ there. For h near 1, $\beta_{c'}(h)$ vanishes modulo p^m, so the equality for ∂k gives $k(gh)=k(g)$ for all g. Hence k is locally constant. It supplies a continuous coboundary on G'.

Finally the open subgroup G' has finite index prime to p. Restriction followed by corestriction is multiplication by its index on H^2. The class under consideration is p^m-torsion, so vanishing after restriction implies vanishing on the full local group: the index is invertible on any p^m-torsion group. Further multiplication by p proves the assertion for all N in the stated range. □

The original formal conclusion is also expressed using the embedding of the finite Tate module as roots of unity into the multiplicative group of the algebraic closure; the finite-module coboundary just obtained gives that conclusion as well. This is the content of `IsScaledCupCoboundary.of_isSafePath`, `phi_smul_sub_mem_smallPart`, `nsmul_modPairing_eq_coboundary`, and `isCupCoboundary_of_eq_coboundary`. The proof shows why two different kinds of safe movement are needed: the four Frobenius moves ensure that the adjusted cochain takes constant values up to small elements, whereas the cocycle motions ensure that its coboundary is the desired cup. Merely showing that the pairing of two Hodge filtrations is small would not directly establish either assertion.

25.11 The walk along the Frobenius permutation

We give the exact order systems for terminal and intermediate points. A cut is a distinguished letter. Let l_i be the least strictly positive integer such that $i[-l_i]$ is a cut, and require that each Frobenius orbit contain a cut. A given cut has an integral shift $t_i\ge 0$ such that the third side has cokernel killer $p^{\kappa}E^{t_i}$ at it. Fix a margin $\eta\ge 1$ and define

$$D(t,\eta,l;a,u)=-u+t\mathbf{1}_{a=0}+\eta\mathbf{1}_{a<l},$$
$$S(\eta,l;a,u)=\begin{cases}\max(0,-u-\eta) & a<l,\\ -u & a\ge l.\end{cases}$$

A standard point at a cut i has orders $\rho_{-\eta}$ and depth $D(t_i, \eta, l_i)$; a stop has orders and depth at least $S(\eta, l_i)$. These are the definitions `standardDepth` and `stopOrders`, and the latter is deliberately the maximum of two order systems. A transported point $\varphi^l x$ has orders shifted by l (zero at the new places $a < l$) and depths shifted by l with lower bound 1 instead of 0 at those new places. This uses the fact that φ^l of a crystalline frame has positive depth at the new divisors; see `phi_pow_smul_mem_pointsOver` in `SidePoints.lean`. We record this integral Frobenius transport property as an input to the walk, exactly in these strengths; translation of the old Hodge orders alone would give only order zero and would not guarantee the positive depth at the new divisors.

At a cut i, put $l = l_i$, $j = i[-l]$ and $l_0 = l_j$. Suppose standard points on the two sides over $p^{e_1}c_i, p^{e_2}c_i'$ have been chosen at all cuts. Write y, y' for the transported points from j to i, and x, x' for the new standard points at i. If a stop q exists on the first side allowing the second to move directly from y' to x', then the three nontrivial moves

$$(y, y') \longrightarrow (q, y') \longrightarrow (q, x') \longrightarrow (x, x')$$

are all safe. We show this by checking the first and last move numerically. The middle move is the stated stop hypothesis. The first movement $y - q$ has order at least the minimum of the shifted $\rho_{-\eta}$ and $S(\eta, l)$. Below l this is zero. The depth of y' there is at least one. Thus the sum is at least 1, meeting the rule from the killer at j. At $a = l$ the moving order is $-u - \eta$ and the resting depth is $-u' + t_j + \eta$, since $l_0 > 0$. Their sum is $t_j - u - u'$, precisely the rule from that killer. At $a > l$ the moving order is $-u$ and the resting depth is at least $-u'$. This meets the remaining part of (25.9). For the final movement $q - x$, the order is at least $-u - \eta$ at $a = 0$, and the resting depth of x' is $-u' + t_i + \eta$ there. Their sum is $t_i - u - u'$, meeting the rule for the killer at i. At $a > 0$, the moving order is at least $-u$ except possibly when $0 < a < l$, where it is at least $-u - \eta$. The depth $D(t_i, \eta, l; a, u')$ is at least $-u' + \eta$ at those places and at least $-u'$ otherwise. Hence the remaining rules hold. If the stop is on the second side instead, switch which of the two intermediate moves is zero; the same inequalities apply. This proves `exists_isSafePath_of_stop`.

There is one additional local check for use in Theorem 25.19. If x, x' are standard at a cut and have been transported by φ^d to another letter, then the move of the second side by its cocycle $g\varphi^d x' - \varphi^d x'$ while the first rests is safe. Each of the moving order and the resting depth is bounded below by the shifted system $D(t_i, \eta, l_i)$ with value 1 at new places $a < d$. Their sum is at least 2 there. At $a = d$ it is at least $-u - u' + 2t_i + 2\eta \geq t_i - u - u'$. Above d it is at least $-u - u'$. Apply the rule from the third side at the originating cut. This proves `isSafeMoveSnd_cocycle`.

Theorem 25.20 (The global Frobenius walk)**.** *Suppose the standard points and third-side killers at each cut exist as above, and at each cut one side has a stop allowing the other side to move directly from its previously transported point to its new standard point. Then*

$$p^{e_1+e_2+\kappa}[\alpha_c \smile_{B_0} \beta_{c'}] = 0.$$

Proof. For each letter i, put $d_i = 0$ if it is a cut and $d_i = l_i$ otherwise. Choose the left point x_i at i to be the transport by φ^{d_i} of the standard point at $i[-d_i]$; define x_i' similarly. Each has the specified projection to $p^{e_1}c_i$ or $p^{e_2}c_i'$. Define y_i and y_i' by the strictly preceding cut (transport by φ^{l_i}). Because

$$l_i = 1 + d_{i[-1]},$$

one has $\varphi(\sum_i x_i) = \sum_i y_i$ and likewise on the other side: the equality follows term by term upon changing the letter index by inverse Frobenius. Off the cuts these transported points already equal the chosen points, so four moves of size zero are safe. At the cuts the previous

three-move calculation gives a four-move path with one zero move. Thus at each letter there is a path from y_i, y_i' to x_i, x_i' with intermediate points w_i, w_i' lying over the respective constants.

The extensions carry the complementary letter decompositions for the bilinear pairing: writing the second using the dual convention, a pairing between different indexed letter factors is zero. For clarity this follows from the balanced coefficient action. If the two idempotents are e_i, e_j, their cross pairing factors through $e_i e_j$, which is zero for $i \neq j$. Equivalently their distinct Teichmüller eigenvalues have unit difference. The Galois elements of G' preserve these letter factors, so the safety conditions for the sums of movements follow by expanding bilinearly: only pairings at the same letter remain and their sum is small. In particular the cocycle movements for the sum of the x_i's are safe by the preceding local check. The projections of the sums of the chosen and intermediate points are precisely the global constants $p^{e_1}c, p^{e_2}c'$. Apply Theorem 25.19 with these four global points. □

This is the proof of `isScaledCupCoboundary_of_stops`, including the minimal-cut condition: it is not sufficient that some cut occurs near each letter; the shift relation for the sums uses the nearest strictly preceding cut. When every letter is a cut, $l_i = 1$ and the same proof reduces to the simpler `isScaledCupCoboundary_of_stops_at_all_letters`.

25.12 Plain cuts and long-gap cuts

The cut theorem is useful only when both types of stops can actually be constructed with the same budget. We prove the two local constructions.

Call a cut with predecessor $l > 0$ steps back plain if among the l letters ending at that cut one side has all jumps $u \leq \eta$ and the opposite side has all jumps $u' \geq -\eta$. We call the former the down side. Suppose the down point x at the cut has not merely the standard depth but the depth

$$D(t_i, \eta, l; a, u) + \eta \mathbf{1}_{a=l}.$$

This is the source's `plainDepth` with equal margins at the two ends (the proof works with a different margin η_0 at the earlier cut, replacing the last η by η_0). On the down side, at $a < l$ the value has depth at least $\eta - u \geq 0$. Its original order is at least $-u - \eta$ and, if that order is positive, it already suffices for $S(\eta, l)$. Otherwise Lemma 25.12 improves the order to zero when the depth is nonnegative. For $a \geq l$ the original order $-u$ already suffices. Consequently x is a stop. Its depth exceeds the stop orders as well.

The up side's direct movement from the transported point at the earlier cut to the current point has order at least the minimum of shifted $\rho_{-\eta}$ and $\rho_{-\eta}$. At $a = 0$ the order is at least $\min(0, -u' - \eta) = -u' - \eta$, since the side is up, and the resting depth at x is $-u + t_i + \eta$. At $0 < a < l$, the moving order is at least $\min(0, -u') \geq -u' - \eta$ and the resting depth is $-u + \eta$. At $a = l$ the moving order is at least $-u' - \eta$ and the resting depth has the added η at that place. Above l the two bounds are $-u', -u$. Thus the rule for the third-side killer at i applies everywhere, proving that the up side moves safely while x rests. This is `mem_stopPoints_of_down` and `exists_mem_stopPoints_isSafeMoveSnd`, together with the version exchanging the first and second sides.

Now fix windows

$$U_i(T) \subset [-\eta, \eta] \cup \epsilon_i[\beta_i, \beta_i + V], \qquad U_i(T') \subset [-\eta, \eta] \cup (-\epsilon_i)[\beta_i, \beta_i + V]$$

with $\epsilon_i \in \{\pm 1\}$ and $\beta_i, V \geq 0$. An isolated far window with an opposing sign may prevent the cut from being plain. The long-gap construction is designed for this case. Suppose the chart

is unramified with $p \geq 5$, the preceding cut $j = i[-l]$ has shift zero, and the pivot letter is $k = i[-t]$ with

$$t + J + 1 = l, \qquad s_0 + \omega(1 + p + \cdots + p^J) < \beta_k.$$

Here $\omega \geq \eta$ bounds all the far window endpoints $\beta_{i[-a]} + V$ for $t < a < l$. Every nonpivot letter $i[-a]$, $a < l$, either has small far window ending at or below η or has sign opposite to the pivot. By choosing which side will rest, arrange that its far window at k is negative. At the other letters before i its jumps are at least $-\eta$. Let it be the side T. All its jumps at k and the J previous letters are at most ω, because the far window at k is negative. If p^κ kills its cokernel at j, E^{s_0} kills it at k, and

$$A = 2b \leq p^{J+1}\delta, qquad\beta_k < 2b,$$

then Theorem 25.16 supplies a point at k over $p^{\kappa+4\delta}c_k$ with order zero below level $J + 1$, except for far negative jumps at level zero where its order is $2b$. Its depth is $2b$ at level zero, zero from level one to level J, and the standard orders above. Transport by φ^t to i. It now has the indicated large orders and depth at level t; below t its order is zero and depth at least one. Use the depth operators for the steps $(a, 0)$, $t < a < l$, to deepen from zero to one there, preserving all existing orders and depths. This costs only the sum of the corresponding $3a$ (the chart is unramified).

The result is a point with depth

$$D_{\mathrm{gap}}(l, t, 2b; a, u) = \begin{cases} 2b & a = t, \\ 1 & a < l,\ a \neq t, \\ -u & a \geq l. \end{cases}$$

At levels below l other than the pivot the side is up, so $u \geq -\eta$, and order zero meets the stop-order requirement. At the pivot either the jump is low or is far negative with $-u \leq \beta_k + V \leq 2b$; the large orders obtained by Taylor extraction meet $\max(0, -u - \eta)$. Above l the standard order $-u$ suffices. The depth is at least the stop orders at all these places.

It remains to check safety during the other side's movement. The actual cut-order system used to construct standard points is slightly stronger than $\rho_{-\eta}$. It is

$$O(s, \eta, l) = \mathrm{imp}\left(\rho_{-s}, D(s, \eta, l + 1)\right). \tag{25.12}$$

At the earlier cut the shift is zero. At a nonpivot level $a < l$, the other side's jump u is at most η. Its cut orders at i are nonnegative there: for $a > 0$ this follows from $-u+\eta \geq 0$ in the depth, and for $a = 0$ from $-u + s + \eta \geq 0$. The transported cut orders are also zero below l. Their minimum is therefore nonnegative, and the resting depth 1 gives the rule for the third side at j. At $a = t$, the far jump of the other side is positive, bounded by $2b$, while its cut orders are at least $\min(0, -u + \eta)$ (also if the pivot lies at $a = 0$, since the shift $s \geq 0$ only increases the depth). Thus the moving order plus $2b$ is at least 1 there, using $\eta \geq 1$. Above and at l the moving order is at least $-u$ (the earlier cut has shift zero), and the resting depth is $-u'$. This proves the remaining rules with $t_j = 0$. Hence the direct movement is safe, and the long-gap point is the required stop. This is the argument of `pointsOver_longGapDepth_nonempty` and `ruleBound_le_inf_add_longGapDepth`.

Notice the role of the positive margin η. The strict 1 needed below the cut predecessor does not follow from mere complementary sign windows; it is supplied either by the depth at nonpivot places or by the positive margin at the pivot. Also, one must keep the stronger orders (25.12) for the long-gap case. Replacing them prematurely by the standard order lower bound $\rho_{-\eta}$ would not establish the nonnegative moving order at the other places.

25.13 Quantitative killers at sparse cut positions

We return briefly to the ideals of cokernel multipliers. Write $J_l(i;r)$ for the multiplier ideal at level l and letter i modulo r, using the notation of Section 24.7. The exact definition and the coefficient retraction imply the Nakayama property: if $p^\kappa \in J_0(i;p^{\kappa+1})$, then p^κ is an exact killer at i. Indeed p^κ lies in the exact multiplier ideal after multiplication by an element of the form $1-pu$ (move the error term divisible by $p^{\kappa+1}$ to the left side); that element is a unit by p-adic completeness. We use the exact one-step scalar calculus of Proposition 23.6, with the moving operator detailed above. Two consequences will be useful. If $p^c z^a \in J_0(j;r)$ and all jumps at j are at most n, then the multiplication by E^n to pass to level one, the depth-step trade of E for z, and the Frobenius return give

$$p^{c+d_0 n+v_p(n!)} z^{\lfloor (a+n)/p \rfloor} \in J_0(j[-1];r). \tag{25.13}$$

Also, with $p^c z^a \in J_0(i;r)$, removing the power of z gives

$$p^{c+a+v_p(a!)} \in J_0(i;r). \tag{25.14}$$

For (25.13) the factorial valuation counts the successive costs of trading E^n for z at level zero (each nonzero step costs $d_0+v_p(k)$), and decreasing a z-exponent to the nearest multiple of p in level one costs zero. For (25.14), first carry forward to level one and power z^{pa}, then lower one z-exponent at a time. The cost is $v_p((pa)!) = a + v_p(a!)$; return by Frobenius. These are `p_pow_mul_z_pow_div_mem_liftIdeal` and `p_pow_mul_z_pow_mem_liftIdeal_of_le` in the readable layout. The preservation of the exact multiplier ideals by the indicated scalar derivations still uses the compatibility of retraction with integral solutions (`mem_liftIdeal_of_eq_mul_smul_sub`); one cannot replace those ideals by arbitrary pointwise conditions on a quotient of the de Rham vector space. Under this ideal compatibility, these two formulas and their arithmetic consequences are now proved.

Suppose all jumps are at most w at all letters and we work modulo $r = p^{\kappa+1}$. The least $a \ge 0$ such that $z^a \in J_0(i;r)$ for all letters satisfies

$$(p-1)a \le e_0(w+\kappa). \tag{25.15}$$

This slightly sharpens Proposition 24.5 at exactly this modulus. The existence of some a follows in the same way, using the height of the original lattice modulo $p^{\kappa+1}$. To get the sharper inequality, go up to level one by multiplication by E^w. The element $z^{e_0(w+\kappa)}$ lies in $(p,E)^{w+\kappa}$, because z^{e_0} belongs to (p,E) in S. Multiplying by z^a, each monomial in the expansion has either an E-exponent at least w or a p-exponent at least $\kappa+1$: if the latter is at most κ, the remaining part is at least w. Therefore $z^{a+e_0(w+\kappa)}$ lies in the level-one ideal. Lower its exponent to the nearest multiple of p at zero cost and return by Frobenius. Minimality gives $a \le (a+e_0(w+\kappa))/p$, which proves (25.15). The small improvement here matters when inserting a minimal exponent into an inequality in which the chosen modulus itself depends on the cost.

Proposition 25.21 (An exact killer from J thin successors)**.** *Suppose all jumps are bounded above by w and at the J successive letters after i by g. Assume $8e_0 \le p^{J+1}$ and*

$$2Jd_0 g + (2J+4)\frac{g}{p-1} + \frac{8e_0 w}{p^{J+1}} \le \kappa. \tag{25.16}$$

Then p^κ kills the cokernel at i.

Proof. Work modulo $r = p^{\kappa+1}$ and begin at exponent a from (25.15). Apply (25.13) repeatedly along the J thin letters. One obtains $p^c z^n \in J_0(i; r)$, with

$$c \le J(d_0 g + v_p(g!)), \qquad n \le \frac{a}{p^J} + \frac{g}{p-1}.$$

The second inequality follows by starting with a and iterating the map $x \mapsto \lfloor (x+g)/p \rfloor$; the geometric sum of g/p is bounded by $g/(p-1)$. Legendre's estimate gives $v_p(g!) \le g/(p-1)$ and $v_p(n!) \le n$. From (25.15), since $p \le 2(p-1)$ and $8e_0 \le p^{J+1}$,

$$\frac{a}{p^J} \le \frac{2e_0(w+\kappa)}{p^{J+1}} \le \frac{2e_0 w}{p^{J+1}} + \frac{\kappa}{4}.$$

Thus the total cost of (25.14) is bounded by

$$\begin{aligned} c + n + v_p(n!) &\le J\left(d_0 g + \frac{g}{p-1}\right) + 2\left(\frac{2e_0 w}{p^{J+1}} + \frac{\kappa}{4} + \frac{g}{p-1}\right) \\ &\le \kappa, \end{aligned}$$

where the last inequality is exactly (25.16) divided by two and added to $\kappa/2$. So $p^\kappa \in J_0(i; p^{\kappa+1})$. Apply the Nakayama property. □

This is `killsCokernel_zero_of_jumps_le`. It explains the bound in terms of p^{J+1} rather than in terms of the full rank or the number of letters. Large jumps far from the cut are admitted into the initial exponent but decreased by a geometric factor during Frobenius return.

We now combine the previous results in the numerical cut scheme of readable `CutExponentBounds.lean`. We state the parameters with slack, so that the construction can be used in both of its principal specializations.

Theorem 25.22 (Cut-window bound)**.** *Suppose the two sides have jumps in the opposite windows of Section 25.12; let $s_i \ge 0$ be the shifts at letters. Fix nonnegative parameters $W, X, \delta, \kappa, b, n, J, s_0, G$ satisfying*

$$\begin{gathered} \eta \ge 1, \quad V + 4\eta \le W, \qquad n + \lfloor \log_p(2b) \rfloor + (J+1+d_0)W + d_0 G \le X, \\ 12((J+1)X + \delta) < \kappa, \qquad 8e_0 \le p^{J+1}, \qquad 2e_0 b \le p^{J+1}\delta, \\ \beta_i + s_i + W \le 2b, \quad s_i + W \le pn, quad G \le (p-1)n, \quad \beta_i + W \le (p-3)s_0 \quad \text{if } p \ge 5. \end{gathered}$$

Assume every letter has a strictly preceding cut within $2J+1$ steps, and every cut i is one of the following:

(i) $s_i = 0$, and each of the J successive letters j satisfies $\beta_j + W \le G$;

(ii) the chart is unramified, $p \ge 5$, $s_i < \beta_i$, and $s_0 \le \kappa + s_i$.

Finally, at each cut, the interval separating it from the nearest preceding cut is plain or is a long gap in the sense of Section 25.12, with s_0 as above and some ω bounding the pre-pivot windows. Then $p^{51\kappa}$ annihilates every paired cup of crystalline classes of the two sides.

Proof. First all three sides have jumps at most $\beta_i + W$. For the individual sides this is immediate from their windows. For the Tate dual, if the two jumps are both far they have opposite signs and $|-u-u'-1| \le V+1$; if both are low the bound is $2\eta+1$; if only one is far the bound is $\beta_i + V + \eta + 1$. All are at most $\beta_i + W$ since $W \ge V + 4\eta$, $\eta \ge 1$. This also covers the jump zero of an extension. At a cut of type (i), Proposition 25.21 applies with

$w = 2b$, $g = G$. Indeed the cost in (25.16) is bounded by $(4J + 4)X + 8\delta < \kappa$. For type (ii), use the exact unramified bound of Proposition 24.6, with total budget $\kappa + s_i \geq s_0$. Thus $p^\kappa E^{s_i}$ kills the cokernel of each of the three sides at every cut.

Fix a cut i and set $l = l_i \leq 2J + 1$. For either side, initial points over $p^\kappa c_i$ exist with order and depth ρ_{-s_i}. Increase the depth to the system $D(s_i, \eta, l + 1)$. The needed steps occur only at $0 \leq a < 2J + 2$ and lie in the two windows of radius

$$r_a + W, qquadr_a = s_i \mathbf{1}_{a=0}$$

centered at zero and the far negative of the sign of the jump window at $i[-a]$. Indeed at level zero the steps range from $-u - s_i$ to $-u + s_i + \eta - 1$ and at the other levels from $-u$ to $-u + \eta - 1$. Since $V + 4\eta \leq W$, the radius is sufficient. Here $r_a + W \leq pn$ and the absolute values of all integers in these windows are at most $2b$, by the stated shifted far bound. At level zero either $r_a = 0$ or $d_0 = 0$ by the cut types, hence always $r_a = 0$ or $d_0 + 3a = 0$. The estimate of Section 25.7 therefore bounds the cost at each level by

$$8n + 2\lfloor \log_p(2b) \rfloor + 5W \cdot 6(J + 1 + d_0) \leq 30X.$$

The cardinal bound uses $4W + 2 \leq 5W$ for $W \geq 4$, and $d_0 + 3a \leq 6(J + 1 + d_0)$. Thus the full sum of costs is at most $60(J + 1)X$. The chosen κ gives

$$\kappa + 4\delta + 60(J + 1)X \leq 25\kappa.$$

This estimate is generous enough for both types of stop. Apply Proposition 25.14 and Lemma 25.12, then multiply by the unused power of p to obtain a point over $p^{25\kappa} c_i$ with strong orders (25.12) and the required depth.

These strong orders imply the standard orders $\rho_{-\eta}$. For $a > 0$, if $-u > 0$ then improved order is at least $-u$; otherwise $D(s_i, \eta, l + 1; a, u) \geq -u$ and $\min(0, D) \geq -u$. At $a = 0$, the same is true with $-u - s_i > 0$ or with $-u - \eta \leq 0$. The remaining possibility $-u - s_i \leq 0 < -u - \eta$ would put a negative jump in $[-s_i, -\eta)$, impossible when $s_i = 0$ or when $s_i < \beta_i$ and the far negative values satisfy $-u \geq \beta_i$. Thus the implication holds. The depth $D(s_i, \eta, l + 1)$ is also the depth required at a plain cut in Section 25.12.

If the cut is plain, take the standard point of the down side as stop; its depth makes the movement safe. At a long gap, the preceding cut has shift zero. The unramified bound with shift s_0 at the pivot kills its cokernel with no p-cost, since all jumps are at most $(p - 3)s_0$. At the preceding cut use the already established p^κ-killer. The radius condition and long-gap hypothesis are exactly those used in Section 25.12; after the Taylor extraction and transport, only the depth steps $(a, 0)$ with $t < a < l$ are needed. These are among the two-window steps, so the full cost, including $\kappa + 4\delta$, is at most 25κ. Thus the required stop exists over the same constant as the standard point, and the move at this stop is safe by the strong order calculation in that section.

All conditions of Theorem 25.20 now hold with $e_1 = e_2 = 25\kappa$, and the third-side cost is κ. The total is 51κ. □

25.14 Two choices of cuts

We explain the combinatorial choice in the sparse-thick-letter case and the two-size case, which have quite different hypotheses. Let T be the set of thick letters (the letter i is thick if $\beta_i = b$ and thin if $\beta_i = g$). Use the direction of Frobenius to say that a letter is followed by another a positive number of steps later.

Sparse thick letters. Assume no three thick letters follow one another with the two successive gaps both between 1 and J, with $J \geq 1$. Mark a cut when the following J letters are all thin. We first prove two combinatorial facts. If a thick letter j is not a cut, then some thick letter follows it within J steps. No thick letter can precede j within J steps, by the no-three hypothesis. Therefore the letter $J+1$ steps before j is a cut. Next, for any letter i, either $i[-(J+1)]$ is a cut or some thick letter $i[-r]$ with $1 \leq r \leq J$ follows it before i. In this second case either $i[-r]$ itself is a cut or $i[-(r+J+1)]$ is one, by the previous fact. Thus every letter has a preceding cut within $2J+1$.

Consider the nearest preceding cut $i[-l]$ of i and a thick letter $i[-r]$ with $0 < r < l$. Since there are J thin letters after the cut, we have $l - r \geq J+1$. The thick letter $i[-r]$ is not a cut (by nearestness), hence $i[-(r+J+1)]$ is a cut, implying $r+J+1 \geq l$. Therefore $r+J+1 = l$. In particular there is at most one interior thick letter, necessarily $J+1$ steps after the cut. These are the combinatorial assertions `hasThinSuccessors_stepBack_add`, `exists_hasThinSuccessors_stepBack`, and `eq_zero_or_add_eq_of_hasThinSuccessors`.

Suppose now the thin windows lie in the low region, i.e. $g + V \leq \eta$. Assume either all signs are equal, or the chart is unramified with $p \geq 7$ and $10p^J\eta \leq b$. In the constant-sign case all cut intervals are plain. Otherwise if an interval is not plain, the only interior thick letter is a pivot of sign opposite to the sign at i. Every other letter than the pivot and i is thin. Put $W = 5\eta$ (here $V \leq \eta$), use $b_* = b + W$ in place of b in Theorem 25.22, let $s_i = 0$ for all i, $n = G = 2W$, and

$$s_0 = \left\lfloor \frac{b + 10\eta}{p-3} \right\rfloor + 1.$$

When the sign is not constant, the pivot satisfies the long-gap inequality with $\omega = \eta$:

$$s_0 + \eta(1 + p + \cdots + p^J) \leq \frac{b+10\eta}{4} + 1 + 2p^J\eta < b.$$

Indeed $p \geq 7$, $p^J \geq 7$ and $b \geq 10p^J\eta$ imply the strict final inequality (the right-hand side of the first inequality is at most $b/4 + b/28 + 1 + b/5 < b$). All the pre-pivot windows are thin. Thus exactly one long gap or a plain interval occurs, and all cuts are of type (i).

Take

$$8e_0 \leq p^{J+1}, \quad 2e_0(b+5\eta) \leq p^{J+1}\delta, \quad 20(J+1+d_0)\eta + \lfloor \log_p(2(b+5\eta)) \rfloor \leq X,$$

where the first inequality is automatic for the unramified $p \geq 7$ case. The cut-size bounds hold: the cost requires $2W + (J+1+d_0)W + 2d_0W + \log$ which is at most $4(J+1+d_0)W + \log$; the shifted windows end at or below $2b_*$, since $\beta_i \leq \max(b, W)$. The bound for $p \geq 5$ follows from the definition of s_0. Choose $\kappa = 12((J+1)X + \delta) + 1$. The theorem gives an annihilation exponent 51κ, bounded by

$$10^3(\delta + (J+1)X)$$

since the latter parenthesis is a positive integer. This is the readable version `isScaledCupCoboundary_sparse_thick`. The compact source uses differently packaged constants and a directly expressed depth budget as in Chapter 23. The exact comparison of the two sets of numerical hypotheses is separate; the cut mechanism is common to both.

Two sizes in arbitrary order. Now assume the chart unramified, $p \geq 7$, $0 < g \leq b$, and suppose

$$b < Kp^2, \qquad 2g < Kp, \qquad W \leq K, qquad V + 4\eta \leq W.$$

Use $J = 1$ and put $q = \lfloor b/p \rfloor$. Set

$$s_i = \begin{cases} \max(0, q - 7K) & \text{the successor of } i \text{ is thick,} \\ 0 & \text{otherwise.} \end{cases}$$

The cuts are the letters for which $s_i < \beta_i$. Every thick letter is a cut ($s_i \le b/p < b$), and every letter with a thin successor is a cut ($s_i = 0 < g$). Thus one of the first two preceding letters of every i is a cut: if $i[-1]$ is thick it is itself a cut, and if thin then $i[-2]$ has a thin successor. If the nearest preceding cut is $i[-1]$, the interval is plain. If it is $i[-2]$, then $i[-1]$ is thin and not a cut, so i itself is thick; moreover $s_{i[-1]} \ge g$, giving $g + 7K \le q$, and $s_{i[-2]} = 0$. If the two signs agree the interval is plain. If not, the pivot is i, $t = 0$, $J = 1$, and with $s_0 = q + 12K$, $\omega = g + W$, we have the long gap

$$s_0 + \omega(1+p) \le q + 12K + (q - 6K)(1+p) = (p+2)q - 6(p-1)K < b.$$

The last inequality uses $pq \le b$, $q < Kp$ and $2Kp < 6(p-1)K$. The other long-gap conditions were just verified.

To verify the cost scheme take

$$b_* = b + W, \quad n = K, \quad G = g + W, \quad X = 3K + \lfloor \log_p(2b_*) \rfloor, \quad \delta = 4K, \quad \kappa = 12(2X + \delta) + 1.$$

Since $d_0 = 0$, the cost bound is immediate from $W \le K$. The inequality $8e_0 \le p^{J+1}$ holds since $e_0 = 1$, $p \ge 7$. Also $2b_* \le 4Kp^2$, and $s_i + W \le pK$ since $q < Kp$. We have $G \le Kp/2 + K \le (p-1)K$ for $p \ge 4$. The shifted far bound $\beta_i + s_i + W \le 2b_*$ is obvious from $s_i \le b$. For the final far bound at $p \ge 5$, write $b = pq + r$, $0 \le r < p$; then

$$b + W - (p-3)(q + 12K) = 3q + r + W - 12K(p-3) < 0.$$

Indeed $3q + r + W < 3Kp + p + K \le 4Kp + K$ since $K \ge 1$, whereas $12(p-3)K \ge 48K$ for $p \ge 7$ and in fact $12(p-3) > 4p + 1$ for such primes. Finally, any cut with thin successor is of type (i), and a cut with thick successor satisfies $s_0 \le s_i + 19K \le s_i + \kappa$ and $s_i < \beta_i$, hence of type (ii). Theorem 25.22 gives

$$51\kappa = 51(12(10K + 2\lfloor \log_p(2b_*) \rfloor) + 1) \le 10^4(K + \lfloor \log_p(2b_*) \rfloor).$$

This is the bound `isScaledCupCoboundary_any_pattern`. In the compact source, the counterpart is `scaledCupBounds_two_sizes`. Its role is precisely to handle a succession of many thick letters without requiring the sparsity used in the previous case: a shift at a cut replaces the unshifted bound from the next J thin letters.

25.15 A multiplicity-sensitive step in the few-low-jumps case

As seen in Section 24.8, the few-low-jumps length estimate, unlike the preceding rank-free bounds, must take account of the multiplicities of the low jumps. We now reconstruct the key correction of relations on the actual constants. In the unramified chart, write $Q(o)$ for the order module on the kernel lattice Y at a fixed letter i, intersected with its localized Kisin module. Its image in the extension is the submodule of vectors over zero with orders o. We use the comparison assertion that the latter submodule identifies with the order module of the kernel on the specified common jump sets, precisely `exists_inclKer_eq_of_kerWithOrders` in readable `SidePoints.lean`; this identification is an integral functoriality assertion, not automatic from splitting an extension as underlying $\mathbb{Z}_p$-modules.

Lemma 25.23 (Primitive relation up to E)**.** *Let m be the Hodge–Tate multiplicity of jump u at $i[-a]$, and suppose $o(a,v) > o(a,u)$ for all smaller jumps $v < u$ there. Put o' equal to o with $o(a,u)$ raised by one. For any $f_0 m + 1$ vectors y_j in $Q(o)$, where $f_0 = [\Lambda : \mathbb{Z}_p]$, a primitive $\mathbb{Z}_p$-linear combination of the y_j belongs to*

$$EQ(o) + Q(o').$$

Proof. Choose the forms generating the degree-u part of the filtered dual space modulo degree $u-1$, at most m forms by the definition of multiplicity. For each y_j evaluate those forms at level a, divide by the common $t^{o(a,u)}$, and take the residue in C_p of the resulting integral values. This gives row vectors in a vector space of dimension at most m. The maximal-minor argument of Section 24.8 gives a subset $y_{j_1},\dots,y_{j_d}$ with $d\le m$ such that each residue row vector is an integral ($\mathcal{O}_{C_p}$-linear) combination of the selected ones. In fact choose a nonzero minor of maximum absolute value among all the minors of size equal to the rank of the rows; after choosing its rows as a basis, Cramer's rule bounds all expansion coefficients by one. If the rank is zero the residues already vanish. Use the surjective $\theta_a : A\to\mathcal{O}_{C_p}$ to lift the coefficients to A (equivalently lift at θ and apply φ^a). Denote the lifts by v_{jk}. Then

$$y_j-\sum_k v_{jk}y_{j_k}\in\mathcal{M}_i^Y(o').$$

Indeed the chosen leading residues cancel, and every form of degree u differs from a combination of the chosen forms by one of lower degree; the lower-degree forms already have order at least $o(a,u)+1$ by the strictness hypothesis. At the other levels and jumps the old orders are preserved by linearity over A. The decomposition of forms by degree here is the filtered basis argument for the Hodge–Tate jumps after the comparison, in the source `exists_span_formsOfType_mod_pred`.

Retract the displayed relation to the localized Kisin module. If $\lambda_0 : A\to S$ is its chosen scalar coordinate, then

$$y_j-\sum_k \lambda_0(v_{jk})y_{j_k}\in Q(o')$$

by the preservation of orders under retraction. Write $\lambda_0(v_{jk})=w_{jk}+Ez_{jk}$ with $w_{jk}\in\Lambda$ and $z_{jk}\in S$, by the unramified division used in the Taylor proof. The vectors $(w_{jk})_k\in\Lambda^d$ are $\mathbb{Z}_p$-linearly dependent since there are $f_0m+1>f_0d$ of them. Choose a relation and divide by a nonzero coefficient of least valuation to arrange one coefficient equal to 1. On taking the corresponding combination the w_{jk} terms cancel; the Ez_{jk} terms lie in $EQ(o)$, and the remainder in $Q(o')$. □

This supplies the relation `exists_sum_smul_mem_locKisinPart_sup` in readable `HodgeExpansion.lean`, including the preservation of the leading coefficients by projection. The regression to the actual constants still requires removing the $EQ(o)$ term; treating that term as zero merely because another filtration coordinate was raised would be incorrect.

For an extension let $X(o)$ denote the no-pole localized Kisin points over some constant and with depth o. Thus they have orders σ_0 and depth o; put $D(o)=qX(o)$ for the corresponding submodule of constants. Let $M(o)$ be the submodule of the extension consisting of the localized kernel vectors with orders o. When $o\ge\sigma_0$, $M(o)\subset X(o)$.

Lemma 25.24 (Division by p in the depth condition)**.** *Assume that no integer divisible by p lies strictly between $o(0,v)$ and $o(0,v)+p^a$ at any jump v. Let x have depth o over a constant. If $w\in EM(o)$ and $px-w$ has depth o' over a constant, where o' raises $o(a,u)$ by one, then there is $m_0\in M(o)$ for which $x-m_0$ has depth o'.*

Proof. Use the common zero-cost operator A_0 for the steps $(0,n)$ lying strictly between $o(0,v)$ and $o(0,v)+p^a$, over all jumps v. All the steps are free by $p\nmid n$ and $d_0=0$. This operator preserves all depths already satisfied and differs from the identity by a kernel correction of the input depth. Since $w\in EM(o)$, it has order o raised by one at level zero, and depth at least that. Consequently A_0w has the same order and has depth o raised by p^a at level zero. We claim its actual order is also at least that latter system. If $o(0,v)+1>0$, Lemma 25.12

gives this. Otherwise $o(0,v)+1 \le 0$ and there is no multiple of p strictly between $o(0,v)$ and $o(0,v)+p^a$; in particular zero is not in that open interval, so $o(0,v)+p^a \le 0$. The same improvement lemma again gives the required order. At other levels the original orders are retained. Hence

$$A_0 w = E^{p^a} w_1, \qquad w_1 \in M(o).$$

In the unramified chart $E^{p^a} - E_a = pr$ for some $r \in S$ (as in (24.13)). Set $m_0 = rw_1 - (A_0 x - x) \in M(o)$. Then

$$p(x - m_0) = A_0(px - w) + E_a w_1.$$

The first term has depth o' since the operator preserves that depth. The second has actual orders o raised by one at level a, which are at least o', hence it also has the required depth. Multiplication by p has no effect on the de Rham orders because p is a unit of B_{dR}^+. Therefore $x - m_0$ itself has depth o', with unchanged projection. □

This is `exists_sub_mem_deepElements_of_smul_sub_mem`. It is not a division inside the integral lattice by a nonunit p without control: the correction m_0 is explicitly integral and over zero.

Theorem 25.25 (Rank-sensitive colength of one step)**.** *Assume $o \ge \sigma_0$, the strict lower-jump hypothesis of Lemma 25.23, and the no-multiple hypothesis of Lemma 25.24. Then for o' raising $o(a,u)$ by one,*

$$\mathrm{length}_{\mathbb{Z}_p}(D(o)/D(o')) \le c(a, o(a,u)) f_0 m.$$

Proof. Consider $X(o)/(X(o') + M(o))$, which surjects onto $D(o)/D(o')$. By the operator of Theorem 25.13, this quotient is killed by p^k with $k = c(a, o(a,u))$: applying it to $x \in X(o)$ gives an element of $X(o')$ differing from $p^k x$ by a vector of $M(o)$. Preservation of the no-pole condition follows from $o \ge \sigma_0$.

Suppose we are given $f_0 m + 1$ elements of this quotient killed by p, represented by $x_j \in X(o)$. Write

$$px_j = x_j' + m_j, \qquad x_j' \in X(o'), \quad m_j \in M(o).$$

Apply Lemma 25.23 on the kernel to the m_j and use its identification with the extension kernel. A primitive combination $\xi = \sum_j b_j x_j$ therefore satisfies

$$p\xi - w \in X(o'), \qquad w \in EM(o).$$

Indeed the combination of the m_j equals $w + v$ with $v \in M(o') \subset X(o')$. Lemma 25.24 gives $m_0 \in M(o)$ with $\xi - m_0$ of depth o'. Since m_0 has no pole ($o \ge \sigma_0$), $\xi - m_0 \in X(o')$. Thus the primitive combination is zero in the considered quotient.

It follows that its p-torsion is an $\mathbb{F}_p$-vector space of dimension at most $f_0 m$. The dimension bound on p-torsion and the annihilator p^k are enough even without finite generation of $X(o)$: for a group killed by p^k, each successive p^j-torsion quotient injects into its p-torsion via multiplication by p^{j-1}. Hence its length is at most $k f_0 m$. The surjective image $D(o)/D(o')$ has no larger length. □

This proves `relLength_deepConstants_raiseOrder_le`. It supplies the rank-sensitive correction left open in Section 24.8. In that section the subsequent numerical sum of single steps yields the factors $(2s+1)n_i$ and $3(W+1)r_{i[-1]} + 6sn_{i[-1]}$, rather than the full rank times the largest gap. The additional initial-depth construction and the precise opposite-window safety inequalities in the last step of the few-low-jumps theorem still require verification: the present proof provides their central one-step colength estimate, but not a substitute for these final case checks.

25.16 Scope of the integral lifting reconstruction

We record exactly what is used from outside this chapter. The integral period-ring inputs are the structural closure, division and comparison properties listed in Sections 25.1, 25.2, 25.3, and 25.6. The Frobenius-basis and scalar inverse claims, the rank-free depth step (including zero), the long-gap extraction after the radius-span estimate, the continuous cup coboundary and the full cut-walk calculation are proofs from those stated hypotheses. The thin-cut annihilation and the two choices of cuts keep their source numerical constants in the readable layout. Where a result is invoked in the compact layout under differently packaged inequalities, the hypotheses must be compared explicitly rather than replacing its bound solely by a name from the other layout.

The special tame-type case uses a chart with a tame generator of inertia and a polynomial q annihilating its action on the crystalline invariant module. The following elementary step in that proof is useful to distinguish from its more specialized integral estimates. If a linear operator T on an extension $0 \to Y \to L \to C \to 0$ acts trivially on the quotient and $q(T)$ vanishes on the kernel, then $q(T)(T-1) = 0$ on L: the image of $T-1$ is contained in Y. At the tame chart the operator T is the linearized action of a chosen inertia element on periods; the same argument applies after the exact invariant-module comparison. Thus the relevant polynomial is $q(X)(X-1)$, of degree $\deg q + 1$. A primitive e_0-th root in the residue field gives e_0 distinct evaluation points; a monic polynomial of degree d has at most d roots among them. This counting is the mechanism that avoids a loss proportional to all e_0 residue classes: the exceptional evaluations are at most $\deg q + 1$. The integral passage from the polynomial relation to annihilation of the Kisin cokernel, however, also needs the construction of the tame chart, a near-identity divisibility estimate for its tame generator, and the scalar ideal recurrence at its exceptional evaluation points. These are the remaining tame inputs of `TameChart.killsCokernel_of_isKilledBy` and `Lift.exists_tameChart`, compact `New/Lift/SparseThickLetters.lean`. A bound for the order of a rational representation alone would not supply them.

The dependence on the Hodge–Tate frame is concentrated in the integral comparisons, not in the elementary stationary or moving arguments. To finish an entirely independent proof one must reconstruct the first half of the radius lemma and the chart comparisons as well as the special tame and few-low-jump obligations just described, or show that a standard theorem provides precisely these integral statements with their constants. Consequently we do not assert that the local crystalline block of the entire GRH argument has become independent of the formal development. The value of the reconstruction is that one no longer needs to import the nonstandard Frobenius walk, the long-gap extraction, or the abstract integral-basis argument as opaque local theorems.

Part V

The arithmetic conclusion and proof audit

Chapter 26

A numerical appendix: asymptotic budgets and the passage to the limit

The compact development checks inequalities on the logarithmic scale for every record satisfying its very large fixed admissibility bounds. A conventional paper need not insist on the same finite threshold for the parameters: to conclude GRH only the existence of admissible data for all sufficiently large x is needed. This chapter collects the estimates along that family, including the flexibility of the second cutoff. It gives ordinary proofs of several numerical comparisons implicit in the use of the formal budgets. No weakening of a non-numerical hypothesis on an automorphic object is intended. In particular a theorem whose proof depends on a precise integral divisor bound cannot be replaced by a numerical big-O estimate of the size of its coefficient matrices.

26.1 Uniform scale asymptotics

Fix N_0, d, and $0 < \varepsilon < 1/100$. We use the family of admissible records of Proposition 4.2 after replacing its second logarithmic cutoff by $u \in [10^6 \log x, (10^6 \log x)^2]$. Put $Y = \log b = x^{\varepsilon}$. In particular $u = o(Y^{\delta})$ for every fixed $\delta > 0$ (uniformly in the indicated window). Notice the identity $2^s = m^{22}$ and the estimates

$$\sqrt{x} \le m \le 2\sqrt{x}, \qquad x^5 \le r \le 3x^5, \qquad k = r + 1 - m, \qquad t = 5r,$$
$$81s < 3^{c_s} \le 243s, \qquad s \asymp \log x, \qquad T_s \asymp x^{11} \log x, \qquad f = \varphi(N_0)T_s.$$

Here $T_s = N_1/N_0$ and $c_s = \lfloor \log_3 s \rfloor + 5$. In particular $x^{11} \le T_s \le x^{11}Y^{1/5}$ eventually. The rounding $a = N_1\lceil xb/N_1 \rceil$ gives $a = xb + O(x^{11} \log x)$, and hence $a \sim xb$. All the estimates in this chapter are uniform in the second cutoff within its window. They also remain valid with the initial choice $\log x_1 = x^{\varepsilon/6}$ for the power-dominance estimates below; closeness of its logarithm to the window is not assumed until we choose the cutoff.

Lemma 26.1 (Power dominance)**.** *For every fixed $A, C, D \ge 0$ and $\eta > 0$, along the family just described,*

$$x^A Y^C x_1^D = o(b^{\eta}).$$

If $P(x)$ is bounded by a fixed power of x, then $P(x) = o(B_0)$, where $B_0 = B/(mk)$.

Proof. The logarithm of the quotient in the first assertion is at most

$$A \log x + C \log Y + D(10^6 \log x)^2 - \eta x^{\varepsilon} \longrightarrow -\infty.$$

If we retain the initial cutoff instead, the third term is $Dx^{\varepsilon/6} = o(x^{\varepsilon})$, so the same proof applies. From $a \ge xb$, $f \ge x^{11}$, $m \le 2\sqrt{x}$, and $k \le 3x^5$ we get $B_0 \ge bx^{13/2}Y^3/6$. Thus the second assertion follows from the first with $D = 0$ and any positive $\eta < 1$. □

This lemma must not be confused with a comparison between powers of x and powers of Y: given ε, a fixed power of x can be much larger than Y. For example, one cannot pay for a factor m merely by a fixed power of Y uniformly for all small ε. The factor m is only harmless when divided by a quantity containing a corresponding power of x or by a positive power of b. This observation explains some of the delicate bookkeeping in the local Hecke repairs and crystalline loss estimates. The three-term avoidance property of the thick places is not an artifact of this elementary dominance calculus.

26.2 Rounding of large coefficients and a Lipschitz estimate

Recall the thick and thin rounded weights $b_1 = N_1\lceil a/(xN_1)\rceil$ and $b_2 = N_1\lceil a/(x_1N_1)\rceil$. Then

$$b_1 = a/x + O(N_1) = b + O(N_1), \qquad b_2 = a/x_1 + O(N_1).$$

Since $N_1, h_2 = r^8$ and m are bounded by fixed powers of x, and b/x_1 grows faster than every fixed power of x, the rounding and gap terms satisfy, uniformly in $0 \le j \le 2m$,

$$0 \le w_j(\tau) - w_{\mathrm{nom}}(\tau) \le N_1 + (4m+1)h_2 = O(x^{81/2}) \le x^{42}.$$

Thus $w_j \le 2b$ eventually and

$$\log w_j - \log w_{\mathrm{nom}} \le x^{42}/w_{\mathrm{nom}} \le x^{-1}.$$

This proves the numerical comparisons of `New/Data/Weights.lean` along our family. We use the following useful estimate for small changes in the weights of a finite prime sum. For two families of real weights w, w' set $A_w(y,\sigma) = \sum_{\nu \le a} \lambda_{d,1}(\nu)(y/\nu - w([-\nu]\sigma))_+$. If $|w - w'| \le \Delta$ and $w, w' \ge y/Z$ pointwise, then

$$|A_w(y,\sigma) - A_{w'}(y,\sigma)| \le \Delta \sum_{\nu \le \min(a,Z)} \lambda_{d,1}(\nu) \le 2\Delta \sum_{\nu \le Z} \log \nu \le 2\Delta Z \log(1+Z).$$

Indeed the positive-part function is 1-Lipschitz and each summand vanishes for $\nu > Z$. The same reasoning at a single prime gives

$$|A_{p,w}(y,\sigma) - A_{p,w'}(y,\sigma)| \le \omega_p \Delta \lfloor \log_p Z \rfloor.$$

The latter is more useful when the bounds need to be summed over local loss functions, because it does not multiply the error by Z. In the case at hand one may take $Z = x_1$ and $\Delta = x^{42}$ for $y \le a$ after increasing the weights. Crude errors polynomial in x times x_1 are absorbed by Lemma 26.1. They must be absorbed before concluding a statement at the prime-power scale: replacing w_j by w_{nom} silently inside a coefficient representation would change the representation.

26.3 The exact level cancellation revisited

In the notation of Chapter 4, let $\bar{w} = f^{-1}\sum_{\tau} \log w_{\mathrm{nom}}(\tau)$, $G = \log(g\,\mathrm{rd}_E/(2\pi))$, and $\kappa_K = \gamma - \Re L'(1,\chi_K)/L(1,\chi_K)$. Then

$$\log(\Omega/g) = \kappa_K - G, \quad J(a) = \log(a!) - a\kappa_K + aG, \quad A_0 = a(\log a - 1 - \kappa_K - \bar{w}).$$

The only large difference between these two quantities is

$$J(a) - A_0 = a(G + \bar{w}) + O(\log(a+1)),$$

by Stirling's formula, or by comparing the sum $\sum_{i=1}^{a} \log i$ to its integral. The level prime is chosen so that

$$\log N \geq m(G + \bar{w}) - G - \tfrac{15}{2}Y.$$

Consequently

$$\begin{aligned} |\mathcal{T}|(m(J(a) - A_0) - a\log N) &\leq |\mathcal{T}|(a(G + \tfrac{15}{2}Y) + O(m\log(a+1))) \\ &= O(faY). \end{aligned}$$

We used $G = O(Y)$, $m/a \to 0$, and $\log a = O(Y)$. The exact source allows the larger error $B = faY^3$; the displayed upper bound is eventually smaller than that allowance. Thus the heights being compared can contain a very large common term H, and the cutoff-dependent excess at the distinguished columns is paid by the rational logarithm of N. Its choice is possible precisely because λ_N lies in the range of logarithmic prime windows in Proposition 4.2. It is not necessary to have a strong pointwise estimate for the prime error at x_1 for this prime choice.

26.4 An elementary bound on the logarithm of the class number

For the global Selmer bounds it is useful to avoid treating the class number as an unspecified enormous constant depending on the growing field E. Let $D_E = [E : \mathbb{Q}] = 2f$ and let Δ_E be the absolute discriminant. Minkowski's ideal class theorem says that every ideal class has an integral representative of absolute norm at most $(4/\pi)^f D_E!/D_E^{D_E}\sqrt{|\Delta_E|}$ (see the geometry-of-numbers treatment of number fields, for instance [11]). For a crude upper bound it suffices to use

$$M_E = 2^{D_E}\sqrt{|\Delta_E|}.$$

For each rational prime p there are at most D_E primes of E above it; fix an arbitrary numbering of these for each p. Map an integral ideal I, supported on these primes with exponents $e_{p,i}$ and residue degrees $f_{p,i}$, to the ordered tuple of positive integers

$$(n_1, \dots, n_{D_E}), \qquad n_i = \prod_p p^{f_{p,i}e_{p,i}}.$$

Fill missing indices with exponent zero. The tuple determines every ideal exponent by ordinary prime factorization, and $\prod_i n_i = N(I)$. Hence the number of integral ideals of norm at most M is at most $(\lfloor M \rfloor)^{D_E}$. In particular

$$\log h_E \leq D_E \log M_E = D_E^2\left(\log 2 + \tfrac{1}{2}\log \mathrm{rd}_E\right).$$

Using $f = O(x^{11}\log x)$ and $\mathrm{rd}_E \leq x^{12}$, we get $\log h_E = O(x^{22}(\log x)^3) \leq x^{23}$ for sufficiently large x. This estimate is far from optimal but adequate for the global budget. It is compatible with the bound used in `exists_isFitFamily`. Notice the difference between a bound for h_E and a bound for its logarithm; the latter is what appears when valuations are summed with log-prime weights.

26.5 Summation of the fit-family costs

We use the same notations as Chapters 20 and 23. Let M_g be the integer clearing all chosen-level group orders, and let X be the integer controlling the odd Tate classes (from the formal-source input concerning those classes). Its size is bounded in that input by

$$\log X \le 4(N_E + t + 3)^{27} \log 2 + \log h_E,$$

where $N_E = N_1 d$ denotes the cyclotomic conductor used in the Selmer calculation. In particular it is bounded by a fixed power of x. We require only this bound, not a proof of the existence of the odd-class annihilator here. Corollary 5.3 gives $\log M_g \le x^{32}$ eventually. Write

$$c_\ell = v_\ell(M_g) + 2^{-\ell}, \qquad D_\ell = h_2^{27}(c_\ell + v_\ell(X)) + 4f \operatorname{loss}_\ell,$$

with loss_ℓ as defined in the crystalline chapter. The series $\sum_{n\ge 2} 2^{-n} \log n$ converges, since $\log n \le n$. By ordinary factorization, for every finite set of primes T,

$$\sum_{\ell \in T} 9 c_\ell \log \ell \le 9(x^{32} + O(1)) \le x^{999}$$

eventually. This establishes the global inverse-ratio cost without any estimate depending on the particular chosen twists.

For the other budget the part

$$h_2^{27}\big(\log M_g + \log X + O(1)\big)$$

is bounded by a fixed power of x, since $h_2 = r^8$ and $\log X$ has the preceding bound. Hence it is $o(B_0)$ by Lemma 26.1. The local-loss part is calculated in Chapter 23. For clarity, let

$$J_* = (2b_1)! \left(\prod_{\max(b_1 - 30h_2^3, 0) < i \le b_1 + 30h_2^3} i \right)^{h_2^3}.$$

The actual local formula is $\operatorname{loss}_\ell = 10^6 (m/k)(b_1 v_\ell(g) + v_\ell(J_*))$. Summing over any finite set of primes gives

$$\sum_{\ell \in T} 4f \operatorname{loss}_\ell \log \ell \le 4 \cdot 10^6 \frac{fm}{k} (b_1 \log g + \log J_*).$$

Now $\log J_* \le 2b_1 \log(2b_1) + 60 h_2^6 \log(b_1 + 30h_2^3) = O(bY)$ by the elementary factorial bound and Lemma 26.1. Thus the preceding expression is $O(fmbY/k)$, and

$$\frac{fmbY/k}{B_0} = \frac{m^2 b}{aY^2} \le \frac{4}{Y^2} \longrightarrow 0.$$

Combining the two pieces yields $\sum_{\ell \in T} D_\ell \log \ell \le B_0$ for any finite set T and all sufficiently large x. In particular the inequalities called `Data.le_of_hasLogBound` in the summation of `exists_isFitFamily` can here be obtained from simple asymptotic inequalities. It is the proof of the block loss and the construction of the odd-class integer X, rather than this last summation, that remain the crucial arithmetic obligations.

26.6 The final numerical inequalities

Two further absorptions used in the reconstruction can be verified without the formal logarithmic size calculus. The absolute size estimate $|H| + m|\mathcal{R}| \le b^2 x^{99}$ follows from the definitions and the crude $\lambda_{d,1}(\nu) \le 2\nu$; its proof is given in Lemma 27.4. The main combination in Proposition 27.5 gives

$$m\mathcal{R} \le 4L + 4x^{99} + 10\log(1+b^2)B.$$

Divide the right side by BY^2. The first and third summands after division are $4/Y$ and $10(2Y+o(1))/Y^2$, both tending to zero; the middle summand tends to zero by Lemma 26.1. In particular the sum is eventually at most one. Similarly the single-lift allowance in Chapter 17,

$$3nrB_0 + n\bigl(10mx^{999} + 6kB_0 - k\mathcal{U}\bigr) \le nk(YB_0 - \mathcal{U}),$$

is equivalent after cancellation and division by nkB_0 to $3r/k + 6 + 10mx^{999}/(kB_0) \le Y$, which holds eventually. Here $r/k \to 1$ and the last quotient tends to zero. This calculation does not require positivity of $\mathcal{U}$.

Finally $|\mathcal{S}| \ge T_s/Y$, so the major arithmetic error after division by $ma|\mathcal{S}|$ is bounded by $\varphi(N_0)Y^6/m \le \varphi(N_0)x^{-1/2+6\varepsilon}$. The window error is x^{-10^6}. The fixed constant $\varphi(N_0)$ can be absorbed by $x^{2\varepsilon}$ for large x, giving $x^{-1/2+8\varepsilon}$. This demonstrates explicitly why the slack in each numerical budget does not affect the target exponent and why the number of lifts is the one important scale that must be retained. All thresholds may depend on the fixed auxiliary modulus and field. The zero-free criterion chooses the auxiliary prime after the hypothetical zero and then uses the eventual bound for that fixed prime, so no uniformity in d has been assumed.

Chapter 27

The arithmetic inequality and conclusion

We assemble the three arithmetic results in their exact quantitative form. The proof of their consequence is independent of the internal choices of kernels.

27.1 Collecting the arithmetic inputs

With the data of Chapter 4, set $n = \dim_{\mathbb{C}} W$, $c = \dim_{\mathbb{C}} C$, and retain

$$H = mk\mathcal{U}, \qquad B = fa(\log b)^3, \qquad L = B \log b.$$

For clarity, we give the particular form of the three theorems that is needed in the calculation. The notation in the three statements below is intentionally the same. The preceding sections give their derivation and identify specialized inputs still requiring independent written proofs.

Theorem 27.1 (A: Short forms, recalled)**.** *There are a rational number $q_0 > 0$ and a set $S \subseteq C^{\mathrm{ram}}$ of integral forms spanning C^{ram} over $\mathbb{C}$ such that, putting $T_s = q_0 P_s$, one has*

$$\|\Theta_1 T_s(v)\|_\sigma^2 \le \exp(H - m\mathcal{R} + L) \quad (v \in S,\ s \in \mathcal{G},\ \sigma \in \mathrm{Aut}(\mathbb{C})).$$

For every finite set $F \subseteq \mathcal{G}$ and every finite family y in the images $T_s(C^{\mathrm{ram}}_{\mathrm{int}})$, $s \in F$, one has

$$\mathrm{IndexLe}(y, |F| cB).$$

Theorem 27.2 (B: Enough forms, recalled)**.** *The spaces C and W have positive dimension and there is a finite set $F \subseteq \mathcal{G}$ such that*

$$|F| c \le 5n \log(1 + b^2), \qquad \dim\left(\sum_{s \in F} \mathrm{im}(P_s)\right) \ge (1 - 4/b^2)n.$$

Theorem 27.3 (C: Integral forms are long, recalled)**.** *Let $y = (y_1, \dots, y_q)$ be q independent row forms with $q \ge (1 - 4/b^2)n$. If $\mathrm{IndexLe}(y, d_{\mathrm{ind}})$ and for every σ one has $\mathrm{covol}_\sigma^2(y) \le e^V$, then*

$$n(H - 3L) \le V + 2d_{\mathrm{ind}}.$$

Together with the crude bound

$$|H| + m|\mathcal{R}| \le b^2 x^{99}, \tag{27.1}$$

these are precisely the inputs in `star_inequality` of `New/Main.lean`. The source proves the latter as `Data.abs_mainHeight_add_remainder_le`. We give the elementary verification along the admissible asymptotic family.

Lemma 27.4 (Crude size comparison)**.** *Along the data in Proposition 4.2, including the chosen second cutoff, inequality* (27.1) *holds for all sufficiently large x.*

Proof. The factor $\lambda_{d,1}(\nu)$ is bounded by 2ν. Thus $0 \le A(z,\sigma) \le 2az$ for $0 \le z \le a$, since each summand is at most $2z$. From $w_j \le 2b$, $|\log(\Omega/g)| \ll_{N_0,d} \log b$ (by (4.2) and the bounds on g), and $\log(j!) \le j \log j$, we have $|J(w_0(\sigma))| \ll b \log b$. The nominal weights range between a/x_1 and a/x, so $|A_0| \ll a \log b$. From the definitions,

$$|H| \ll mkf(b \log b + ab), \qquad m|\mathcal{R}| \ll mf(a^2 + a \log b).$$

Here $a \le 2xb$, $k \ll x^5$, $f \ll_{N_0} x^{11} \log x$, $m \ll x^{1/2}$, and $\log b = x^{\varepsilon} \le x$. These upper bounds are $O_{N_0,d}(b^2 x^{20})$ for large x. Enlarging the threshold gives (27.1). The proof in the compact source instead uses explicit logarithmic inequalities to have the same bound on the full admissible record, a stronger quantification not needed for the eventual inequality here. □

27.2 From A and B to an almost-basis of short forms

Choose F as in Theorem 27.2 and use the set S from Theorem 27.1. Since q_0 is nonzero and S spans C^{ram}, the set

$$\{T_s(v) : s \in F,\ v \in S\}$$

spans the sum of the images of the P_s, $s \in F$. Choose independent vectors $y_1, \dots, y_q$ from this set which form a basis of its span. In particular,

$$(1 - 4/b^2)n \le q \le n.$$

By A and the bound on the number of translates we may take

$$d_{\mathrm{ind}} \le 5n \log(1 + b^2)B.$$

For each conjugated pairing Hadamard's inequality gives

$$\begin{aligned} \mathrm{covol}_\sigma^2(y) &\le \prod_{i=1}^{q} \|\Theta_1 y_i\|_\sigma^2 \\ &\le \exp\left(q(H - m\mathcal{R} + L)\right). \end{aligned}$$

Thus $V = q(H - m\mathcal{R} + L)$ is a valid upper bound. This is the conclusion of `Data.exists_family_covolumeSq_le` in `New/Endgame/LongForms.lean`. The step is only linear algebra, but realizes the crucial purpose of Theorem B: it converts estimates on vectors arising from a particular automorphic construction into an estimate on almost the entire determinant of the row space.

27.3 Comparison with the lower bound

Proposition 27.5 (The main arithmetic inequality)**.** *For all sufficiently large x along the admissible family of Proposition 4.2, the three arithmetic inputs imply*

$$m\mathcal{R} \le fa(\log b)^5. \tag{$\star$}$$

Proof. Write $h = H - m\mathcal{R} + L$ and apply C to the family just constructed. Then

$$\begin{aligned} n(H - 3L) &\le qh + 2d_{\mathrm{ind}} \\ &= nh + (n - q)(-h) + 2d_{\mathrm{ind}}. \end{aligned}$$

We cannot simply replace q by n in qh, because h may be negative. Instead, use (27.1) and $L > 0$ to get $-h \le b^2 x^{99}$, and the dimension estimate to get $n - q \le 4n/b^2$. Consequently

$$n(H - 3L) \le n(H - m\mathcal{R} + L) + 4nx^{99} + 10n \log(1 + b^2)B.$$

Divide by $n > 0$, cancel the common H, and rearrange:

$$m\mathcal{R} \le 4L + 4x^{99} + 10 \log(1 + b^2)B. \tag{27.2}$$

The scale inequalities imply that the right-hand side is at most $(\log b)^2 B$. For the eventual assertion this is immediate from $a \ge xb$, $B = fa(\log b)^3$, and $\log b = x^{\varepsilon}$, since $x^{99}/B \to 0$ and $\log(1 + b^2) = O(\log b)$. The Lean version uses its explicit admissibility bounds to make this last comparison for every admissible datum, not just sufficiently large x along this family. Substituting the value of B gives ($\star$). □

It is worth looking closely at the size of the three errors in (27.2). The $4L$ is the sum of the error per vector in the upper bound and the three units of L in the lower bound. The $4x^{99}$ is the penalty for the missing dimensions, which are so few that even a very crude bound on h suffices. The last term is twice the index. All are smaller than the final budget by a factor of order $\log b$ or more along the asymptotic family. The leading term H has disappeared completely. Without the matching of its exact leading coefficient in A and C, this argument would not give the desired bound.

27.4 Extracting the analytic remainder

Here is the short algebra behind (4.9). Use the notation of Chapter 2 and write

$$S_M(z) = \sum_{\nu \le z} \lambda_{d,M}(\nu)(1/\nu - 1/z).$$

For a fixed integer ν, the number of $\sigma \in \mathcal{T}$ such that the place $s_\nu \sigma$ lies in $\mathcal{S}$ is $|\mathcal{S}|$ if $\nu \equiv 1 \pmod{N_0}$ and zero otherwise. Indeed, multiplication by s_ν carries the whole coset $\mathcal{T}$ onto the coset of class $-\nu$ when ν is invertible modulo N_1. The thick places have class -1 modulo N_0. Every $\nu \equiv 1 \pmod{N_0}$ is a unit modulo N_1 in this tower. In the remaining nonunit cases the source's place convention places the product back in $\mathcal{T}$, so it is not thick. It follows that replacing the weight a/x_1 by a/x inside the sum over the tops gives the identity

$$\frac{1}{a} \sum_{\sigma \in \mathcal{T}} A(a, \sigma) = |\mathcal{S}|\big(S_{N_0}(x) - S_{N_0}(x_1)\big) + |\mathcal{T}| S_1(x_1).$$

Here S_M denotes the sum before subtraction of its main terms. By the analytic normalization,

$$S_M(z) = R_M(z) + \frac{\log z}{\varphi(M)} + c_{d,M}.$$

The constants for $M = N_0$ cancel in the difference. For $M = 1$, we have $c_{d,1} = -1 - \kappa_K$. Moreover

$$\log a - \overline{w} = \log x_1 + \frac{|\mathcal{S}|}{f}(\log x - \log x_1).$$

Since $f = \varphi(N_0)|\mathcal{T}|$, we have exactly

$$\frac{|\mathcal{T}|A_0}{a} = |\mathcal{T}|(\log x_1 - 1 - \kappa_K) + \frac{|\mathcal{S}|}{\varphi(N_0)}(\log x - \log x_1).$$

Subtract these main terms from the previous sum identity. This proves (4.9), including its constants. In particular, no estimate involving the Chowla–Selberg constant is needed at this final stage.

27.5 Choosing the second cutoff and the power saving

Fix $0 < \varepsilon < 1/100$ and choose admissible data by Proposition 4.2 for every sufficiently large x. The quantity $X = |\mathcal{T}|/|\mathcal{S}|$ satisfies $0 < X \le \log b$. Write $T = 10^6 \log x$. For large x we have $X \le e^T$. Apply the window lemma of Chapter 2 to the two functions $u \mapsto R_{N_0}(e^u)$ and $u \mapsto R_1(e^u)$. It supplies $u \in [T, T^2]$ such that

$$R_{N_0}(e^u) - \frac{|\mathcal{T}|}{|\mathcal{S}|} R_1(e^u) \le e^{-T}.$$

Take $x_1 = e^u$ in the auxiliary data, as permitted by Proposition 4.2. This is the decisive use of the windowed argument: there is no need for a strong pointwise estimate at the initially specified, extremely large, second scale.

Solve (4.9) for the desired term and use ($\star$):

$$\begin{aligned} R_{N_0}(x) &= \frac{\mathcal{R}}{a|\mathcal{S}|} + R_{N_0}(x_1) - \frac{|\mathcal{T}|}{|\mathcal{S}|} R_1(x_1) \\ &\le \frac{f(\log b)^5}{m|\mathcal{S}|} + e^{-T} \\ &\le \frac{\varphi(N_0)(\log b)^6}{m} + x^{-10^6} \\ &\le \varphi(N_0) x^{-1/2+6\varepsilon} + x^{-10^6}. \end{aligned}$$

Absorbing the fixed constant using 2ε of additional slack gives, for all sufficiently large x,

$$R_{N_0}(x) \le x^{-1/2+8\varepsilon}. \tag{27.3}$$

This is the precise exponent in the compact development's `eventually_primeError_le_rpow`. Replacing ε by a smaller positive number proves a bound with any fixed exponent greater than $-1/2$.

Remark 27.6 (Where the critical line enters). Everything up to ($\star$) concerns norms, divisibilities, dimensions and heights. The saving $x^{-1/2}$ arises from dividing the error by $m \ge \sqrt{x}$. The large field degree, large weights and rank do not themselves account for it. One cannot hope to improve the analytic conclusion by choosing m arbitrarily larger: the arithmetic theorems and their errors depend crucially on the admissible relationship between the length, the weights, ranks and local losses. A proof of stronger short/long-form estimates on other ranges would be an additional mathematical result.

27.6 GRH

To finish, suppose a Dirichlet character χ modulo q had a zero ρ to the right of $1/2$, excluding the pole convention of Theorem 1.1. Choose $N_0 = 18480q$ and regard χ as a character of level

N_0 by induction; the induced L-function differs by finite Euler factors with no zeros in this region. The class $N_0/2 - 1$ is a unit modulo N_0 and forces $d \equiv 7 \pmod 8$. For every prime d in this class, the above arithmetic argument gives (27.3). Choose ε positive and sufficiently small that $\beta = 1/2 + 8\varepsilon < \Re(\rho)$. The one-sided criterion of Proposition 2.9 forbids the zero ρ. This proves the deduction of Theorem 1.1 from the arithmetic inputs; a complete conventional proof requires independently establishing the remaining formal-source inputs listed in the final section. The standard functional equations imply the statements about the critical strip. This completes the deduction.

Chapter 28

Glossary and concordance

The following glossary is intended for returning to one chapter after reading another. Every symbol has a local role; the numerical estimates only acquire their meaning after the objects and all normalization factors have been fixed. We list central Lean declarations for the compact edition. Where the same symbol is used in a different context within a local calculation (for example the rank of a single lattice), the text specifies the local meaning.

Symbol or term	Meaning	Source or chapter
N_0	Fixed auxiliary modulus divisible by 18480; chosen to be a multiple of the given Dirichlet modulus.	`Parameters.N₀`; 4
d	Auxiliary prime $\equiv N_0/2 - 1 \pmod{N_0}$; varied in the zero-free criterion.	`Parameters.d`; 2
K	Imaginary quadratic field $\mathbb{Q}(\sqrt{-d})$.	`Parameters.K`
E/F	CM extension $E = K(\zeta_{N_1})$, $F = E^+$; field changes with the scale.	`Data.E`; 4
Φ, f	CM-type embeddings over K, and their number $f = [E : \mathbb{Q}]/2$.	`Data.Φ`, `Data.f`
s, T_s, N_1	Tower exponent, $T_s = 3^{\lfloor \log_3 s \rfloor + 5} 2^s$, $N_1 = N_0 T_s$.	`towerDeg`, `Data.N₁`
$\mathcal{T}, \mathcal{S}$	Tops of cyclotomic class 1 modulo N_0; thick places in class -1 with a multiple-progression avoidance property.	`Data.tops`, `Data.thick`
x, x_1, b	Prime error at x, auxiliary second cutoff, much larger order scale b with $\log b = x^{\varepsilon}$.	`Data.x`, `Data.x₁`, `Data.b`
m, r, k, t	Number of lifts, row rank, column rank $k = r + 1 - m$, row increment $t = 5r$.	`Data.m`, `Data.r`, `Data.k`, `Data.t`
a, b_1, b_2	Rounded order $a \geq xb$ and rounded weights $\geq a/x, a/x_1$, all multiples of N_1.	`Data.a`, `Data.b₁`, `Data.b₂`
h_2, w_j	Half-gap r^8, integral weight with increment $2h_2$ per stage.	`Data.h₂`, `Data.w`
w_{nom}	Effective weight a/x at thick places, a/x_1 elsewhere; enters the exact prime sum.	`Data.nominalWeight`
g_0, g_1, g	Two chosen split auxiliary primes and their product, $g = g_0 g_1$.	`Data.g`, `Data.gProd`
N	Extra prime level on column groups; selected to pay for the main archimedean excess.	`Data.N`; 26.3
M_ω	Auxiliary integer $(N_0 d)^{s+6}$ specifying tame-order characters at chosen places.	`Parameters.Mω`
X_j, Y_j	Definite Hermitian spaces of rank $k + j(t+1)$, row spaces of rank $r + jt$.	10

Symbol or term	Meaning	Source or chapter
C, C^{ram}, W	Column form space, image of the ramifying operator, and good row form space; $c = \dim C$, $n = \dim W$.	`Column`, `ramifiedColumns`, `rows`
$\mathcal{G}$	Product of the finite GL_r at chosen places above g_0, g_1; indexes projected translates.	`Data.G`; 14
e_0	Central orthogonal projector onto the specified good finite types.	`Data.e₀`
P_s	Projected Poincaré maps, with the common normalization adopted in Theorem A.	`Data.projectedPoincare`
θ, Θ_θ	Weight-zero tame m-fold twist; composite row theta lift with the prescribed powers of the twist.	10
ι	Field isomorphism to a p-adic algebraically closed coefficient field; never assumed continuous from $\mathbb{C}$.	`PadicIso`
$\mathrm{IndexLe}(y, d_{\mathrm{ind}})$	Logarithmic cost of integral combination matrices for an independent or dependent row family at all finite-place embeddings and all prime-to-p twists.	`Data.IndexLe`; 5.10
V	Log upper bound on the squared covolume of a lifted family, uniform in its coefficient conjugates.	17
$\lambda_{d,M}(\nu)$	$\mathbf{1}_{\nu \equiv 1 \bmod M}\Lambda(\nu)(1 - \chi_K(\nu))$.	`lambdaClass`; 2
$R_M(z)$	Smoothed normalized prime remainder, after subtracting $\log z/\varphi(M)$ and the constant $c_{d,M}$.	`primeError`
$A(y, \sigma)$	Unnormalized finite prime sum with weights $w_{\mathrm{nom}}([-\nu]\sigma)$ and cutoff a.	`Data.primeSum`
$A_p(y, \sigma)$	Local positive-part sum for the live powers of a nonsplit p, with the factor $(1 - \chi_K(p))/2$.	`Data.localAsk`; 9
g_p	Desired p-adic gain over the entire chain, m times the sum of the distinguished long-row gain and k ordinary-direction gains.	`Data.gain`
$\Omega, J(j)$	Period scale $2\pi \exp(\kappa_K)/\operatorname{rd}_E$ and derivative cost $\log j! - j \log(\Omega/g)$.	`Parameters.bigOmega`, `Data.jetSize`
$\mathcal{U}, H$	Net ordinary direction cost and common main term $H = mk\mathcal{U}$.	`Data.dirCost`, `Data.mainHeight`
$\mathcal{R}$	Arithmetic remainder after subtracting the common mean term from the sum at all tops.	`Data.remainder`
B, L, B_0	$fa(\log b)^3$, $fa(\log b)^4$, and $B/(mk)$, respectively.	`Data.budget`; 26
$Z(\theta)$	Squared volume of the row-chain map at a twist, from its Gram determinant in the given metric.	17
Fit family	A choice of twist at each embedding, controlling determinants of tested projections and ratios of Hecke values, for every power through m.	`Data.IsFitFamily`; 20
Local block loss	Cost of imposing a finite/crystalline relation on torsion cocycles in the row and Tate-string blocks at a rational prime.	`loss`; 23
$\star$	Main inequality $m\mathcal{R} \le fa(\log b)^5$.	`star_inequality`; 27

28.1 Reading the main theorem in the Lean source

To compare the monograph with the formal proof, begin with the five exported statements in `Main.lean` and the real statement definitions in `Statements.lean`; these do not depend on the specialized definitions of the auxiliary construction. Then read `New/Main.lean`, where the endgame is expressed in only a few declarations. The upper and spanning inputs are `Data.short_forms` and `Data.enough_forms`, in `New/Endgame/ShortAndEnough.lean`. The lower input is `Data.integral_forms_long`, in `New/Endgame/LongForms.lean`. The final

chosen data come from `exists_data` in `New/Data/DataExist.lean`. The analytic objects and criteria are in `New/Endgame/PrimeError.lean` and `New/Endgame/OneSidedCriterion.lean`. This is the shortest route through the dependency logic. One cannot read the code at that level to verify the arithmetic input theorems themselves, but one can check the precise mathematical interface against the deductions in the final part of the volume.

At a second level, the modules `New/Curve/`, `New/Columns/`, `New/TwistedHecke/` and `New/Seesaw/` implement A; `New/Filling/`, `New/RealPlace/` and `New/DiscreteSeries/` implement B; `New/Rows/`, `New/Exchange/`, `New/Selmer/`, `New/Patching/` and `New/Lift/` implement the principal pieces of C. Definitions and auxiliary lemmas in the `Lib/` tree are shared or specialized at these interfaces. The readable edition of the repository is often easier for understanding motivation or naming a mathematical package: consult its `GUIDE.md`, `INDEX.md` and `CONTENTS.md`. The compact edition is generally easier to use for extracting exact constants. The two should not be conflated when a particular theorem has been reworked mathematically, notably the analytic choice of the second cutoff.

Chapter 29

Dependency ledger and the status of the reconstruction

29.1 Three levels of mathematical justification

The formal source and a paper proof are not interchangeable objects. The source artifact [1] reports a kernel-checked proof of the five statements in its `src/Statements.lean` and `src/Main.lean`, using only the standard Lean axioms `propext`, `Classical.choice`, and `Quot.sound`. This volume is a mathematical reconstruction of that argument. We have not rebuilt or replayed the entire formal development for this exposition. A correct paper proof of the final claim would moreover require that every step in the mathematical translation, as opposed to its formal counterpart, is correct and justified. In particular the presence of a proof of a Lean declaration does not certify a paraphrase of its statement or proof in a book.

The reader can distinguish three different kinds of material here.

1. **Mathematical derivations in the text.** These may be elementary but nontrivial, such as the finite Hecke elementary-divisor calculation, or substantial analytic and integral arguments. Their proofs refer only to the hypotheses used, other proved statements, and the classical or specialized inputs explicitly specified. A proof of a conditional implication is not a proof of its hypotheses.

2. **Standard external mathematics.** A research monograph need not prove every foundational theorem. It can use, for example, class field theory, a classical mass formula, or the usual rational p-adic Hodge comparison formalism, if the hypotheses and conclusion really match the application. A specialized estimate uniform in a growing number field, degree, weight and level is not automatically a consequence of a classical qualitative theorem.

3. **Specialized formal-source inputs.** A statement described in the text as such, or an explicitly retained comparison or existence assertion, is a mathematical interface with a theorem in the supplied Lean source whose complete independent proof has not been reconstructed in this volume. This status represents a gap in the conventional exposition, not an extra axiom asserted for the underlying formal proof. A complete translation would have to discharge it, by a conventional proof or a genuinely applicable external theorem. Simply collecting several inputs into a larger proposition does not discharge any of them.

The fourth exposition has been designed to be substantially more autonomous than its predecessors: it supplies the classical analytic prerequisites and derives the shared global and local algebraic normalizations, then develops many detailed mathematical proofs that

were previously only interfaces. The point of this ledger is to allow a reader to identify the remaining mathematical obligations, rather than count pages as a substitute for verification. A positive assessment of the text of one lemma does not verify the complete chain of all of its inputs.

29.2 Where to find the added derivations

The next table is a reading guide and a record of the principal additional work. The third column indicates the types of specialized assertions which cannot simply be inferred from the derivations in the second column. Further exact anchors are given in the indicated chapters themselves.

Area	Mathematical derivations	Remaining dependencies or standard inputs
Analysis (Chs. 2, 3)	Character sums, reciprocity and quadratic mean-square bounds, functional equations, continuation and classical boundary nonvanishing, reciprocal prime divergence, uniform prime-window lemma, analytic zero detector and high-moment Landau window.	The arithmetic estimates A, B, C are hypotheses of the final analytic deduction; no strong prime number theorem or uniform Siegel–Walfisz premise is needed.
Arithmetic models (Chs. 4, 5, 6, 7)	Parameter construction; class-sum and coefficient structures, splitting-character extension, self-dual different idele, local norms and unramified integral Gram–Schmidt, framed signed global spaces, finite group orders, norm-one class sum, uniform mass estimate, finite lattice multiplier and normalized Heisenberg cocycle.	Classical global reciprocity and norm subgroups, Landherr's theorem, Hermitian mass formula, integral Schur basis and standard Heisenberg representation theory. Real phase and theta-functional comparison with the nonpolar source model remain separate tasks.
CM, theta chains and Hecke (Chs. 9, 8, 10, 11, 12)	Arithmetic and formal CM calculations, normalized section identities, local weighted gains, abstract Frobenius root and elementary-divisor arguments, explicit character and lattice computations in the seesaw and its unfolding.	Classical CM and Weil foundations in their precise normalizations, and any remaining specialized support, comparison and uniform coefficient estimates identified in Chapter 12.
Spanning (Chs. 13, 14, 15, 16)	Formal degree via matrix beta integral, elliptic support and finite trace algebra, local kernel and Bessel identities, finite group counting, analytic moments and Fourier estimates, uniform numerical inequalities and the finite sieve, reduction of global matrix estimates.	Stable elliptic orbital comparison and rigid finite-label hypotheses in the transfer, and the geometric real-place and filling comparisons not proved in Chapter 16. All uniform estimates must have the specified constants.

Area	Mathematical derivations	Remaining dependencies or standard inputs
Determinant and periods (Chs. 17, 18, 19)	Complementary frames and period matching; integral gluing, finite Fourier and exchange algebra; determinant and norm comparisons; elimination and reduction of the remaining global period input.	Specialized analytic transfer, normalized global period/nonvanishing and comparison statements not independently established in Chapter 19. The classical Siegel–Weil and Rallis theories must be applied with the same levels, characters and measures.
Selmer and congruences (Chs. 20, 21, 22)	Auxiliary Kummer and triangle calculations, row and attachment reductions, finite twisting algebra, Stickelberger and congruence estimates, integral determinant lifting and the assembly of families.	Genuine global attachment, realization, duality and reciprocity comparisons not proved in Chapter 22; the specialized lattice index requirements of the fit family are not supplied by the mere existence of a twist.
Crystalline estimates (Chs. 23, 24, 25)	Crystalline rank and index, frame and Frobenius algebra, Kummer charts, depth and long-gap arithmetic, safe exponent and length estimates, tame/cut and window reductions.	Period-ring and integral comparison hypotheses stated in Chapter 25, and any specialized window/spacing comparisons still retained there.
Assembly (Chs. 26, 27)	Exact weighted-prime remainder, cancellation of the main term, extensive asymptotic loss budget and final reduction.	Coherence of all the preceding constructions and completeness of A, B and C.

29.3 The exact logical role of the remaining work

The deductions in Chapter 27 take three simultaneous statements about a single choice of the admissible data. Assertion A produces a scaled family of short lifts with bounded integral index; assertion B provides a family of translates spanning almost the required dimension using few source vectors; assertion C gives a determinant lower bound for every suitable independent family. An apparently close estimate in only one of these directions might not suffice: the three statements have a common leading term and a common row space, and their use of the same levels and integral coefficient coordinates is essential.

Specific reductions achieved in this edition. The following is a short checklist of what can now be read as a conventional proof (occasionally conditional on a specified comparison), and not merely as an imported interface. The stated qualifications are essential.

- The normalized splitting character and the self-dual idele are constructed in Proposition 6.4 and Lemma 6.6; the uniform mass ratio follows in Theorem 6.13 from the classical mass formula. Theorem 7.6 computes the finite vacuum phase, including its sign at inert primes, so no unnamed finite splitting normalization is needed at that point.
- Theorem 12.5 derives the soft normalized CM reciprocity law via marked elliptic curves, their sections and simultaneous torsion translations. Theorems 12.12 and the argument of Section 12.15 show which one-stage orbit, support and integral calculations now enter the iterated seesaw. The finite cuts and neighbour-lattice evaluation give the actual denominator propagation, conditional on the normalized geometric realization.

- Proposition 13.1 calculates the formal degree; the trace and row-dimension deduction in Chapter 13 explicitly separates the stable elliptic orbital identity and the rigid-label hypothesis from finite projection and dimension algebra. Thus the inequality (13.8) has a conventional conditional derivation; it is not an unconditional assertion of transfer for arbitrary automorphic spectra.

- The local matched-factor and stable synthesis arguments (Propositions 19.2 and 19.4), the square move of Section 19.8, and the integral gluing and gauge bounds of Propositions 19.13 and 19.14 supply mathematical results used in the exchange. The global normalized theta identity remains a separate obligation.

- Theorem 22.14 derives the explicit uniform odd-class bound using the finite Stickelberger and Kolyvagin calculations, rather than treating it as an extra hypothesis. The bounded-pencil result of Proposition 22.15 and the specialization in Section 22.10 reduce the common-twist problem to the precisely stated balance, mirror and integral value-measure hypotheses.

- Theorems 25.13, 25.16, 25.20 and 25.22 give the rank-free depth, long-gap Taylor, Frobenius safety and cut calculations from the indicated integral period-ring and radius inputs. These do not replace the first half of the radius-span comparison or the exact tame chart.

For the upper bound, the major dependency is not a routine application of the theta correspondence to an arbitrary vector. The exact compatible kernels, adjointness and integral support in `Data.seesaw` must realize the same short vectors as the Poincaré family. Chapter 12 develops the concrete CM and orbit algebra underlying this step. The residual input is the exact part of the compatible geometric and integral construction not established there; in particular a qualitative seesaw identity cannot replace uniform integral coefficient bounds. Principal source files are `New/Seesaw/SeesawTheorem.lean`, `UnfoldingStages.lean`, `GoodDatumExists.lean` and the local coefficient results in `New/Columns/` and `New/TwistedHecke/`.

More precisely, Section 12.15 identifies the real joint-harmonic type, the concrete finite projector and arithmetic Weil-vector realization, the exact even-stage coefficient formula `Data.unfoldStage_even`, and the final branching factor `Data.exists_last_stage_map`. The geometric orbit and cut estimates are not among these remaining black boxes; the use of the correct real form and integral coefficients is.

For spanning, the energy bounds are quantitative. The Bessel system must be the one whose real shape, finite flag and covariance produce the normalized global map. The projected moments must have the actual errors of Chapter 14; for example, pointwise rapid decay for fixed weights does not give the required uniform bound as the family grows. The relevant specialized proofs are in `New/Filling/FillingData.lean`, `Moments.lean`, `GlobalSieve.lean` and `New/RealPlace/`. The first-moment restriction must use the opposite factor to the translating central projector, as corrected in this edition. Section 16.18 distinguishes the new local shape and compact correction, marked-lattice and unipotent-covolume computations, finite flag propagation, sublattice-zeta and global frequency bounds, central projection and energy deductions from the residual real and finite floor-gauge and Fourier estimates, uniform joint-pair and saturation bounds, global quadric and top-packet limits, and the remaining simultaneous size and admissibility conditions. These latter assertions are needed to apply the quantitative sieve to the actual Poincaré family.

The determinant uses actual exchange and normalized period identities: the local gluing combinatorics alone does not imply the equality of global period norms. The precise anchors are `New/Exchange/`, `New/Rows/RowVolume.lean`, `LiftVolumeHeight.lean`, `HeckeValueHeight.lean` and the global comparison with the same coefficient lattice. The existence of the fit

family requires the actual weak Galois attachments, twisting and congruence-to-cohomology maps, together with the specified integral local estimates. These are not consequences of qualitative crystalline comparison alone. Chapters 19, 22 and 25 replace many of the constituent assertions with proofs and give the relevant remaining interfaces.

In particular the remaining period comparisons include the normalized top-period and line constant-term identities and convergence, the geometric centre coefficient vanishing, the real square law with its kernel height and type identities, the quantitative row restriction and carrier separation, the actual tower-vanishing and Hecke transfer, and depth-indicator attainability. The finite matching and square move do not prove these. For the congruence part see Section 22.14: a residual geometric seed and exact local–global type compatibility, global mirror and even-part reciprocity extraction, precise integral ramified-pair value measures, and the actual patched Ribet-to-Hecke lifting remain in the denominator step. In the crystalline part Section 25.16 isolates the structural integral period-ring comparison, the initial radius-span and derivation compatibility, the final few-low-jump conditions, and tame near-identity and polynomial-to-ideal recurrence. The compact and readable versions package some numerical hypotheses differently; their correspondence has to be verified before transporting a bound without change.

In summary, the final statement is a rigorous implication from the premises established or explicitly imported in the body of this volume. To regard the mathematical reconstruction itself as a standalone proof of the theorem one must verify those premises and the exact coherence conditions. We do not use the claimed formal verification as a premise about the truth of a new mathematical translation; the Lean source is the reference for the remaining mathematical interfaces.

29.4 What the compact and readable editions contribute

The compact edition specifies the shortest mathematical dependency path, whereas the readable edition sometimes separates intermediate claims into files of substantial standalone interest. Both have been used to identify the intended exact statements and recover the calculations. The compact repository's `REFERENCES.md` reports checks of 196 citations: 57 compared directly with precise locations, 30 identified at second hand, 102 located works without precise theorem identification and 7 not compared. We treat these as the source's reports, not as verification of the references by this exposition. In particular a theorem name in a comment does not cause a stronger quantitative version to follow from the published theorem.

The principal standard references used explicitly in this exposition are recorded at the end. Some mathematical arguments come from the source's more detailed proof steps rather than a previously published exposition of the GRH argument. The distinction between the standard background and the specialized statements matters even when both are proved formally in `Lib` and `New`, respectively. None of the remaining specialized input declarations should be cited as an independently proved theorem of this book merely because its counterpart is reported as accepted by Lean.

The intended contribution of these four expositions is therefore a succession of increasingly detailed mathematical reconstructions, with a common notation and increasingly explicit interfaces. The present expanded reference edition should make it easier to read the formal argument as mathematics, and to isolate specific obligations for independent verification, rather than treat the entire development as an indivisible black box.

Bibliography

[1] L. Alpöge, with Claude, and the authors of the shortened development, *GRH for Dirichlet L-functions, in Lean 4*, supplied snapshot `grh-lean-main`, November 2026. Compact and readable editions (Apache 2.0); the root README links the announcement https://alpo.ge/remark.pdf. Main statements: `src/Statements.lean`. Main mathematical interfaces: `src/New/Main.lean`, `src/New/Endgame/ShortAndEnough.lean`, `src/New/Endgame/LongForms.lean` in the compact edition.

[2] H. L. Montgomery and R. C. Vaughan, *Multiplicative Number Theory I: Classical Theory*, Cambridge Studies in Advanced Mathematics 97, Cambridge University Press, 2007. In particular Chs. 9–11 and Cor. 11.19, and Ch. 15 for Landau arguments. https://doi.org/10.1017/CBO9780511618314.

[3] M. Jutila, On mean values of Dirichlet polynomials with real characters, *Acta Arith.* **27** (1975), 191–198. https://doi.org/10.4064/aa-27-1-191-198.

[4] F. A. Behrend, On sets of integers which contain no three terms in arithmetical progression, *Proc. Natl. Acad. Sci. USA* **32** (1946), 331–332. https://doi.org/10.1073/pnas.32.12.331.

[5] L. C. Washington, *Introduction to Cyclotomic Fields*, 2nd ed., Graduate Texts in Mathematics 83, Springer, 1997. https://doi.org/10.1007/978-1-4612-1934-7.

[6] B. H. Gross, Algebraic modular forms, *Israel J. Math.* **113** (1999), 61–93. https://doi.org/10.1007/BF02780173.

[7] M. Greenberg and J. Voight, Lattice methods for algebraic modular forms on classical groups, arXiv:1209.2460. https://arxiv.org/abs/1209.2460.

[8] J. W. S. Cassels, *An Introduction to the Geometry of Numbers*, Springer, 1959; reprint of the 1971 edition in Classics in Mathematics, Springer, 1997. https://doi.org/10.1007/978-3-642-62035-5.

[9] C. De Concini, D. Eisenbud and C. Procesi, Young diagrams and determinantal varieties, *Invent. Math.* **56** (1980), 129–165. https://doi.org/10.1007/BF01392548.

[10] K. Akin, D. A. Buchsbaum and J. Weyman, Schur functors and Schur complexes, *Adv. Math.* **44** (1982), 207–278. https://doi.org/10.1016/0001-8708(82)90039-1.

[11] J. Neukirch, *Algebraic Number Theory*, Grundlehren der mathematischen Wissenschaften 322, Springer, 1999. https://doi.org/10.1007/978-3-662-03983-0.

[12] M. Kirschmer, *Definite quadratic and hermitian forms with small class number*, Habilitationsschrift, 2016. In particular Proposition 4.2.7 and Lemma 4.2.9 for the unitary mass formula and its local factors. https://www.math.uni-bielefeld.de/~mkirschm/papers/herm.pdf.

[13] W. Landherr, *Äquivalenz Hermitescher Formen über einem beliebigen algebraischen Zahlkörper*, Abh. Math. Sem. Univ. Hamburg **11** (1936), 245–248.

[14] A. Raghuram, Notes on the arithmetic of Hecke L-functions, 2022, https://arxiv.org/abs/2207.03399. General background on algebraic Hecke characters.

[15] A. Weil, Sur certains groupes d'opérateurs unitaires, *Acta Math.* **111** (1964), 143–211. https://doi.org/10.1007/BF02391012.

[16] S. S. Kudla, Splitting metaplectic covers of dual reductive pairs, *Israel J. Math.* **87** (1994), 361–401.