A research monograph

The Geometry and Arithmetic of Iterated Theta Lifts

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

Third exposition

Working mathematical manuscript – November 2026

Scope. These notes develop the specialist manuscript into a research-level monograph. They explain the adelic and integral coefficient structures, normalized CM theta series, local divisibility and Hecke repair, Bessel and Poincaré estimates, determinant exchange, Selmer-theoretic twists, and integral crystalline estimates that lead to the weighted prime-sum inequality. The mathematical proofs reconstructed here are distinguished from specialized results still taken from the formal 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 purpose of the present volume is to go one level deeper. It adds substantial developments of the principal coefficient, local, analytic, and congruence calculations, and introduces the common background at the places where a reader needs it. 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, complex analysis and the main classical ideas in analytic number theory. 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. The volume is not intended to be a first textbook in any one of these subjects; in particular it does not rebuild the full theories of unitary automorphic representations or crystalline cohomology from first principles. Background that is standard and applicable in the required form can be cited. Where a nonstandard integral variant is essential, the exact statement is recorded and its proof developed as far as the reconstruction permits.

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. A useful way to approach the book is to 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 chapter next. The upper-bound, spanning and lower-bound parts can then largely be studied separately; their common definitions and notational interfaces were selected precisely for this purpose. In each part a technical chapter complements the main chain of deductions.

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 The admissible data and coefficient spaces 16
3.1 The CM extension and the different . . . 16
3.2 The scales, and the reason for so many parameters . . . 17
3.3 Admissibility and construction of the data . . . 18
3.4 Derived weights, level and integral structures . . . 20
3.5 Two kinds of automorphic forms . . . 21
3.6 Precise weights and levels of the source and target . . . 21
3.7 Integral gauges and the precise uniform index . . . 23
3.8 The theta chains and their relation . . . 24
3.9 Norms, squared covolume and index . . . 25
3.10 The common leading term and the prime remainder . . . 25

4 Adelic coefficient models and integral linear algebra 27
4.1 Hermitian spaces, places and adelic lattices . . . 27
4.2 Why the extra level removes stabilizer denominators . . . 29
4.3 Normalization of class sums and change of level . . . 30
4.4 A useful elementary form of weak approximation . . . 31

4.5 Algebraic modular forms and the finite-dimensional model 32
4.6 The integral Weyl model 32
4.7 Fischer pairing and the highest-weight hook constant 34
4.8 Examples of the coefficient normalizations 35
4.9 The class-sum metric and the real matrix-coefficient model 35
4.10 Transported lattices and gauges 36
4.11 Finite-place Hecke operators on explicit lattices 38
4.12 Class correspondences and adjoints 39
4.13 Where this elementary model stops 40

II Upper bounds: CM jets, theta chains and Hecke repair 41

5 The CM differential, theta normalization and integral frame 42
5.1 An elliptic curve normalized at a point of order two 42
5.2 Theta sections and their integral size 44
5.3 Isogenies, their multipliers, and Tate's inequality 45
5.4 The integral formal parameter and the Gauss principle 46
5.5 Logarithms, norm-pair limits, and the formal height 48
5.6 A unit-adjusted Chowla–Selberg formula 50
5.7 The finite idele and the period scale 52
5.8 Constructing arithmetic tight frames of prescribed length 53
5.9 Frobenius on the frame matrix 56
5.10 Theta-group orthogonality and the finite realization 57
5.11 From an arithmetic frame to a theta section 59
5.12 Integral vacuum and the remaining reciprocity law 60
5.13 What the arithmetic normalization proves 62

6 Local CM theta analysis and Schur jets 63
6.1 The arithmetic local scales 63
6.2 The CM period and the differential coordinate 64
6.3 Integral and logarithmic theta identities 65
6.4 The regularized CM parameter at nonsplit primes 66
6.5 The exceptional exponential factor 67
6.6 Tight frames and theta vectors 68
6.7 Frobenius pairings in tight-frame products 69
6.8 Schur matrices and determinant losses 71
6.9 Application to the prescribed weights 74
6.10 From theta products to the standard integral kernel data 75

7 The definite chain, collisions and Hecke repair 78
7.1 Coefficient lattices, class spaces and index 78
7.2 The groups and the chain of theta kernels 80
7.3 The geometric seesaw and its denominator control 82

7.4 Ramification and short generators on the definite group 85
7.5 The collision loss and the marked set 86
7.6 Moving the marked Hecke operators through the chain 88
7.7 The Frobenius-twisted elementary-divisor theorem 91
7.8 Application to the actual column lattice 98
7.9 The prime N: moving denominators to a Fourier operator 99
7.10 The simultaneous local divisibility statement 101
7.11 Archimedean growth and the level normalization 102
7.12 Projection, product formula scaling and Theorem A 104

8 The local Hecke estimate in full detail, and the integral seesaw 106
8.1 The cyclic coefficient representation and the two filtrations 106
8.2 The root estimate, including its coefficient extraction 107
8.3 The dimension bound from flags and convex volume 109
8.4 Deep coefficients: a weighted Nullstellensatz and its application 113
8.5 Integral diagonal pairings and class means 115
8.6 Frobenius products of matrices and the approximate projector 117
8.7 Primitive half-depth matrices 120
8.8 The Smith pairing lemma and the local Hecke theorem 122
8.9 The finite Fourier operator at the level prime 124
8.10 Geometry and integrality in passing to a seesaw 127

III Spanning: Bessel and Poincare series 130

9 Quantitative spanning by projected Bessel–Poincaré series 131
9.1 Statement and finite-dimensional framework 131
9.2 Finite level, central projection, and the transfer bound 132
9.3 Hermitian geometry of the Bessel subgroup 134
9.4 The archimedean Bessel system 135
9.5 Construction and exact unfolding of the Poincaré map 138
9.6 Energy, pairs of congruent terms, and removal of lower cells 141
9.7 The sharp coherent main term 144
9.8 A finite-field positivity argument for singular minors 148
9.9 First moment and bad finite-type loss 150
9.10 From the moment bounds to the covering theorem 151

10 The analytic and arithmetic mechanism of the Bessel sieve 154
10.1 Norm balls in the local groups 154
10.2 A rank-one convolution estimate and its differentiated form 156
10.3 Threshold projectors and a lowest-type coefficient 158
10.4 Two-factor Laplace estimate 160
10.5 Interpolation, isotropic frames, and the Bessel peak 162
10.6 A peak-to-cut inequality with all mass factors 165

10.7 From an L^2 reproducing estimate to the bandlimit tower 167
10.8 The cell L^2 estimate and its limited differentiation cost 169
10.9 A finite-subset version of adelic Poisson summation 170
10.10 An explicit induction for the sieve system 171
10.11 Frequency summation and the product formula 173
10.12 Finite packets, their normalization, and the global comparison 174
10.13 Degenerate terms and the accumulation of energy errors 176

IV Lower bounds: periods, congruences and crystalline defects 178

11 The determinant lower bound 179
11.1 Conventions for norms of conjugate coefficient vectors 180
11.2 Heights with varying finite-place entries 181
11.3 The determinant lemma with index 182
11.4 Special-value data and fit twists 183
11.5 The correspondence of complementary rank and the finite square 185
11.6 The two-grid exchange and the crude rows 188
11.7 The triangle of weighted lifts 191
11.8 Comparison forms for one lift 194
11.9 Compressed Gram determinants 197
11.10 The height saving in a rank-one special value 199
11.11 Combining the special values and closing the lower bound 202

12 Period identities, finite matching, and the geometry of the lower bound 204
12.1 Complementary frames and matched data 204
12.2 Fourier transformations of the matching identity 207
12.3 The graph term in doubling 210
12.4 Sharp support, level factors, and the common scalar 212
12.5 A rank-one Euler factor test 214
12.6 Hyperbolic-plane removal and the signed product 215
12.7 What makes an exchange an identity on all partners? 216
12.8 Level orders in the paths 218
12.9 Why a false lift with one common root contributes no period 219
12.10 Two finite ingredients of integral comparison 221
12.11 Limits of the reconstruction 223

13 Congruences, auxiliary primes, and the fit family 224
13.1 The integral projection problem 224
13.2 Labels and Galois representations of lifted eigensystems 225
13.3 Selmer modules over finite coefficient rings 228
13.4 New level characters and tied cocycle conditions 229
13.5 Shrinking several tied Selmer modules 231
13.6 Finite presentations and interpolation of lengths 233

13.7 Degree-free selection of roots of unity . . . 234
13.8 The global twisting theorem and value families . . . 236
13.9 The Tate characters and triangular Taylor–Wiles places . . . 238
13.10 The first-order Ribet construction . . . 240
13.11 Hecke operators realizing the first-order relations . . . 243
13.12 The denominator bound from the local length budgets . . . 245
13.13 Global numerical summation and proof of the fit theorem . . . 247

14 The foundations of the congruence argument 249
14.1 From weak Hecke compatibility to a trace identity . . . 249
14.2 Tame labels and full matrix span . . . 252
14.3 A duality calculation at the mirror . . . 254
14.4 Anticyclotomic primes from Kummer extensions . . . 256
14.5 Odd laws and the triangular removal of even classes . . . 257
14.6 The algebra of adjoining relations . . . 261
14.7 Passing from scalar rows to matrix blocks . . . 262
14.8 Two applications of Cramer's rule over a family . . . 263
14.9 A transfer calculation at one auxiliary level . . . 264
14.10 Separation of the arithmetic and algebraic budgets . . . 267

15 Local crystalline defects and the integral cup-product estimate 268
15.1 Lattices, crystalline cocycles and imperfectly dual reductions . . . 268
15.2 Why this defect bounds the local discrepancy of congruent cocycles . . . 271
15.3 The Kummer chart, letter decomposition and height of a frame . . . 272
15.4 Exponent lattices and the cost of increasing depth . . . 274
15.5 Cokernel estimates: divisibility by p and by the Eisenstein element . . . 276
15.6 Cup-product vanishing from the three sides . . . 277
15.7 The window estimates: the three ways to obtain safe points . . . 280
15.8 The tame type bound, uniform in the ramification degree . . . 283
15.9 Case analysis for the primes and the explicit local loss . . . 284
15.10 Insertion of the global weights and the sum over all blocks . . . 285

16 Integral period frames and the algebra of the local loss 288
16.1 The rank calculation and duality over a finite coefficient ring . . . 288
16.2 The coordinates in a Kummer chart . . . 290
16.3 A coefficient retraction and why it is used . . . 292
16.4 From crystalline periods to an integral frame . . . 293
16.5 The radius estimate and its fourth power . . . 296
16.6 Exponent lattices, Frobenius letters and the small part . . . 298
16.7 Exact cokernel multipliers via finite differences . . . 300
16.8 A rank-sensitive bound for raising filtered conditions . . . 302

V The arithmetic conclusion and proof audit 306

17 A numerical appendix: asymptotic budgets and the passage to the limit 307
17.1 Uniform scale asymptotics . . . 307
17.2 Rounding of large coefficients and a Lipschitz estimate . . . 308
17.3 The exact level cancellation revisited . . . 308
17.4 An elementary bound on the logarithm of the class number . . . 309
17.5 Summation of the fit-family costs . . . 310
17.6 The final numerical inequalities . . . 311

18 The arithmetic inequality and conclusion 312
18.1 Collecting the arithmetic inputs . . . 312
18.2 From A and B to an almost-basis of short forms . . . 313
18.3 Comparison with the lower bound . . . 313
18.4 Extracting the analytic remainder . . . 314
18.5 Choosing the second cutoff and the power saving . . . 315
18.6 GRH . . . 315

19 Glossary and concordance 317
19.1 Reading the main theorem in the Lean source . . . 318

20 Mathematical status and proof dependencies 320
20.1 What has changed at monograph level . . . 320
20.2 Priority for the remaining reconstruction . . . 322
20.3 The use of published mathematics . . . 322

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 3. 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

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 classical prime number theorem for a fixed arithmetic progression; 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

The admissible data and coefficient spaces

3.1 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^+. It 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

$$|\tau(\varpi)| = \sqrt{d}|P'(\tau\zeta)|, \qquad \mathrm{rd}_E^f = \prod_{\tau\in\Phi} |\tau(\varpi)|, \tag{3.1}$$

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,

$$x^{10} \le |\tau(\varpi)|, \mathrm{rd}_E \le x^{12}. \tag{3.2}$$

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 retain exactly those choices of the compact source, since different choices of theta normalizations can change the period and volume identities used below.

3.2 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 $\log b$.

3.3 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

$$10^{12} c_* \le \log Y, \quad 100 \log Y \le \log x, \quad 10^6 \log x \le U_1, \quad 10^6 U_1 \le Y, \quad 6 \log U_1 \le \log Y. \tag{3.3}$$

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

$$\sigma, \tau, \rho \in \mathcal{S}, \quad \phi(\tau/\sigma)^j = \phi(\rho/\sigma)^i \implies \phi(\tau/\sigma) = 1. \tag{3.4}$$

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

$$\forall U \in [100Y, 2^s] \quad \exists p \text{ prime}, \quad p \equiv 1 \pmod{M_\omega^2}, \quad U \le \log p \le U + \sqrt{Y}. \tag{3.5}$$

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

Lemma 3.1 (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 3.2 (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*

$$Y = x^{\varepsilon}, \qquad U_1 = x^{\varepsilon/6} \ge (10^6 \log x)^2, \qquad |\mathcal{S}| \ge \frac{T_s}{Y}. \tag{3.6}$$

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 (3.6) satisfies (3.3): 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 (3.2).

To choose $\mathcal{S}$, apply Lemma 3.1 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. Here a standard theorem replaces a substantial segment of the formal background: the Siegel–Walfisz theorem, in the usual form uniform for moduli

at most a fixed power of the logarithm of the argument ([2, Cor. 11.19]). We use exactly the following consequence. For each fixed $A > 0$, if $Q_0 \le Y^A$, $U \ge 10Y$ and $Y \to \infty$, then there is a prime $p \equiv 1 \pmod{Q_0}$ with $U \le \log p \le U + \sqrt{Y}$, uniformly in U. Indeed, with $v = e^{U+\sqrt{Y}}$, $u = e^U$, the two asymptotics for the numbers of primes in the class 1 mod Q_0 have main terms $\mathrm{Li}(v)/\varphi(Q_0)$ and $\mathrm{Li}(u)/\varphi(Q_0)$. Their ratio tends to infinity uniformly, while their relative errors are $o(1)$, by choosing an exponent in Siegel–Walfisz larger than $A + 2$. This applies because $Q_0 = M_\omega^2 = \exp(O_{N_0,d}(s)) \le Y^A$ for a fixed A depending on ε, N_0, d. It establishes (3.5), 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 (3.3), 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 a classical uniform prime-distribution theorem, rather than reproducing the source's weaker prime-window lemma `exists_prime _modEq_one_log_mem`, in `Lib/Analytic/Primes/AuxiliaryPrimes.lean`. The latter avoids invoking a quantitative prime number theorem for progressions. It is not the same analytic use as the pointwise prime number theorem at the second cutoff: the latter is avoided altogether by the window lemma.

3.4 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

$$w_{\mathrm{nom}}(\sigma) \le w_j(\sigma) \le w_{\mathrm{nom}}(\sigma) + x^{42}, \quad w_j(\sigma) \le 2b, \quad \log w_j(\sigma) \le \log w_{\mathrm{nom}}(\sigma) + 1/x. \tag{3.7}$$

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

$$\lambda_N = m(w_* + G_*) - G_* - \tfrac{15}{2}\log b, \quad \lambda_N \le \log N \le \lambda_N + \sqrt{\log b}, \quad N \equiv 1 \pmod{M_\omega^2}. \tag{3.8}$$

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 (3.5), 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.

3.5 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.

3.6 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 $e_j(\tau) = \lfloor n^X_{j-1}/2 \rfloor - w_j(\tau)$ for $j \geq 1$. 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/SchurM odule/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/Cus pForms/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 9 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 $i = 0, \dots, r-1$, 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 7).

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.

3.7 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 $\sum_{p\in S} e_p \le d_{\mathrm{ind}}$ 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

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

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.

3.8 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.

3.9 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

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

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 3.7.

3.10 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

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

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:

$$\mathcal{R} = a\{|\mathcal{S}|(R_{N_0}(x) - R_{N_0}(x_1)) + |\mathcal{T}|R_1(x_1)\}. \tag{3.9}$$

This is `Data.remainder_eq`. We explain its mechanism in Chapter 18: 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 4

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 13 and 15.

The reader familiar with algebraic modular forms may skip most of this chapter, but should retain the normalization of the coefficient pairing in (4.6), the freeness criterion 4.2, and the distinction between the Weyl lattice and the lattice of polynomials in Lemma 4.7. 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.

4.1 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

$$A \longmapsto (A, A^{-T}), \qquad A \in \mathrm{GL}_l(F_v). \tag{4.1}$$

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 4.1 (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.

4.2 Why the extra level removes stabilizer denominators

Although every stabilizer in Proposition 4.1 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 4.2 (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*

$$(2l)^{[E:\mathbb{Q}]}<\prod_{i=1}^{a}\#k_{w_i}, \tag{4.2}$$

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 $z=l-\mathrm{tr}\,\gamma\in\mathcal{O}_E$. 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, $\Re(\mathrm{tr}\,\gamma)=l$ 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 (4.2). □

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 (4.2) 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 4.3 (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.

4.3 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 4.4 (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.

4.4 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 4.5 (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.

4.5 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 4.6 (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 4.4 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 7). 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.

4.6 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

$$S_\kappa(A)\,\mathrm{ev}_X = \mathrm{ev}_{AX}\,. \tag{4.3}$$

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 (4.3) yields the contravariance identity

$$G_\kappa S_\kappa(A) = S_\kappa(A^T)^T G_\kappa. \tag{4.4}$$

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.

4.7 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 $\big(\mathsf{F}(b_i, b_j)\big)_{ij}$ 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 4.7 (The two integral pairings)**.** *One has*

$$G_\kappa H_\kappa = H_\kappa G_\kappa = h_\kappa I. \tag{4.5}$$

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 (4.4), 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`.

4.8 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.

4.9 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,

$$\langle \phi, \psi \rangle_{K'} = \sum_{[g]} \frac{\overline{\phi(g)}^{T} G_\kappa \psi(g)}{|\Gamma_g|} \tag{4.6}$$

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 11 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 9).

4.10 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 4.8 (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 4.5.*

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 4.9 (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 7.

4.11 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 (4.1) 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 7. 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.

4.12 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 4.4 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.

4.13 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 9); the last takes up Chapters 7, 11, and 13. The elementary computations in this chapter isolate the exact places where their more advanced statements enter.

Part II

Upper bounds: CM jets, theta chains and Hecke repair

Chapter 5

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.

5.1 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

$$Y^2 + a_1 XY = X^3 + X. \tag{5.1}$$

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

$$X' = 2Y + a_1X, \qquad Y' = 1 + 3X^2 - a_1Y, \qquad X'' = 6X^2 + a_1^2X + 2. \tag{5.2}$$

Proposition 5.1 (Unit discriminant in the chosen normalization)**.** *For every prime p and embedding of coefficients into an algebraically closed p-adic field, the normalized cubic* (5.1) *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 (5.1) 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_a₁_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.

5.2 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

$$(U_y^h F)(z) = \exp(-h\pi(z\bar{y} + |y|^2/2)/A)F(z+y). \tag{5.3}$$

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,

$$\vartheta(z+w)\vartheta(z-w) = \vartheta(z)^2(\vartheta(w)^2 - X(z)\sigma(w)^2). \tag{5.4}$$

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

$$\vartheta(2z) = \vartheta(z)^4 - \sigma(z)^4, \tag{5.5}$$

$$\sigma(2z)^2 = \sigma(z)^2\vartheta(z)^2(4\vartheta(z)^4 + a_1^2\sigma(z)^2\vartheta(z)^2 + 4\sigma(z)^4), \tag{5.6}$$

$$G^2 + a_1\sigma^2\vartheta G = \sigma^6 + \sigma^2\vartheta^4. \tag{5.7}$$

For the first, specialize (5.4) 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 (5.1) 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 5.2 (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 (5.7) 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 5.1. 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 (5.5) and (5.6), 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 5.3 (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 (5.1), 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.

5.3 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 5.4 (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*

$$|(U_y^h F)(az)|_p \le H_p(z)^{hn} \quad (z \in \mathbb{C}). \tag{5.8}$$

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 5.3 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 5.2 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 (5.8). □

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.

5.4 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 (5.2) 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 6.3.

Lemma 5.5 (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 5.6 (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 5.3. 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 5.1 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 5.5 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 (5.1) 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 5.4 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 5.7 (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 (5.3) identifies it with the function in the statement. Proposition 5.4 bounds its size by H_p^n. Apply Proposition 5.6. □

This proves the assertion (6.5) 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.

5.5 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 (5.4), 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

$$\mathscr{L}'' = L_2 + X, \qquad L_2 = -\vartheta''(0). \tag{5.9}$$

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 (6.3), 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 6.3 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 5.8 (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 6 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 5.9 (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 (5.1) is a good Weierstrass model by Proposition 5.1; 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 6 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 5.10 (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

$$(X^nX')' = (4n+6)X^{n+2} + (n+1)a_1^2X^{n+1} + (4n+2)X^n. \tag{5.10}$$

Indeed use $(X')^2 = 4X^3 + a_1^2X^2 + 4X$ and (5.2). 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 (5.9) 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 5.10 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.

5.6 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 5.6 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

$$\prod_c A_c^\gamma = (A_C^\gamma)^{h_K}|\gamma(u_*)|^2. \tag{5.11}$$

The curves C_c^γ have the same underlying ideal-class family, with their prescribed individual differential scalings. In particular the product formula (5.11) 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

$$(s-1)\sum_{l\in L\setminus\{0\}} \left(|l|^2/(k_0A_L)\right)^{-s} = \pi k_0 + \pi k_0(s-1)\left(2\gamma_{\mathrm{E}} - \log\frac{A_L}{k_0\pi^2} - \tfrac{1}{6}\log|\Delta_L|\right) + O((s-1)^2). \tag{5.12}$$

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 (5.12). 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 (5.12) 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

$$\exp\left(\frac{1}{h_K}\sum_c \log A_c^\gamma + \frac{1}{6h_K}\sum_c \log|\gamma(\Delta_c)|\right) = \frac{2\pi^2}{\sqrt{d}}\exp(\kappa_K). \tag{5.13}$$

Proposition 5.11 (CM normalization by a single algebraic unit)**.** *There exists an algebraic unit u_0 such that for every $\gamma \in \mathrm{Aut}(\mathbb{C}/K)$,*

$$A_C^\gamma|\gamma(u_0)|^2 = \frac{2\pi^2}{\sqrt{d}}\exp(\kappa_K). \tag{5.14}$$

Proof. By Proposition 5.1, 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 (5.11), the expression on the left in (5.13) 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 (5.14). □

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.

5.7 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

$$\chi_{\mathrm{fin}}(y) = \prod_{\tau \in \Phi} \tau(\bar{y}). \tag{5.15}$$

This is the type of the normalized finite character, not the unitary Hecke character prior to the absolute-value factor. The existence of the compatible splitting character and its use in the Weil representation is a formal-source input here; 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 (5.15) 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 (5.15) 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

$$|\iota(\chi_{\mathrm{fin}}(z))|_p = \prod_{\tau\in\Phi} |z|_{\iota\tau\circ\mathrm{conj}}. \tag{5.16}$$

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 (5.15), the squared ordinary absolute value of $\chi_{\mathrm{fin}}(\alpha)^h$ is this rational ideal norm. Since the value on the right of (5.15) is in K, any field automorphism γ over K fixes it. Thus, taking positive roots,

$$|\gamma(s_0)|^2 = \mathrm{rd}_E/\sqrt{d}. \tag{5.17}$$

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 (5.16) 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 5.11 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 $y\in\mathcal{I}$,

$$|\iota(s_1\overline{\tau(y)})|_p \le 1, \qquad |\iota(s_1)|_p = 1 \quad (p\nmid N_1). \tag{5.18}$$

This is exactly the finite normalization required in Proposition 6.2, 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 (5.17) and (5.14) give

$$|t_\gamma|^2 = \frac{\mathrm{rd}_E}{2\pi e^{\kappa_K}} = \Omega^{-1}. \tag{5.19}$$

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.

5.8 Constructing arithmetic tight frames of prescribed length

We now prove the arithmetic frame assertion of Proposition 6.10. 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

$$\sum_i m_i \sigma(w_i)\overline{\tau(w_i)} = M\mathbf{1}_{\sigma=\tau} \quad (\sigma, \tau \in \Phi). \tag{5.20}$$

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 5.12 (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 $0 \le k, i < \ell$:

$$w_{\ell,j,k,i} = \zeta^j \left(\mathbf{1}_{k=0} A_\ell^i + \frac{b_0 M_\ell}{N} (A_\ell^i - A_\ell^k) \right). \tag{5.21}$$

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 5.13 (Length of the modified cyclotomic frame)**.** *The family* (5.21)*, with multiplicities one, is a tight frame of length*

$$C_\ell = N + M_\ell(1 + 2n_0 z_0 + 2\ell n_0 M_\ell e). \tag{5.22}$$

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 (5.21) 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 (5.22).

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 5.12. Equation (5.18) 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 5.14 (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 (5.18), and the vector in Lemma 5.13 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.

5.9 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 5.15 (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 5.16. *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 5.15.

Proof. Factor out β^{p^e+1} and apply Lemma 5.15 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 (6.13). This proves Lemma 6.12, 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.

5.10 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 (5.3)). 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 5.17 (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,

$$m \sum_{x \in D} \ell(R_x v) R_{-x} F = |D| \ell(F) v. \tag{5.23}$$

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 (5.23). 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

$$\Theta\big(\varrho(\operatorname{finImage}\gamma)F\big) = U_{-\operatorname{realImage}\gamma}(\Theta F). \tag{5.24}$$

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 (5.24). 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 (5.24). By Lemma 5.17 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 6.11. Transport (5.23) 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.

5.11 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

$$A = \frac{|a|^2}{D_\delta} M \frac{\sqrt{d}}{2}. \tag{5.25}$$

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 (5.25). 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 (5.24) 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 6.11.

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 (5.25). 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.

5.12 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. 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 (5.25) 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 (5.19). 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 $R_x = \varrho(\alpha x)$,

$$\langle v, R_x v\rangle \in \{0\} \cup \mu_\infty. \tag{5.26}$$

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 5.18 (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 (5.26) 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

$$N^{4fh}C_1^{N^2h} \le \exp((4+\log C_1)N^2h),$$

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 $s_{\mathrm{vac}}^h v$ 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 6.20 and Proposition 6.22. 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`.

5.13 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 6

Local CM theta analysis and Schur jets

The preceding chapter (Chapter 5) supplies expanded proofs of the CM arithmetic, local section identities, frame and Frobenius calculations used here. The statements retained as interfaces in this chapter are understood with the derivations and residual hypotheses in that chapter.

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.

6.1 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

$$A_p(y, \sigma) = \omega_p \sum_{e \text{ live}} \left(y/p^e - w_{\mathrm{nom}}(\mathfrak{p}_e(\sigma))\right)_+, \qquad 0 \le y \le a \tag{6.1}$$

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 6.1. *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

$$F(p) = \sum_\sigma A_p(w_{\mathrm{nom}}(\sigma),\sigma), \qquad g_p = m\left(\sum_{\tau\in\mathcal{T}} A_p(a,\tau) + kF(p)\right). \tag{6.2}$$

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 (6.1).

6.2 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 6.2 (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 5, retaining the precise normalized Weil-character input.

6.3 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 6.3 (Normalized CM identities (see Chapter 5))**.** *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*

$$\mathscr{L}''(z) = L_2 + \xi(z), \qquad |\iota(\sqrt{-d})L_2|_p \le 1. \tag{6.3}$$

If (c, D) is in the closure of the pairs $(\iota\delta, \iota(\delta\bar{\delta}))$ for $\delta \in \mathcal{O}_K$, then

$$\phi(cz), \quad \exp(D\mathscr{L}(z) - \mathscr{L}(cz)), \quad (D\mathscr{L}(z) - \mathscr{L}(cz))' \ \in \mathfrak{R}_\phi. \tag{6.4}$$

If $0 \neq \delta \in \mathcal{O}_K$, $\delta y \in \Lambda$ and $D = \delta\bar{\delta}$, the twisted theta translate satisfies

$$(U_y\vartheta)(\iota(\bar{\delta})z) = \vartheta(z)^D V(z), \qquad V \in \mathfrak{R}_\phi. \tag{6.5}$$

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 5.

Lemma 6.4 (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 (6.5), 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.

6.4 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 6.3 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 (6.4) 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

$$Z'(u) \in \mathfrak{o}[[u]], \tag{6.6}$$

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 6.5 (Degree of the reduced isogeny (see Proposition 5.9))**.** *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

$$|\varpi|_p = p^{-\rho_p}, \qquad |\pi|_p = p^{-\omega_p}. \tag{6.7}$$

Lemma 6.6. *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,

$$T(\varpi z) = \varpi S(z), \qquad S - S^{q_p} = z + \pi E_0(z). \tag{6.8}$$

Proof. The first two nonzero coefficients are as stated by the choice of ϖ. In degree $n = 1 + Q_0 j$, $j \geq 2$, (6.6) 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.

6.5 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 6.7. *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

$$1 + \omega_p I - v_p(N) - v_p(i) + \omega_p \mathbf{1}_{j=1}. \tag{6.9}$$

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 (6.9) 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 (6.9) 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 (6.8), 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 (6.3), $v_p(c_2) \geq -1/2$ (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

$$\Psi_0(z) = u_0 z^2 + u_1 S(z)^{p+1} + u_2 S(z)^{p+q_p} + \pi R_0(z), \tag{6.10}$$

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 (6.3) 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 6.8 (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 (6.4) 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 (6.5) 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 (6.8), (6.10) and the proposition above give the nonsplit part of `CMLattice.exists_formal_parameter` in `New/Curve/FormalParameter.lean`.

6.6 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 6.9. 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 6.2.

Proposition 6.10 (Large frames with a unit vector (see Proposition 5.14))**.** *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 5.14 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

$$[z^\beta]\mathcal{T}_c F = \frac{c^{|\beta|}}{\beta!}\vartheta_V(v_\beta \otimes F) \qquad (\beta \in \mathbb{N}^{\Phi\times\{1,\dots,h\}}), \tag{6.11}$$

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 6.9. 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 6.11 (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.

6.7 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 6.8, after substituting ϖz the product has the form

$$F(\varpi z)=\exp(\kappa_n\Psi(z))V(z), \quad V\in\mathfrak{o}[[z]], \quad \Psi(0)=0, \tag{6.12}$$

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 6.2 gives $r_{\tau i}\in\mathfrak{o}$.

Lemma 6.12 (Frobenius orthogonality (see Proposition 5.16))**.** *In characteristic p, modulo $\pi\mathfrak{o}$, the frame matrix satisfies, for any live exponent e and for $e=0$ when $p=d$,*

$$\sum_i r_{\sigma i}^{p^e}r_{\tau i}=(N\iota(b'))^{p^e+1}M\,\mathbf{1}_{\tau=\mathfrak{p}_e(\sigma)}. \tag{6.13}$$

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

$$\sum_{\tau,c}Z_{\tau,c}^{p^e}Z_{\mathfrak{p}_e(\tau),c}=0 \quad (e\in I_D) \tag{6.14}$$

hold after reduction modulo π and terms of degree $>D$. We shall prove they kill the exceptional exponent in (6.12) and render the exponential integral through degree D.

Lemma 6.13 (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 6.14 (Null-cone integrality)**.** *Whenever the integral zero-constant series $Z_{\tau,c}$ satisfy (6.14), 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 $t_{ic} = S(L_{ic}(Z))$, (6.8) gives $t_{ic} - t_{ic}^{p^e} = L_{ic}(Z)$, and t_{ic} is nilpotent by degree. Using (6.13) 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 6.13 shows that the two S-powers in (6.10) 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 (6.12) 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 (6.10), $u_0 = 0$ 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.

6.8 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 6.15 (Formal-source input: integral Schur expansion)**.** *For a weight satisfying the rank condition `FitsRanks` the Schur matrix $S_\kappa(X)$ satisfies*

$$b_{i'}(XY) = \sum_i S_{i'i}(X) b_i(Y) \prod_\sigma \det(Y_\sigma)^{d_\sigma}, \tag{6.15}$$

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

$$\mathfrak{C}_{p,e}(\sigma,\tau) = \min(a_\sigma^{ev}/p^e, c_\tau + a_\tau^{ev}) + c_\sigma/p^e + \min(D_\sigma/p^e, D_\tau). \tag{6.16}$$

This is `couplingCount` in `New/Curve/LiveSum.lean`. Its three terms distinguish the three types of pairs of columns.

Lemma 6.16 (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

$$\prod_{\sigma,h} |\det M_{\sigma,h}|_p^{\kappa_\sigma(h)} \geq p^{-\omega_p \sum_{e\in I_D}\sum_\sigma \mathfrak{C}_{p,e}(\sigma,\mathfrak{p}_e(\sigma))}. \tag{6.17}$$

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 (6.17). □

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 6.17 (Schur-jet bound)**.** *Suppose F satisfies the null-cone property of Proposition 6.14 for degree $D > 0$ and the scales (6.7). Then*

$$\begin{aligned} v_p(\mathscr{J}_\kappa(F)_{i'i}) \geq (1/(p-1) - \rho_p)D - \omega_p \sum_{e \in I_D} \sum_{\sigma} \mathfrak{C}_{p,e}(\sigma, \mathfrak{p}_e(\sigma)) \\ - fN^3(\lfloor \log_p D \rfloor + 1) - G_{i'}. \end{aligned} \tag{6.18}$$

Proof. Apply the jet functional in X to (6.15). 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 6.16. 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 (6.17). This proves (6.18). □

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 (6.16) are essential: replacing them by the more expensive alternatives would eliminate the positive-part terms in (6.1).

6.9 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 6.18 (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

$$E_{p,j}(ev) = \left(\frac{1}{p-1} - \rho_p\right)D_j - \omega_p\sum_{e\in I_{D_j}}\sum_\sigma \mathfrak{C}_{p,e}(\sigma,\mathfrak{p}_e(\sigma)) - fR_j^3(\lfloor\log_p D_j\rfloor + 1). \tag{6.19}$$

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 6.19 (Prime-power gain of one jet)**.** *For nonsplit $p < x_1$ and $0 \le j < m$,*

$$E_{p,j}(ev) \ge \frac{g_p}{m} - \mathrm{coll}_p(ev) + \omega_p|\mathcal{T} \cap ev|\, a/x_1 - \frac{fa(\log b)^2}{pm}. \tag{6.20}$$

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 6.18. 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 6.1. 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 (6.2). □

The exact uniform source declaration is `Data.le_jetExponent`, in `New/Columns/JetExponentBound.lean`. Formula (6.20) 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 (6.20) compensates part of the charge incurred in doing so.

6.10 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 6.20 (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 $\exp(CN_1^2h)$ 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 6.21 (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 6.10. Repeat its vectors so that all multiplicities are one. Use Proposition 6.11 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 6.14, Theorem 6.17, and Lemma 6.21 bound the Schur jets on the entire module whenever the scalar multiplier has integral norm. For general $\delta \neq 0$, Lemma 6.4 applied using Proposition 6.2 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 6.22 (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 (6.18), with any extra weights charged box by box. This completes the local differential calculation needed for the Hecke and seesaw argument.

Chapter 7

The definite chain, collisions and Hecke repair

The purely algebraic local Hecke estimates, including the Frobenius root and coefficient arguments, are developed in Chapter 8. Reading them separately is useful: they 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

$$g_p = m \left(\sum_{\tau \in \mathcal{T}} A_p(a, \tau) + k \sum_{\tau \in \Phi} A_p(w(\tau), \tau) \right). \tag{7.1}$$

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.

7.1 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 7.1. 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 7.2 (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.

7.2 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

$$(\mathsf{R}v)(u) = \left(\prod_{\tau\in\Phi} \tau(\bar\pi)^{-e_{\mathrm{br}}(\tau)} \right) D_{\mathrm{br}} v(\mathrm{diag}(u,1)). \tag{7.2}$$

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 (7.2) 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:

$$\|\sigma\mathsf{R}v(u)\|^2 \le \exp\left(x^{70} - \textstyle\sum_\tau e_{\mathrm{br}}(\tau)\log N\right) \|\sigma v(\mathrm{diag}(u,1))\|^2. \tag{7.3}$$

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 7.3 (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.

7.3 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 7.4 (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 7.5 (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_reducedVacuum`, `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 7.4.

7.4 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 7.6. *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:

$$\operatorname{span}_{\mathbb{C}}\{u \in C^{\mathrm{ram}} : u \text{ is everywhere integral, } \|\sigma u(g')\|^2 \leq e^B \ \forall \sigma, g'\} = C^{\mathrm{ram}}. \tag{7.4}$$

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 (7.4) with the ordinary complex norm squared, not its logarithm, bounded by e^B.

7.5 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 (7.1) 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 $o_0 \le s_p$. These are the definitions `Data.IsShortPrime`, `Data.loss`, `Data.ceilLogM`, `Data.deltaP`, `Data.IsServed` and `Data.collisionSum` in `New/Curve/LiveSum.lean`.

Lemma 7.7 (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:

$$m \mid \mathrm{ord}(u). \tag{7.5}$$

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

$$l(u^e\tau) = \overline{l(\tau)} \quad (e \text{ odd}), \qquad l(\tau') \in \{l(\tau), \overline{l(\tau)}\} \iff \tau' = u^e\tau \text{ for some } e. \tag{7.6}$$

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 (7.5) 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 7.8 (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 (7.6). 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 (7.6) 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

$$\sum_{\tau} G(w_j(\tau), l(\tau)) = \sum_{\tau} G(w_j(\tau), \overline{l(\tau)}). \tag{7.7}$$

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.

7.6 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 7.9 (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 (7.7) 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 (7.7). 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 7.10 (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 7.9(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 7.3 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

$$J_{p,j}(E_0) \ge \frac{g_p}{m} - C_p(E_0) + \omega_p \#(\mathcal{T}\cap E_0)\frac{a}{x_1} - \frac{fa(\log b)^2}{pm}. \tag{7.8}$$

For marked primes the other bound required is

$$|\Lambda|\delta_p + s_p \sum_{\tau\in E_0(\Lambda)} L_j(\tau) \le \frac{\ell_p}{m} + \#(\mathcal{T}\cap E_0(\Lambda))\frac{a}{x_1}. \tag{7.9}$$

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 (7.9). 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 (7.9) could not simply be discarded: for p near x_1 it would be too large. It is exactly matched by the positive term in (7.8).

For a full marking, $C_p(E_0(\Lambda)) = 0$ and $\omega_p = 1$. Combining (7.8), (7.9) and Lemma 7.10 gives

$$\|\mathsf{B}_\theta^\Psi(Y^m v)\|_\iota \le p^{2\ell_p - g_p}. \tag{7.10}$$

There is also a useful estimate with no marking at all. One may take $\Lambda = \varnothing$ in Lemma 7.10, with no nontrivial operators. Then (7.8) gives

$$\|\mathsf{B}_\theta^\Psi(v)\|_\iota \le p^{mC_p(\varnothing) + \ell_p - g_p}. \tag{7.11}$$

Writing

$$j_p = \left\lceil \frac{2mfa}{p} \right\rceil,$$

the estimate $C_p(\varnothing) \le 2fa/p$ and scalar linearity imply

$$\|\mathsf{B}_\theta^\Psi(p^{j_p} v)\|_\iota \le p^{\ell_p - g_p}. \tag{7.12}$$

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.

7.7 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

$$\begin{aligned} &0 < 2r_p \le k, \quad \epsilon \le 1 \le y, \\ &\nu a k(\lambda+2) + a + k < P_0^\lambda, \quad 6c_* + 2 \le y, \\ &2(\nu+2)k/y \le \epsilon, \\ &(y + \nu(k+1))kP_0^{3r_p} + k + 2 \le a/y, \\ &\nu a\left(P_0^{-2r_p} + \frac{r_p^2}{kP_0}\right) \le z. \end{aligned} \tag{7.13}$$

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

$$\mathrm{exc}_i(\alpha; b, b') = \frac{1}{Q_l} \sum_s \left(P_0^{1+[s-1-i]_\nu} \alpha(y_s) - P_0^{[s-i]_\nu} \alpha(z_s) \right), \tag{7.14}$$

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 7.11 (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 8.1, `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

$$[\mathrm{GL}_k(R_l) : \mathcal{P}_{\mathrm{lev}}] \dim U_y \le (1+\epsilon) D_{\mathcal{P}}, \tag{7.15}$$

$$\sum_{0 \le j < J} \dim V_j \le (1+z) D_{\mathcal{P}} \quad (J \ge 0). \tag{7.16}$$

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 8, 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 (7.14), 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

$$\frac{\mathrm{vol}(\Sigma(D))}{\displaystyle\prod_{j=1}^{k} \frac{Q_l^j - 1}{Q_l - 1}}. \tag{7.17}$$

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 $\beta_b \geq Q_l\beta_{b+1}$. 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 (7.17). 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 (7.15) 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 $u \in \{0, \ldots, D-1\}^{k\nu}$. 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 8.3, 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 (7.13). 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 (7.17). The exact inequalities including the roundings are those of `index_mul_finrank_excessCond_le`. They supply (7.15).

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 8.6, `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

$$\frac{a}{k}\frac{\nu}{2}\left(2r_pP_0^{-1}+(k-2r_p)P_0^{-2r_p}\right)\le \nu a\left(P_0^{-2r_p}+\frac{r_p^2}{kP_0}\right)\le z.$$

This proves (7.16), using only the lower bound on cost in V_j and the last size condition in (7.13).

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 8.9, 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 8.10, 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 8.8, 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

$$(z + \epsilon s + y)|X||I| \quad \text{modulo } p^s \tag{7.18}$$

when (7.13) holds. Its algebraic proof is in Section 8.8; 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 (7.15). 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 $Qu - Mv = a_0 wF$ 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 (7.15) and (7.16) now account for the three terms ϵs, z, and y in (7.18): 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 7.2. 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.

7.8 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 (7.18), 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 7.12 (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 (7.13). 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 (7.13),

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 7.13 (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 7.2 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

$$\frac{2m^2 f}{x^{14}(\log b)^2} + \frac{mp}{ax^{14}(\log b)^2},$$

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 $m^2 \leq 4x$, and the third using $p < x_1$ and $x_1^{10^6} \leq b$. 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.

7.9 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 (7.3). 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 8.14, 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 7.3,

$$\|\mathsf{B}_\theta^\Psi(v)\|_\iota \le N^{3mfb} \tag{7.19}$$

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 7.2 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 (7.19) 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.

7.10 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 $\|\mathsf{B}_\theta^\Psi v\|_\iota \le 1$ is `Data.IsGoodDatum.gaugeLe_one` in `New/Columns/LocalStatement.lean`. It is the application of the crude jet bound and Lemma 7.3 to the exact unmodified data.

Proposition 7.14 (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 $e \ge 0$ such that*

$$I(M, M'_p; p^e), \qquad e \le \epsilon_p c.$$

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 7.13. The two terms have gauges at most $p^{2\ell_p - g_p}$ and $p^{\ell_p - g_p}$ respectively by (7.10) and (7.12); the same bound $p^{2\ell_p - g_p}$ holds for their sum. Since $\ell_p \ge 1$, the condition on β gives $\beta + 2\ell_p - g_p \le 0$, 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 7.7, and (7.11). At N use the preceding frame-change index and (7.19); the condition $\beta < 1 - 4mfb$ ensures $\beta + 3mfb \le 0$ since $mfb \ge 1$. At the other primes integrality is unconditional and the hypothesis $\beta < 1$ gives $\beta \le 0$. 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:

$$\sum_p \epsilon_p \log p \le B. \tag{7.20}$$

Here is a direct calculation, with generous constants for the scale comparison. If $p \equiv 1 \pmod m$ is prime, $p = um + 1$ with $u \ge 1$, so $m/p \le 1/u$. If $p \equiv -1 \pmod m$, then $p = um - 1$ with $u \ge 1$, so $m/p \le 2/u$ since $m \ge 2$. Dropping the primality conditions and bounding $\log p$ by $\log x_1$, one obtains

$$\sum_{p<x_1} \frac{\log p}{p} \begin{cases} m, & p \text{ short}, \\ 1, & \text{otherwise} \end{cases} \le 4 \log x_1 (1 + \log x_1).$$

Thus the small-prime contribution is bounded by $12fa(\log b)^2 \log x_1(1 + \log x_1)$. Since $(\log x_1)^6 \le \log b$ and $\log b$ is large, this is less than $B/2$. At N, one has $4mfb \log N \le 120fm^2 b \log b \le 480fa \log b$, using $m^2 \le 4x$ and $a \ge xb$. 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.

7.11 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

$$h_0 + 2 \sum_p g_p \log p = m \left(\sum_{\tau \in \mathcal{T}} J(a) + k \sum_{\tau \in \Phi} J(w_0(\tau)) \right) - \#\mathcal{T}\, a \log N. \tag{7.21}$$

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

$$\|\sigma \mathsf{B}_1^\Psi v(g')\|^2 \le \exp\left(\eta + h_0 + 2 \sum_p g_p \log p + B \right). \tag{7.22}$$

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 `diagramWeigh t_le_by_index` checks the corresponding inequalities with the fixed scale constants. The other terms in the complex triangle inequality of Lemma 7.3 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 (7.3), 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 (7.21), plus at most B of error. The exact inequality (7.22), 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 $a \ge 1$ is

$$\log(a!) \le a\log a - a + 1 + \tfrac{1}{2}\log a.$$

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

$$J(a) - M_0 \le a(\overline{w} + G_0) + 1 + \tfrac{1}{2}\log a.$$

By construction of the level window

$$\log N \ge m(\overline{w} + G_0) - G_0 - \tfrac{15}{2}\log b.$$

Consequently

$$\begin{aligned} h_0 &= H - m\mathcal{R} + \#\mathcal{T}\{m(J(a) - M_0) - a\log N\} \\ &\le H - m\mathcal{R} + \#\mathcal{T}\left\{m(1 + \tfrac{1}{2}\log a) + a(G_0 + \tfrac{15}{2}\log b)\right\} \\ &\le H - m\mathcal{R} + B. \end{aligned} \tag{7.23}$$

The last estimate uses $\#\mathcal{T} \le f$, $G_0 = O(\log b)$ and $m \le 2\sqrt{x} \ll a$, 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.

7.12 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 7.4. 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 7.15 (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 (7.4). 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 7.14. Therefore on M'_p the factor q times the restricted lift is integral for every admissible twist prime to p. The upper bound (7.22) 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 7.14 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 7.2 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 (7.20) 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 8

The local Hecke estimate in full detail, and the integral seesaw

The exceptional-prime correction of Chapter 7 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.

8.1 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

$$E_i(\alpha; b, b') = \frac{1}{q}\left(\sum_s P^{1+[s-1-i]_\nu}\alpha(y_s) - \sum_s P^{[s-i]_\nu}\alpha(z_s)\right), \tag{8.1}$$

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`.

8.2 The root estimate, including its coefficient extraction

The short argument using multiplicity of zeros in Chapter 7 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 8.1 (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

$$\Phi \neq 0, \qquad \deg \Phi \leq M. \tag{8.2}$$

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 $c(z_s)+\omega qP^{t_s} \geq c(y_s)$, 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 (8.2) 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 (8.1) 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 8.2. 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.

8.3 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 8.3 (Amplification by a full Frobenius period)**.**

$$\dim U_h \le F_L \operatorname{vol}(\mathscr{C}_{h+\nu(k+1)} \cap \Sigma(n_0 + k + 1)).$$

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 $0 \le v_{bs} < M$, 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 $F \ne 0$. 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 $E_i(\beta) \le h$ and $E_i(\eta) \le \nu(k+1)M$ because all variable excess weights are at most one. Hence $E_i(\alpha/M) \le h + \nu(k+1)$ 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

$$M^{k\nu} \dim U_h \le F_L \#((M^{-1}\mathbb{Z}^{k\nu}) \cap \mathscr{C}_{h+\nu(k+1)} \cap \Sigma(n_0 + k + 1)).$$

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

$$E_i(\beta; b, b') = (1 - q^{-1})\big(\beta(y_i) - (v_{b,i} - v_{b',i})\big). \tag{8.3}$$

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 8.4 (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 $k, k-1, \dots, 1$. 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 (8.3) shows $\sigma \in \mathscr{C}_{-1}$. At an even letter the inner expression is $P^2u_{b+1} - u_b$, 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 8.3 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 (7.13) 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 8.3 and 8.4 therefore give the needed bound

$$[\mathrm{GL}_k(R) : \mathcal{P}] \dim U_h \leq (1+\varepsilon)M_0. \tag{8.4}$$

This is the quantitative count `index_mul_finrank_excessCond_le` with exactly the needed leading constant.

8.4 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 8.5 (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 8.6 (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 8.5 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 8.5 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 8.6. Each monomial is counted at most 1 plus its cost when one sums these dimensions over $j \geq 0$. 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 (7.13). 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,

$$\sum_{0\le j<J}\dim V_j\le (1+z)M_0. \tag{8.5}$$

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.

8.5 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

$$\sum_{j\ge1}\left(\left\lfloor\frac{a_0}{p^j}\right\rfloor-\sum_i\left\lfloor\frac{\alpha_i}{p^j}\right\rfloor\right).$$

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

$$q_\alpha[X^\alpha]L_G(X^\beta)=q_\beta[X^\beta]L_{G^t}(X^\alpha). \tag{8.6}$$

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 8.7 (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 (8.6) for $|S| = k^2$. Define

$$\Lambda_T(F) = \mathrm{Tr}_{U_T}(S(K(F))|_{U_T}).$$

Equation (8.6), applied to the matrix substitution under conjugation, gives covariance of $S(K(F))$ under the tensor action of a Levi element. With $C_g(F)(X) = F(g^{-1}Xg)$ 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 $Y \mapsto gYg^{-1}$. 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.

8.6 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 8.8 (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

$$\widetilde{x} \cdot \beta - \nu n_0 k(\lambda+2)u. \tag{8.7}$$

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 (8.7). □

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 8.7.

Now let $g \in L$ and replace the generic matrices used above with

$$\Psi_{s,d} = \iota_s(g)\big(\iota_{s-d}(g^{-1})X_{s-d}\iota_{s-d}(g)\big)^{(P^d)}\iota_s(g^{-1}). \tag{8.8}$$

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 8.9 (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 (8.8) 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 (8.8) 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 8.8 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 (8.7) 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 (7.13), 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 3; 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 (8.6).

Theorem 8.10 (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,

$$y_0^t BZF \equiv p^{a_*} y_0^t F \pmod{w^{c_*}}. \tag{8.9}$$

Proof. In every letter take the means of Proposition 8.7 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 (8.8) 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 8.9 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 (8.9) 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 8.7, 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 8.9 applies again, so the entire contribution is zero modulo w^{c_*}. This proves (8.9). □

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.

8.7 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

$$D_w\rho(h)D_w^{-1} \in \mathrm{GL}_{M_0}(O) \qquad (h \in \mathcal{P}). \tag{8.10}$$

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 (8.10): 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 8.11 (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

$$w^{x_0} \mid w^{d(\alpha)}[X^\alpha]\rho(g)F'_x \tag{8.11}$$

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 (8.11) 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 $U_h \bmod w^c$ 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.

8.8 The Smith pairing lemma and the local Hecke theorem

We recall that the index condition $I(M, M'; z)$ of Chapter 7 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 7.2. 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 8.12 (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 $u_0 e_j \bmod u_0 w$. 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 7. 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 (7.18) 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 8.11; 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 8.10. Then by the transpose identity and $D_wB = BD_w$ one has an exact matrix identity, with B also repeated over the domain:

$$TB = J_2^tU\mathbb{B}J_1 = d'^t(Q'^tU\mathbb{B})Qd. \tag{8.12}$$

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 8.10 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 8.12 the direct sum of copies of U_h in every frame and spectator coordinate. The flag estimate (8.4) gives

$$\dim S = |X||J_0||I_0| \dim U_h \leq (1 + \varepsilon)D.$$

Lemma 8.12 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 (8.12) 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 (8.5); 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.

8.9 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 7.

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

$$e(u_{j+1}^{-1} X u_j) = b(X). \tag{8.13}$$

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

$$\pi_{j+1}(B)L_j[f(\overline{B}\cdot\overline{A})]\pi_j(A) = L_j[f], \qquad A \in \mathrm{GL}_{k_j}(\Lambda), \tag{8.14}$$

$$\pi_{j+1}(u_{j+1})L_j[e] = \gamma_j L_j[b]\pi_j(u_j). \tag{8.15}$$

The notation in (8.14) 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 (8.13) 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 (8.14) 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

$$\mathcal{F}_j = \pi_j(u_j) \sum_{\substack{y \in F_S^{k_j} \\ y_{\mathrm{last}} \in F_S^\times}} \psi_N(y_{\mathrm{last}})\pi_j(K_y). \tag{8.16}$$

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 8.13 (Sliding the finite Fourier operator)**.** *For a ramified vector v and a lift law as above,*

$$\pi_{j+1}(u_{j+1})L_j[c_{e_j}]v = \mathcal{F}_{j+1}L_j[b]v = \gamma_j L_j[b]\mathcal{F}_j v. \tag{8.17}$$

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 (8.14),

$$L_j[b_y] = \pi_{j+1}(K_y)L_j[b], \qquad L_j[e_z] = L_j[e]\pi_j(K_z),$$

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 (8.14) gives exactly the two displayed lift identities. Applying $\pi_{j+1}(u_{j+1})$ to the column expansion gives the middle member of (8.17). Applying it to the row expansion and using (8.15) 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 8.14 (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,

$$\mathcal{P}_{h-1}\cdots\mathcal{P}_0 v = \left(\prod_{0\le i<h} d_i q_i \prod_{0\le l\le 2i+1} \gamma_l\right) B_{2h}\mathcal{F}_0^h v, \qquad q_i = \prod_{w\in S}(N^{k_{2i+1}} - 1). \tag{8.18}$$

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 (8.17) 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 (8.14) 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 (8.17) 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 (8.18). 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 8.15 (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 (8.13) 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 8.16. 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 8.13 cannot be replaced by a sum over all vectors regardless of ramification.

8.10 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 7. 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 7. 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 (8.18) 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.

Part III

Spanning: Bessel and Poincare series

Chapter 9

Quantitative spanning by projected Bessel–Poincaré series

Chapter 10 supplies the expanded kernel, Schur, Laplace, finite-volume, Poisson and sieve arguments underlying the estimates used here. Statements that also require uniform geometric input retain that input; 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.

9.1 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 9.1 (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

$$\dim(P_s D_N)^{-1}(J) \le ((1+\alpha)/2 + b^{-2})c \quad \text{for some } s \in \mathfrak{G}. \tag{9.1}$$

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.

9.2 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 9.2 (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

$$n_0 = |\mathfrak{G}| d_\infty \,\mathrm{mass}_r(E). \tag{9.2}$$

Proposition 9.3 (Formal-source input: dimension transfer)**.** *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 (9.2). The nontrivial passage between signatures is not a formal consequence of dimension equality and cannot be omitted; the indicated trace-transfer theorem is the imported assertion. The same trace equality is used again to compare auxiliary lifts with the main chain in the lower-bound argument.

9.3 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 9.4 (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/FillingAdm`

`issible.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

$$1 \le |\kappa_v| \le b^{23}, \quad |v(\varpi)| \le b, \quad N \le b^{60s_0}, \quad b^{10} \le g_i \le b^{11}.$$

Also $h \ge 10^5$, $1000(2h+2)^9 \le k \le f$, $s_0^7 \le k$ and $(200fk\,\mathrm{rd}_E)^{1000} \le b$. 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

$$\mathfrak{d}(\gamma) = D_{h+1}\overline{D_{h+1}} - N^2 D_h \overline{D_h}. \tag{9.3}$$

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.

9.4 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 9.5 (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
$$\int_{U\backslash G_1} \left\| \int_{G_1} \phi(y) W_{v,\alpha}(\gamma y)\, dy \right\|^2 d\gamma \le P^{2|\alpha|} \frac{D_v}{d_v} \|\phi\|_2^2. \tag{9.4}$$

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 9.6 (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 (9.6) below, and the open-cell estimate of (9.7) 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 9.7 (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

$$T_v(\sigma, p) = \frac{\sqrt{c_G}\,(|p_h|^2-1)^{(k-1)/2}}{D_v}\prod_{i<h}|p_i|^{k+1+2(h-i-1)}\,J_{W_v^\flat}(w_p\sigma). \tag{9.5}$$

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 (9.5), 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

$$\beta_v(\xi)\Lambda_i(p)\Lambda_i(q), \qquad \beta_v(\xi) = c_{\mathrm{band}} b^{2000 s_0^9} \lambda_{k-s_0^7}(|\xi|^{-1}). \tag{9.6}$$

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

$$M_v(\rho_0,\gamma) = \left(\frac{\rho_0^{9s_0}}{\mathfrak{d}_v}\right)^{10^5 s_0^4} \frac{D_v^2}{d_v c_G (\mathfrak{d}_v/2)^{k-1}}. \tag{9.7}$$

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 (9.5) is intentionally different: in the sharp term one preserves cancellation between Fourier phases, while in (9.7) one uses upper bounds allowing losses. The large exponent on $\mathfrak{d}_v$ is compensated globally by discriminant bounds and congruence counts.

9.5 Construction and exact unfolding of the Poincaré map

Using the coordinate isomorphism of Proposition 9.4, 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

$$\sum_{v \in U_u(g)} \beta_u(vy) = 1. \tag{9.8}$$

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

$$\nu = \frac{(g^{2d_U+h}/N^k)^f}{v_f^{d_U}} \frac{D_E^{h/4}}{\mathcal{L}}. \tag{9.9}$$

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 (9.9), 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 (9.8) 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

$$Qz(g) = \sum_{u,i} z_{u,i} \sum_{\gamma \in G(E^+)} \phi_{u,i}(\gamma g). \tag{9.10}$$

When summing first over unipotents, the periodization replaces the cutoff β_u by 1 on the permitted finite coset, by (9.8) and the Bessel character being trivial on $U_u(g)$. This identifies (9.10) 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 $I \times I$)

$$A(\gamma; u, u') = \int_{G_\infty} \beta_u(y) \beta_{u'}(\gamma_\infty^{-1} y) W(y) \otimes W(\gamma_\infty^{-1} y)^*\, dy, \tag{9.11}$$

$$A_T(u, s, u') = \sum_{\gamma \in T \cap \mathcal{E}(u,s,u')} A(\gamma; u, u'). \tag{9.12}$$

Define $A_T^\flat$ with $W^\flat$ in place of W.

Proposition 9.8 (Formal-source input: convergence and reproduction)**.** *For every u, s, u' the sums in (9.12) 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 (9.11), is at most $\nu b^{-k^3 f}/2$. The sum of the norms in the inner periodization in (9.10) 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 9.5(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 9.9 (Exact coefficient identity)**.** *In the Petersson normalization there is a constant $v_0 = \mathrm{vol}(K(g)) > 0$ such that*

$$\langle Qe_{u,i}, \rho(s) Qe_{u',j} \rangle = v_0 A_{\mathrm{all}}(u, s, u')_{i,j}. \tag{9.13}$$

Here and below the pairing uses the convention in which the order of i, j agrees with (9.11); 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 9.8 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 (9.11); the two class cutoffs remain in it. Thus the integral is precisely (9.13). 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 9.10 (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 (9.8) 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 9.7. This proves the claimed identity (`amplitudeSum_besselS_one` in `New/Filling/AmplitudeSums.lean`). □

9.6 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

$$\mathcal{E}(A_T;\Omega) = \frac{1}{\nu^2}\sum_{u,u'\in\mathcal{C}}\sum_{s\in\Omega}\|A_T(u,s,u')\|_{\mathrm{HS}}^2. \tag{9.14}$$

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

$$\tau^{44}D_1 \le D, \qquad \tau^{18}\mathfrak{c}D \le D_1, \qquad \mathfrak{c} \le 1. \tag{9.15}$$

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 (9.2) shows $D/|\mathfrak{G}| = c^2/n_0$.

For two compatible collections of amplitudes A, B, any $\delta > 0$ gives

$$\mathcal{E}(A+B;\Omega) \le (1+\delta)\mathcal{E}(A;\Omega) + (1+\delta^{-1})\mathcal{E}(B;\Omega). \tag{9.16}$$

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 9.8 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 (9.14) 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

$$\psi_U(u) = \psi_U(\gamma^{-1}u\gamma) \quad (u \in U(E^+) \cap \gamma U(E^+)\gamma^{-1}). \tag{9.17}$$

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 (9.8). 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 (9.17) 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 (9.14), 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 (9.7) 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 9.11 (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

$$\mathcal{E}(A_T;\Omega) \le (1+b^{-4})^2 w + b^{-12}C_0. \tag{9.18}$$

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 (9.16) (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 (9.16), now using $\delta = b^{-4}$ and the extremely small truncation error, gives (9.18). 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 9.12 (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 9.11 with $\mathcal{L} = \varnothing$, $w = 0$, $C_0 = D_1$; the scale inequality (9.15) 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.

9.7 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:

$$\int_A J_{i+1}(p[i : c], q[i : c])\, dc = J_i(p, q), \qquad |J_0(p, q)| \le \mathfrak{e},$$

$$\left| \int_A \psi(\xi c) J_{i+1}(p[i : c], q[i : c])\, dc \right| \le \beta(\xi)\Lambda_i(p)\Lambda_i(q). \tag{9.19}$$

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 $\left(J_{h+1}(p_a,p_b)\right)_{a,b}$ 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 (9.19) 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 every required leading frame minor is a unit at all chosen places above g_0 not in Σ. This is the set denoted `fixingSet` in the formal proof. The omitted-minor test at places not in Σ is indispensable. 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

$$(1 + b^{-6}) b^{-6k|\Sigma|}. \tag{9.20}$$

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 9.4, and at the finite ones the systems of (9.20). 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 9.13 (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 (9.5), 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 9.13, 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 9.14 (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 9.11. 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 (9.15). 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.

9.8 A finite-field positivity argument for singular minors

The open-set estimate of Proposition 9.14 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 9.15 (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 9.14.

Proposition 9.16 (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 9.14. 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 9.15 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 9.14 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 9.9. With $\Sigma = \varnothing$ Proposition 9.16 therefore gives

$$\frac{1}{|\mathfrak{G}|}\sum_s \|P^*\rho(s)P\|_{\mathrm{HS}}^2 \le (1+9b^{-4})\frac{c^2}{n_0}. \tag{9.21}$$

This is where both local and global sharpness are indispensable.

9.9 First moment and bad finite-type loss

The estimates needed in addition to (9.21) are not general consequences of that second moment. They use the diagonal and central projections at the two different level primes. In fact

$$\sum_{i=1}^{c}\|(P^*P-I)e_i\| \le b^{-4}c, \qquad \|(I-R)P\|_{\mathrm{HS}}^2 \le 4b^{-4}c. \tag{9.22}$$

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_i, note that this set is invariant under right multiplication by the embedded Bessel-compatible translations in the proof; more directly one only needs the support of the functions on the set in which their i-component is the Bessel reduction. The convolution estimate and the covariance under the Bessel subgroup give

$$|\bar{S}_B|\|M_z(1)\|_{\mathrm{HS}}^2 = \sum_{s\in\bar{S}_B}\|M_z(s)\|_{\mathrm{HS}}^2 \le \mathcal{E}(A_{G(E^+)\setminus S_B}; \Omega_i).$$

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-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 9.12). Taking $z = 1$ and using the identity term gives the first inequality of (9.22); 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

$$\Re \sum_{a\in\mathfrak{G}_i} z(a) \sum_{u,j} A_{S_B}(u, \jmath_i(a)^{-1}, u)_{j,j}/\nu \le b^{-4}c. \tag{9.23}$$

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 (9.23). Applying this with $z = 1 - e_i$ and Lemma 9.2, 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 (9.22).

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.

9.10 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 (9.22), but the first bound in (9.22) 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 (9.21), 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 (9.1). 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 (9.1). 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 9.1. 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 9.1, 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 10

The analytic and arithmetic mechanism of the Bessel sieve

In Chapter 9 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 9.3. 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.

10.1 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 10.1 (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

$$\int_S \left(\frac{2}{\|j(s)y\|}\right)^c d\nu_S(s) \leq 16^c Q. \tag{10.1}$$

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 (10.1). 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 (10.1) 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 10.2 (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*

$$|\psi_l(g)|^2 |\det a_g|^c \leq 1, \qquad |\psi_l(g)| \leq (2/\|g\|)^c, \tag{10.2}$$

where a_g is the block diagonal part relative to the signs of J_*. Also

$$|\psi_l(xg) - \psi_l(x)| \leq 100(r + L_0 + 1)^5\|g - 1\|. \tag{10.3}$$

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 (10.2) 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 10.3 (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 (10.3) 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 10.1 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.

10.2 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

$$\int_G \psi(a^{-1})\psi(gay)\, da = \frac{\psi(g)\psi(y)}{d}. \tag{10.4}$$

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 10.4 (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,

$$\|F * x\|_2^2 = \frac{\|x\|_2^2}{d}\|F\|_2^2, \tag{10.5}$$

$$\left\| \int_G \phi(a)(F * x)(\cdot\, a)\, da \right\|_2^2 \le \frac{\|x\|_2^2}{d^2}\|F\|_2^2\|\phi\|_2^2. \tag{10.6}$$

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, (10.4), 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 (10.5). 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 (10.5) to w proves (10.6). 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

$$\|u_w\| \le \big((r+1)^4(L_0 + r + 1 + j)\big)^j. \tag{10.7}$$

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 (10.7) 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 10.4 to x_w to conclude that for every integrable ϕ,

$$\left\| \int \phi(a) F_w(\cdot\, a)\, da \right\|_{U\backslash G,2}^2 \le P^{2j} \frac{\|F\|_{U\backslash G,2}^2}{d} \|\phi\|_2^2, \qquad P \ge (r+1)^4 (L_0 + r + 1 + 2s_0^2), \quad j \le 2s_0^2. \tag{10.8}$$

This is exactly the constant in (9.4). 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 (10.8); the estimate concerns the *full* Bessel function. This distinction is used later in the energy comparison between full and cut functions.

10.3 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

$$\sum_t \|[P_t, X]\|^2 = Q_w(X), \qquad \sum_t \mathrm{tr}(P_t X) = \sum_i (w_i + a) X_{ii}, \tag{10.9}$$

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 10.5. *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 (10.9) 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

$$\left|\frac{\psi_l(g)}{\exp(\operatorname{tr}(\Lambda(l)\log g))} - \kappa\right| \le \kappa/10 + e^{-100T}. \tag{10.10}$$

   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 (10.9) 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 10.5, 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 (10.10) 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 (10.10).

10.4 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 $\|u_i(s)\| \le 1/2$,

$$\operatorname{tr}(\Lambda_1 \log j_1(s)) - \operatorname{tr}(\Lambda_2 \log j_2(s)) = \varphi(s). \tag{10.11}$$

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 (10.11). 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 10.6 (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,

$$\mu\{s : \|j_1(s) - 1\| \leq 1\} \leq \beta^{20N} \left| \int_S \theta(s)\, ds \right|. \tag{10.12}$$

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 (10.10), the off-diagonal decay and `exists_near_one` give two auxiliary bounds:

$$|\theta(s)| \le (2/\|j_1(s)\|)^a \qquad (0 \le a \le 18N), \tag{10.13}$$

$$4|\theta(s)| \le 8\left(\operatorname{Re}\theta(s) + 16\beta^{-50N}(2/\|j_1(s)\|)^{8N}\right). \tag{10.14}$$

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 (10.11). 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 (10.13) 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 10.6 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 (10.12). □

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 (10.14) which prevents cancellation. Rescaling the Haar measure multiplies both sides of (10.12) 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.

10.5 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 10.6. We reproduce the algebra of this embedding.

Lemma 10.7 (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 10.6 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

$$Y(y) = \int_S \psi_{l_1}(x^{-1} s_1 y)\, e^{-\varphi(s)} \rho(s_2^{-1}) v \, ds. \tag{10.15}$$

Conjugating the two group maps by their chosen frames, Proposition 10.6 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 (10.15) 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

$$\|Y(y)\| \le \int_S |\psi_{l_1}(x^{-1}s_1y)|\,ds, \qquad \int_S (2/\|x^{-1}s_1y\|)^{8r^2+1}\,ds \le b^n\|Y(x)\| \tag{10.16}$$

for every y in the normalized frame, with the specified exponent n, and $Y(x) \neq 0$. The second inequality follows from Proposition 10.6 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.

10.6 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 (10.16) cancels $b^n\|Y(x)\|$ and proves

$$\Phi(y) \le \nu\{g : \mathrm{bs}(yg) \le R\}. \tag{10.17}$$

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 (10.16) 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 (10.17):

$$\|Y(x)\|^2 \le d^2 I \int_{U\backslash G} \Phi(y)\|Y(y)\|^2\,dy \le d^2 I \nu(G)\|Y\|_2^2. \tag{10.18}$$

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 (10.18), 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

$$\begin{aligned}\|V-F\|_2^2 &\le b^{-2A}c^2 \int_{U\backslash G} \Phi(y)^2\,dy \\ &\le b^{-2A}c^2\nu(G)I \le D_v b^{-10r^2}.\end{aligned} \tag{10.19}$$

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 (10.17) 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 (10.19) is exactly the price for passing between the two, in a setting where no derivatives of the cutoff are required.

10.7 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 9.4. 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 10.6 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.

10.8 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 (10.8). 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 9.5. 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.

10.9 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 10.8 (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`.

10.10 An explicit induction for the sieve system

For the precise kernels defined in Section 9.13, 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

$$\sum_{y \in X} \beta(\operatorname{ratPoint}(c_v / l_{vi}, y)) \leq \varepsilon, \tag{10.20}$$

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

$$\left(\sum_{p \in A^i} \lambda_i(p) \right)^2 \leq \Gamma \prod_{a<i} c_a. \tag{10.21}$$

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 10.8 and the Jacobian computation give

$$S_i(V) \leq c_i \operatorname{Re} \sum_{p,q \in A^i} J_i(p,q) + c_i \varepsilon \Gamma \prod_{a<i} c_a. \tag{10.22}$$

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 (10.20), 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.

10.11 Frequency summation and the product formula

The bound (10.20) 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 (10.21) 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 9.13. 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.

10.12 Finite packets, their normalization, and the global comparison

We expand the finite-place mechanism in the same notation as Section 9.13. 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 10.11. 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 10.10 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.

10.13 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 $\delta > 0$,

$$\|A + B\|^2 \leq (1+\delta)\|A\|^2 + (1+\delta^{-1})\|B\|^2. \tag{10.23}$$

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 (10.23) 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 (10.19); 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 (10.23) 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 9.11 and the subsequent `energy_fixingSet_le` calculation, with the arithmetic numerical comparisons of the specialist chapter left unchanged.

Part IV

Lower bounds: periods, congruences and crystalline defects

Chapter 11

The determinant lower bound

Chapter 12 develops the complementary frames, finite square formulas, Rallis unfolding 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

$$\mathcal{U} = \sum_{\tau \in \Phi} (J(w_0(\tau)) - A(w_{\mathrm{nom}}(\tau), \tau)). \tag{11.1}$$

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 11.1 (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,*

$$\mathrm{covol}_\sigma^2(\Theta_1 y_1, \dots, \Theta_1 y_q) \leq e^V \quad (\sigma \in \mathrm{Aut}\,\mathbb{C}).$$

Then

$$n(H - 3L) \leq V + 2d_{\mathrm{ind}}. \tag{11.2}$$

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.

11.1 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

$$\sigma\langle u, v\rangle = \sum_h a_h\, \overline{\sigma^*(u(h))}^{\,T} G \sigma(v(h)). \tag{11.3}$$

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 11.2 (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 (11.3) 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

$$|\phi\langle u, v\rangle|_\ell \le 1. \tag{11.4}$$

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 (11.4).

11.2 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 11.3. 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 11.4 (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.

11.3 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

$$\det G_\theta(v) = Z_\theta \det \mathrm{Gram}(v), \qquad \frac{Z_\theta}{Z_1} = \frac{\det G_\theta(v)}{\det G_1(v)}. \tag{11.5}$$

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 11.5 (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 (11.4). Then

$$\mathrm{Ht}_{\le \alpha + 2d_{\mathrm{ind}} + hg_0}\left(\frac{Z_{\theta_\iota}}{Z_1}\right). \tag{11.6}$$

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 (11.5), 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 11.4. 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 (11.6).

11.4 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

$$\rho_p(\vartheta, s) = \frac{|L^{M\cup Q}(z, (\nu\vartheta^s)^{-1})|}{|L^M(z, \nu^{-1})|}, \qquad s \in \{-1, 1\}. \tag{11.7}$$

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

$$\frac{\Pi_{M\cup Q}((\nu\vartheta^s)^{-1},z)}{\Pi_M(\nu^{-1},z)} = N(Q)^{z^2}\prod_{v=1}^{z}\rho_{p_v}(\vartheta,s)^2. \tag{11.8}$$

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,

$$\left|\iota_\ell \det(\langle u_i, p_Z v_j\rangle)\right|_\ell \le e^{d_Z r\beta}. \tag{11.9}$$

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 (11.9); 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 (11.9) 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 11.6 (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 11.7. *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 11.4. 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 (11.8) 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`. □

11.5 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 11.8 (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 11.8. 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.

11.6 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

$$N(z, i, j) = r + z + i(t + z) + j(t - z). \tag{11.10}$$

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`; (11.10) 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 11.9 (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 11.7. 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 (11.10) 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 11.10 (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 11.11 (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 11.9 and the route calculation with Proposition 11.10, 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 11.5 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 11.2 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 11.5 and using $n - q \le 4n/b^2$ gives

Proposition 11.12 (Forward ratio)**.** *For the indicated family of short rows and fit family θ_ι, one has*

$$\mathrm{Ht}_{\le V + 2d_{\mathrm{ind}} + nb}\left(\frac{Z(\theta_\iota)}{Z(1)}\right). \tag{11.11}$$

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. □

11.7 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

$$w_v a_{i+1,v} d_{i,v} = \overline{\beta_{i,v}}^{-1} w_{v+2} d_{i,v+1} a_{i,v+1}. \tag{11.12}$$

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`; (11.12) 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

$$A_{i+1,v}d_{i,0}a_{i,0} = \frac{w_{v+1}}{w_1 w_0 \overline{\beta}}\, d_{i,v}A_{i,v+1} \tag{11.13}$$

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 (11.12) 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

$$\langle z, a_{i,v}A_{i,v}j_i u\rangle = \kappa\langle d_{i,v}z, A_{i+1,v}j_{i+1}u\rangle \tag{11.14}$$

with the product $\kappa\beta$ costing at most v squares. In the general injectivity case only (11.13) 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 11.13 (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 (11.13): $A_{i+1,v}d_{i,0}a_{i,0}j_i u = 0$. 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 11.14 (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 (11.14) on one side and (11.13) 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 (11.13)–(11.14). □

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 11.15 (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 11.13. Taking determinants of the resulting Gram comparison and applying Proposition 11.15 stage by stage yields the following identity.

Proposition 11.16 (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*

$$Z(\theta) = |c(\theta)|^n \prod_{j=1}^{m} \lambda_j(\theta) > 0. \tag{11.15}$$

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 (11.15) 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 11.7. 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.

11.8 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

$$\Gamma_j(\vartheta) = \prod_{i=0,+,-} \Pi_{M\cup Q}((\nu_i \vartheta^{s_i})^{-1}, e_i), \quad s_0 = s_+ = 1, \quad s_- = -1. \tag{11.16}$$

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 11.17 (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*

$$u_\vartheta^* g_\vartheta = \gamma_\vartheta A, \tag{11.17}$$

$$|\gamma_\vartheta/\gamma_1|^2 = \Gamma_j(\vartheta)/\Gamma_j(1). \tag{11.18}$$

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 11.17; the quotient of the two norm identities for κ gives (11.18). 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.

11.9 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 11.17 give

$$h_j u_\vartheta^* g_\vartheta g_\vartheta^* u_\vartheta = h_j|\gamma_\vartheta|^2 AA^*.$$

Taking determinants and using $u_\vartheta^* p_\vartheta = u_\vartheta^*$ gives

$$\det(u_\vartheta^* u_\vartheta) D_\vartheta = h_j^n |\gamma_\vartheta|^{2n} \det(AA^*). \tag{11.19}$$

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

$$\frac{\lambda_j(1)}{\lambda_j(\theta_\iota)} = \frac{D_{\theta_\iota^j}}{D_1}\left(\frac{\Gamma_j(1)}{\Gamma_j(\theta_\iota^j)}\right)^n. \tag{11.20}$$

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 11.18 (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 11.18, 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 (11.9) to these two families. The matrix of their pairings

$$(\langle g_\vartheta f_{a_i}, p_\vartheta(h_j g_\vartheta g_\vartheta^* z_h)\rangle)_{i,h}$$

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

$$|\iota D_\vartheta|_\ell \le e^{nr\beta_\ell}. \tag{11.21}$$

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

$$\|p_\vartheta y\|_\sigma \le \|y\|_\sigma. \tag{11.22}$$

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 (11.22). 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 (11.3), contraction (11.22), and the two table norm estimates in Proposition 11.17 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

$$|\sigma D_1| \leq e^{2nrB_0}. \tag{11.23}$$

Using witness D_1 in the height definition, the local estimate (11.21), the global sum of β_ℓ, and (11.23), we obtain

$$\mathrm{Ht}_{\leq 3nrB_0}\left(\frac{D_{\theta_\iota^j}}{D_1}\right). \tag{11.24}$$

11.10 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

$$\rho_p(\vartheta, 1) = \frac{\mathcal{A}_p(\vartheta)}{\mathcal{A}_p(1)} = N(Q)^{-(2z-1)} \frac{\|\Xi_{p,\vartheta}\|^2}{\|\Xi_{p,1}\|^2}. \tag{11.25}$$

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 11.19 (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

$$\mathrm{Ht}_{\leq B_0 - \sum_\tau J(w_0(\tau))}(\|\Xi_{p,1}\|^2). \tag{11.26}$$

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 11.20 (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 $\vartheta = \theta_\iota^j$,

$$|\phi Q|_\ell \, \ell^{-F_\ell} \le \left| \phi \left(\frac{N_c^2 \langle u_\phi, \Xi_{p,\vartheta} \rangle}{\rho_{p_2}(\vartheta, 1)} \right) \right|_\ell . \tag{11.27}$$

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 (11.27), 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 (11.27) 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 (11.26) and by the square of the inverse ratio of the ramified flat pair p_2 (height at most $2x^{999}$ from fit). Using (11.25), 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 (11.1). Since $x^{999} \le B_0$ by the admissible magnitudes, we conclude

$$\mathrm{Ht}_{\le 6B_0 - \mathcal{U}}(\rho_p(\theta_\iota^j, 1)^{-1}). \tag{11.28}$$

This is `heightLe_inv_ratio`.

11.11 Combining the special values and closing the lower bound

We now have precisely the estimate for the long string in (11.16). 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 (11.28). 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 (11.8) are local units. Consequently

$$\mathrm{Ht}_{\le 10mx^{999}+k(6B_0-\mathcal{U})}\left(\frac{\Gamma_j(1)}{\Gamma_j(\theta_\iota^j)}\right). \tag{11.29}$$

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 (11.29), the compressed determinant estimate (11.24), and identity (11.20). 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 11.16 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 11.21 (Reverse ratio)**.** *For every fit family,*

$$\mathrm{Ht}_{\le 2n(2L-H)}\left(\frac{Z(1)^2}{Z(\theta_\iota)^2}\right). \tag{11.30}$$

This is `Data.heightLe_volume_ratio_sq` (`New/Endgame/LongForms.lean`). Finally choose a fit family, apply Proposition 11.12 and square its ratio. The product of its ratio squared and that of (11.30) is identically one by injectivity. Lemma 11.4 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 11.1 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 (11.28). 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 12

Period identities, finite matching, and the geometry of the lower bound

The proof of Theorem C in Chapter 11 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 11 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.

12.1 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

$$XY = 0, \quad TY = I_{b_0}, \quad XZ = I_{a_0}, \quad YT + ZX = I_N. \tag{12.1}$$

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 12.1 (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 (12.1) 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 12.1, and for every compatible pair furnished by k,

$$\Phi(X) = \chi(\det k)\Psi(Y). \tag{12.2}$$

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 12.2 (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

$$\det X_2 \det k = \det Y_1 \tag{12.3}$$

follow from (12.1); 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 (12.3) gives (12.2). 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 (12.2). 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 12.1. 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

$$\Delta_\chi(X, Y) = \frac{\chi(\det k)}{|\operatorname{GL}_{b_0}(R)|}, \tag{12.4}$$

and put it equal to zero otherwise. The determinant is well-defined by Lemma 12.1. 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

$$\Delta_\chi(X, Y) = \chi(\det g/(\det f \det q))\Delta_\chi(f^{-1}Xg, g^{-1}Yq). \tag{12.5}$$

For a covariant Ψ as in Proposition 12.2, summation of (12.4) against Ψ on column frames for a fixed X gives exactly its partner Φ. Thus the denominator in (12.4) is the orbit size, not a multiplicative loss in the integrality estimate.

Lemma 12.3 (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*

$$\sum_Y \Psi(Y)L'(Y) = c\sum_X \Phi(X)L(X) \tag{12.6}$$

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 (12.6) 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 (12.1). 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.

12.2 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

$$\sum_B \psi(\operatorname{tr}(CB)) = \begin{cases} |R|^{uv} & C = 0, \\ 0 & C \neq 0, \end{cases} \tag{12.7}$$

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 12.4 (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

$$\mathcal{F}_\psi F(A) = \chi(\det A)^{-1} G_d(\chi, \psi) F(I_d), \tag{12.8}$$

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 (12.8). 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

$$Y_* = \begin{pmatrix} Y \\ I_{a_0} \\ 0 \end{pmatrix}, \qquad X_* = \begin{pmatrix} I_u & -Y & V \\ 0 & 0 & I_d \end{pmatrix}. \tag{12.9}$$

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 12.3,

$$|\operatorname{GL}_{a_0}(R)|L'(Y_*) = c|\operatorname{GL}_{u+d}(R)|L(X_*). \tag{12.10}$$

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 (12.10) 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 (12.10) by $p(Y)$, summing over Y and averaging over V proves

$$|\operatorname{GL}_{a_0}(R)|\sum_Y p(Y)L'\begin{pmatrix} Y \\ I_{a_0} \\ 0 \end{pmatrix} = c|\operatorname{GL}_{u+d}(R)|\sum_X p^+(X)L\begin{pmatrix} I_u & X_{11} & X_{12} \\ 0 & X_{21} & X_{22} \end{pmatrix}. \tag{12.11}$$

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 (12.11) remains valid without primitivity of χ, because it uses only covariance and the matched-pair relation.

Proposition 12.5 (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 (12.7) 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 12.4 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 (12.10). 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 (12.11) 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 12.4. □

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.

12.3 The graph term in doubling

To identify the analytic input in the Gram factors of Proposition 11.8, 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 12.6 (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 12.6, 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 12.7 (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.

12.4 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 12.8 (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 12.9 (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.

12.5 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.

12.6 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 12.10 (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 12.11 (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 12.10 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.

12.7 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

$$C|c|^2 = \Pi_\nu(\dim V - \tfrac{1}{2}\dim W)\ell_\nu(z), \tag{12.12}$$

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 (12.12) 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 11.8.

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 12.12 (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 12.13 (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 (12.12). 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.

12.8 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 11.11 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 12.4 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.

12.9 Why a false lift with one common root contributes no period

The purpose of comparison forms in Section 11.8 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 12.14 (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 12.15 (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*

$$q^{m+1}|X| \leq 1,$$

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

$$\|F\|_\infty \leq \left(\prod_{j=0}^{m-1} (1 + q^j |X|) - 1 \right) \|F\|_\infty,$$

by Lemma 12.14. Using $1 + u \leq e^u$ for $u \geq 0$,

$$\prod_{j=0}^{m-1} (1 + q^j |X|) \leq \exp\left(|X| \frac{q^m - 1}{q - 1} \right) \leq \exp\left(\frac{1}{q(q-1)} \right) \leq e^{1/2} < 2.$$

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 12.15. 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.

12.10 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 11.17. 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 12.16 (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 12.17 (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 12.17. 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.

12.11 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 11.5.

Chapter 13

Congruences, auxiliary primes, and the fit family

Chapter 14 reconstructs the auxiliary-prime and triangular relations, the finite duality and matrix-span arguments, 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.

13.1 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 11, their test value is denoted X_w. The condition imposed on a form v is

$$v \perp (Z^{\perp} \cap \operatorname{im} e \cap \ker H_w(X_w)) \quad \text{for every split } w \text{ outside the level and its conjugate.} \tag{13.1}$$

It is essential that the subspace in (13.1) 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 (13.1) 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

$$\left|\iota\left(\det\left(\langle u_a, P_Z v_b\rangle\right)_{a,b=1}^{q}\right)\right|_{\ell} \le \exp(qr'\beta) \tag{13.2}$$

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 (13.1). 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 (13.1) 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 13.1. 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 (13.2) 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 13.2 (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.

13.2 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:

$$\operatorname{tr} U(g) = \psi(g) \operatorname{tr} T(g) + \sum_{e=0}^{t'-1} \epsilon_\ell(g)^{-e} \qquad (g \in G_E). \tag{13.3}$$

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

$$g_i^a - 1 \mid M_g \quad (i = 0, 1,\ 1 \le a \le rt). \tag{13.4}$$

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 13.3 (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 (13.3) 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 (13.3) 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.

13.3 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 13.4 (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

$$\begin{aligned}\sum_{v \in S} \inf_{A|v} \lambda_{\mathcal{O}}(Z^1_v(\mathcal{L}, \mathcal{O}/(a))) + r' v_\pi(\alpha - \epsilon_\ell(g^c)\alpha') \\ \leq (|S| + f)\lambda_{\mathcal{O}}((\mathcal{O}/(a))^{r'}) + \beta_0.\end{aligned} \tag{13.5}$$

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 13.5 (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*

$$\frac{\beta_0}{e_{\mathcal{O}}} + v_\ell \det(R(z) - 1) + v_\ell \det(R((z^c)^{-1}) - 1) \leq 10 f r' c_\ell. \tag{13.6}$$

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 (13.5) to identify exactly where local conditions enter the subsequent use of global duality.

13.4 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 13.6 (One anticyclotomic prime (see Proposition 14.9))**.** *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:

$$h_{P'}(\varphi_P) = \begin{cases} 1, & P' = P_0, \\ 0, & P' \neq P_0, \end{cases} \quad (P \neq P_0), \tag{13.7}$$

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

$$\begin{aligned} c(\tau_P) &= t_{P,i}\, c(\varphi_P), \\ c((\tau_P^c)^{-1}) &= t_{P,i}\, c((\varphi_P^c)^{-1}). \end{aligned} \tag{13.8}$$

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

$$\psi_y(g) = \prod_{P \in J} y_P^{h_P(g)}. \tag{13.9}$$

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 (13.7) 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 (13.7) 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 13.7 (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 (13.9): 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 13.7 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.

13.5 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 13.8 (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*

$$\lambda_\mathcal{O}\big(Y \cap \ker(s\Phi_0 - \Phi_1)\big) \le \lambda_\mathcal{O}\big(Y \cap \ker \Phi_0 \cap \Phi_1^{-1}(\Phi_0 Y)\big). \tag{13.10}$$

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 (13.10). □

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 13.9 (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 (13.5). 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 13.10 (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 13.8 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.

13.6 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 13.11 (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:

$$\ell \mid (y-1)^{\ell^k}, \qquad \ell^M \mid y^{\ell^N} - 1 \quad (N \ge k + M). \tag{13.11}$$

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.

13.7 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 13.12 (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 13.13 (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

$$v_\ell(F_{i,j}(\zeta^j))\leq B_i+\delta,\qquad v_\ell(D_h(\zeta_k))\leq\delta\quad\text{for every } k,h. \tag{13.12}$$

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 13.12 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 (13.12). □

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.

13.8 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

$$V(Q,\psi) = \sum_{a \in A_Q} c_a \psi(g_a), \qquad |c_a|_\ell \le 1. \tag{13.13}$$

In particular, on a character given by (13.9), 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

$$v_\ell V(Q \cup \{P\}, \psi) \le v_\ell V(Q,\psi) + 8c. \tag{13.14}$$

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 (13.13) 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

$$v_\ell(\mu_i(z)-1) \le v_\ell(q_w-1) \le c_\ell \tag{13.15}$$

for its tame generator z. The last inequality follows from (13.4). 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 13.14 (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 13.6. 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 13.10, 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 (13.11) 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 13.11 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 13.7: 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 (13.14), using φ_{P_0}

close to z and the bound (13.15). On inserting the other primes successively, their Frobenius values are the helper value by (13.7). 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 (13.11) all these specializing points satisfy the near-1 power condition with one common exponent. Lemma 13.13 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 (13.9) 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 13.15 (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*

$$\sum_{P \in Q} v_\ell(\psi(\tau_P) - 1) \le c_\ell, \tag{13.16}$$

$$\lambda\big(\mathfrak{S}_M(\psi T_i; S \cup Q \cup Q^c, \mathcal{L}) \cap \ker \mathrm{ev}_{\gamma_i}\big) \le (10 f r_i + 1) c_\ell, \tag{13.17}$$

$$\left|\iota(\rho_p(\vartheta, s)^{-1})\right|_\ell \le \exp(9 c_\ell \log \ell) \quad \textit{for every ramified } p, s. \tag{13.18}$$

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 13.5 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 13.14, using (13.15) 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 (13.14) and the ramified-pair norm relations in that file gives $v_\ell(\iota\rho_p(\vartheta, s)) \le 9c_\ell$, equivalent to (13.18). This is exactly `exists_mFoldTwist_selmerLengthLe` in `New/Selmer/TwistForRows.lean`. □

13.9 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 13.16 (Formal-source input: one integer for all odd classes)**.** *There exists an integer $X > 0$ with*

$$X \le 2^{4(N_E+t+3)^{27}} h_E, \tag{13.19}$$

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 (13.19), 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 13.17 (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 13.17; 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 (13.4); 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 13.17 arise from the Hecke string lemma discussed below. In particular the global odd term in (13.19) and the bounded valuations of the separating products are sufficient even if the even-class modules are very large.

13.10 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 (13.3), 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

$$\mathcal{N} = [\gamma a] - s[a] + p_0 D \sum_g n_g[g], \qquad \beta_y = \det\left(\sum_g \mathcal{N}_g(U_y(g) - 1)\right), \tag{13.20}$$

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)

$$\operatorname{tr}((U_y(\tau_j) - 1)U_y(g)) - X_j \operatorname{tr}((U_y(\phi_j) - 1)e_j(U_y(\phi_j))U_y(g)), \tag{13.21}$$

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

$$F\beta \in \mathfrak{J}, \qquad \beta_y = 0 \ (y \in \mathfrak{X}), \qquad [X^d]F_y = 0 \ (y \in \mathfrak{X},\ d < \mathbf{1}), \qquad [X^{\mathbf{1}}]F_{y_0} = c_0, \tag{13.22}$$

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 (13.20), 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 13.18 (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 (13.22) exists with*

$$v_\pi(c_0)\le b_0+\sum_e b_e+t'\left(\sum_v v_\pi(o_v(1))+2|s|v_\pi(D)\right)+(|J'|+1)v_\pi(D).\tag{13.23}$$

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 (13.23). 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 13.16, without any even-class length term.

13.11 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 (13.20) 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 13.19 (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*

$$S(X) \prod_{i<t'} (X - \epsilon_\ell(\phi)^{-i}) \quad \textit{with } |S(1)|_\ell \geq 1,$$

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 13.17 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 (13.21), 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 (13.22) 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 (13.1), after inserting the chosen label projector. It follows that the diagonal operators satisfy the required integrality hypothesis on the projected tested forms.

13.12 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 13.16. 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

$$\sum_{i \text{ blocks}} \lambda(\text{loss of crystalline relations on block } i) \le 4r' f \operatorname{loss}_\ell, \tag{13.24}$$

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 13.20 (Formal-source input: arithmetic Ribet element)**.** *Use the apt columns and rows of an auxiliary pair and a twist satisfying* (13.16) *and* (13.17) *at a depth $M > \lfloor rD_\ell \rfloor$. Suppose their normalized matrices satisfy* (13.3) *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* (13.20) *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 (13.24). The Tate costs in (13.23) contribute at most $v_\ell(X)$ per block plus their shares in (13.24). 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 (13.16). 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

$$\mathrm{DenominatorLe}(\iota, Z, D_\ell \log \ell), \tag{13.25}$$

that is (13.2) 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 (13.25) marks the end of the congruence argument.

13.13 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

$$\sum_{\ell \in T} 9 c_\ell \log \ell \le x^{999}, \qquad \sum_{\ell \in T} D_\ell \log \ell \le B_0. \tag{13.26}$$

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 13.2, using the stated formal-source inputs. For each ℓ and ι, make the finite list of rows in Proposition 13.3, including all rank-one pairs, before choosing any twists. Apply Corollary 13.15 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 (13.25) 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 (13.26) is its required finite-sum bound. The special-value inequality (13.18) and the first inequality of (13.26) 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 13.1. 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 14

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 13.1. All claims imported from the formal proof without a reconstructed argument are marked as such.

14.1 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

$$\left|[X^i]\left(\det(X - R(g_w)) - \iota P_w(X)\right)\right|_\ell \le \eta \qquad (0 \le i \le n). \tag{14.1}$$

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 (14.1).

Lemma 14.1 (Trace extraction from attachment)**.** *Let R be attached to P as in (14.1). 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 (14.1) 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 14.2 (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 14.3 (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 14.4 (Two cancellations of auxiliary factors)**.** *Suppose R_i are as in Lemma 14.3 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.

14.2 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 14.5 (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 14.6 (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 (13.4)). 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 14.7 (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 14.5 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 14.2 and the particular shifts of the algebraic weight in `exists_hasRow_auxPair` of `New/Selmer/RowsExist.lean`.

14.3 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 14.8 (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

$$x(zh)(ay(zh)) = x(c(zh))(ay(c(zh))). \tag{14.2}$$

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 (14.2) 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

$$x(b)(aR(k)y(\tau)) = x(c\tau)(a\chi(ck)R(ck)^{-1}y(cb)). \tag{14.3}$$

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 (14.3) 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.

14.4 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 14.9 (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.

14.5 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

$$Z(kuk^{-1}) = \varepsilon_\nu(k)^{e+1} Z(u) \qquad (u \in K). \tag{14.4}$$

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,

$$2Z(f_j) \mid b \text{ in } A_\nu \quad \Longrightarrow \quad bd(\tau_j) = 0. \tag{14.5}$$

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 14.10 (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 14.11 (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 14.10 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

$$d(f_{(e,i)}) = 2d(b_i) \tag{14.6}$$

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 (14.5) 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 (14.6), $2d(b_i) = 0$. 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.

14.6 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 14.12 (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 14.12 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 14.13 (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 14.12. 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.

14.7 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 14.5 explain the mechanism behind the following statement.

Lemma 14.14 (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.

14.8 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.

14.9 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 14.15 (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

$$\mathcal{N}\Theta B(1) = nTB(1), \qquad L\Theta B(1) = 0. \tag{14.7}$$

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 14.16 (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 (14.7), take

$$Z_g = \Theta B(1)\sum_x v(a(h(g,x)))\rho(x^{-1}).$$

The second part of (14.7) 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 13.11. Proposition 14.16 reconstructs the transfer calculation used in these inputs; it does not establish their global automorphic hypotheses.

14.10 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 14.14). 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 14.12 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 13.12 and 13.13 to yield the fit family.

Chapter 15

Local crystalline defects and the integral cup-product estimate

Chapter 16 supplies further proofs of the classical cohomological rank and index calculation, the chart algebra 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.

15.1 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

$$\mathrm{rl}_p(S,T) = \frac{\mathrm{length}_{\mathbb{Z}_p}(S/(S\cap T))}{[K_O:\mathbb{Q}_p]}. \tag{15.1}$$

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

$$j_{T^*(1)}(\gamma, -u-1) = j_T(\gamma, u). \tag{15.2}$$

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

$$Z_f(T)/p^a Z_f(T) \hookrightarrow Z^1(\Gamma, T/p^a T). \tag{15.3}$$

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 15.1 (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

$$g\psi - \psi = p^a q(g). \tag{15.4}$$

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 15.2 (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 15.1 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`. □

15.2 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 (15.4) 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 15.3 (Local congruence loss)**.** The local Tate-duality index estimate, derived in Lemma 16.2 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

$$p^{Md(2+d_F)}, \qquad d = \mathrm{rk}_{\mathbb{Z}_p} T_v. \tag{15.5}$$

Indeed, by (15.3) their orders are the powers of p with exponents M times their $\mathbb{Z}_p$-ranks. The dual jump relationship (15.2) 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 (15.5) 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.

15.3 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

$$\delta_0 = \begin{cases} 0, & e_0 = 1, \\ v_p(e_0) + 1, & e_0 > 1. \end{cases} \tag{15.6}$$

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

$$e_i = \frac{1}{q_0} \sum_{j=0}^{q_0 - 1} \left(\omega(i)^{-1} \zeta \right)^j . \tag{15.7}$$

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 15.4 (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.

15.4 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

$$\mu(\varphi^{-a} v) \in \mathrm{Fil}^{c(a,u)} B_{\mathrm{dR}} \tag{15.8}$$

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

$$\mathrm{st}_l(a, u) = \begin{cases} 0, & a < l, \\ -u, & a \geq l. \end{cases} \tag{15.9}$$

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

$$c_p(a, n) = v_p(n) + \delta_0 + 3a, \tag{15.10}$$

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 15.5 (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 $p^{\delta_0 + 3a}$ 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 $v \in \mathfrak{N}_{\widetilde{T}}^{\mathrm{loc}}$ 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 15.5; 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 $d \ge 0$,

$$\sum_{j \in I(n,c)} (v_p(j) + d) \le (4M+2)\lfloor \log_p H \rfloor + (4n+2)d. \tag{15.11}$$

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.

15.5 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 15.6 (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 15.5, 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.

15.6 Cup-product vanishing from the three sides

Each pairing in (15.4) 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 (15.4), 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 $g, h \in \Gamma'$:

$$P(D_g, gT_h), \quad P(W - X, T_g), \quad P(D_g^W, gC), \quad P(\varphi X - W, g\varphi Y - \varphi Y + b_g). \tag{15.12}$$

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

$$o(a,u) + o'(a,v) + \mathbf{1}_{a \geq l} o_2(a-l, -u-v-1) > 0 \tag{15.13}$$

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. This is *formal-source input* `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 (15.12).

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 (15.13), 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 $\mathfrak{e}_a$.

15.7 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 (15.13) 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 (15.12), 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 15.6(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 15.6(iv). At cuts the resulting points can improve their depth in the two windows, paying the bound (15.11) 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

$$10^4 K(\lfloor \log_p(2b) \rfloor + 1) \tag{15.14}$$

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 15.6(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 15.4 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

$$\beta_i + W_0 + \eta \leq (p-3)A_0, \qquad 6(A_0 + \eta + W_0) < \beta_i, \qquad \mathrm{NM}(p, 6(A_0 + \eta + W_0), \beta_i). \tag{15.15}$$

Then the following estimate holds.

Theorem 15.7 (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

$$(2A_0 + 1)n_i + 3(W_0 + 1)r_{i[-1]} + 6A_0 n_{i[-1]}. \tag{15.16}$$

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 (15.16). The verification of (15.16) 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 (15.16) 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 (15.1) proves the displayed normalized length bound from the integral depth statements. The zero-loss part of Theorem 15.7 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`.

15.8 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

$$p^{3\kappa_0 + 150(\eta + W_0 + 1)} \tag{15.17}$$

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 (15.17), 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 (15.17); 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.

15.9 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

$$\lambda_p(h_*, b_*) = v_p\left((2b_*)!\left(\prod_{\max(b_*-30h_*^3,0)<j\le b_*+30h_*^3} j\right)^{h_*^3}\right). \tag{15.18}$$

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

$$\Delta_p(T; \mathcal{Q}) \le d_F\, 10^5 \left(\mathrm{rk}_O\, T \left\lfloor \frac{2\lambda_p(h_*, b_*)}{Q} \right\rfloor m + \mu\lambda_p(h_*, b_*)\right). \tag{15.19}$$

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 15.2 then suffices, at cost bounded by the first term of (15.19).

- If $p > Q$ and $v \geq h_*^3$, use the two-sizes exponent (15.14). 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 15.2 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 (15.18) 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 15.7 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 (15.19).

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.

15.10 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 (15.2), while the rank-one blocks have all jumps small and nonpositive. The local substitution in (15.19) 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 (15.19). 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:

$$\mathrm{loss}_p = 10^6 \frac{m}{k}(b_1 v_p(g) + \lambda_p(h_2, b_1)), \qquad g = g_0 g_1. \tag{15.20}$$

At the selected split level primes use (15.17), 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 (15.19). With the preceding scale and spacing verifications the result is the following exact statement.

Proposition 15.8 (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 15.3 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 16

Integral period frames and the algebra of the local loss

We now examine the steps behind the local estimate in Chapter 15. 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.

16.1 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 15.1. 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 15.1.

Proposition 16.1 (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 16.2 (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 16.2 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 15.1 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.

16.2 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

$$E = z^{e_0} a(z) + p. \tag{16.1}$$

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 (16.1) 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 (16.1) 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`).

16.3 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 16.3 (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 16.3 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

$$a \in A, \quad p^m \mid a \text{ in } A_{\max} \quad \Longleftrightarrow \quad a \in (p,E)^m A. \tag{16.2}$$

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

$$\widehat{\lambda}(xy) = x\widehat{\lambda}(y) \qquad (x \in S_B,\ y \in B^+). \tag{16.3}$$

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 (16.3) 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.

16.4 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:

$$MY = W_0 \operatorname{diag}(t^{k_1}, \ldots, t^{k_d}), \qquad -a \leq k_i \leq b.$$

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 16.2. 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

$$\varphi U = v^a U F_1, \tag{16.4}$$

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:

$$P = \{x \in A^d : \ \exists y \in S_B^d, \ Uy = E^h x\}. \tag{16.5}$$

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 16.3 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 16.3, 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

$$QL = \tau^h I_d. \tag{16.6}$$

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 (16.6) 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 (16.6) 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

$$E_m^a \lambda(\varphi^m(n)x) = \mathrm{val}(\lambda)\, \psi_m(M)\, \varphi^m(Z)x, \tag{16.7}$$

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; (16.7) is their explicit linear algebra.

16.5 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

$$\varphi^{\ell_0} Z = E_{\ell_0}^{b}\, C\, V, \tag{16.8}$$

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 16.3 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

$$p^{2D} H(X) = H_0(z) + z^{p^{\ell_0} D} H_1(X), \qquad H_0 \in S, \quad H_1 \in S[[X]]. \tag{16.9}$$

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 (16.9) 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

$$p^{4D}x = (p^{4D}I - MN)x + M(Nx) \tag{16.10}$$

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 (16.8). 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.

16.6 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 $I = \mathbb{F}^{\times}_{p^{f_0}}$ 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 15.6.

Lemma 16.4 (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 15.6. 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.

16.7 Exact cokernel multipliers via finite differences

We expand the arithmetic in Proposition 15.6. 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

$$\alpha([\epsilon]^v - 1) = E\tau(p^{v_p(v)} + \beta z). \tag{16.11}$$

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 (16.11), 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 (16.11) 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 15.6(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 16.5 (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 15.6(i), $E^n z^a$ lies in every level-one ideal $J_1(i; r_0)$. From (16.1), $z^{e_0} \in (p, E)$ 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 15.6 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 16.6 (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 16.5 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 $0 \le m \le L$,

$$p^{L-m}z^m \in J_0(i;r_0). \tag{16.12}$$

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 15.6 now proves (16.12). 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 (16.12). 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 15.6. 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`.

16.8 A rank-sensitive bound for raising filtered conditions

We now give further details of the length mechanism for Theorem 15.7. 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 (15.1) only at the end.

Lemma 16.7 (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

$$E^{p^a} - E_a \in pS \tag{16.13}$$

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 (16.13) 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 16.7. 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 (15.15), 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 (15.16). 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 16.7 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 (15.16), 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 15.7, 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.

Part V

The arithmetic conclusion and proof audit

Chapter 17

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.

17.1 Uniform scale asymptotics

Fix N_0, d, and $0 < \varepsilon < 1/100$. We use the family of admissible records of Proposition 3.2 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 17.1 (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.

17.2 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 17.1. 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.

17.3 The exact level cancellation revisited

In the notation of Chapter 3, 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 3.2. It is not necessary to have a strong pointwise estimate for the prime error at x_1 for this prime choice.

17.4 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.

17.5 Summation of the fit-family costs

We use the same notations as Chapters 13 and 15. 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 4.3 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 17.1. The local-loss part is calculated in Chapter 15. 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 17.1. 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.

17.6 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 18.4. The main combination in Proposition 18.5 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 17.1. In particular the sum is eventually at most one. Similarly the single-lift allowance in Chapter 11,

$$3nrB_0 + n\big(10mx^{999} + 6kB_0 - k\mathcal{U}\big) \le nk(YB_0 - \mathcal{U}),$$

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 18

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.

18.1 Collecting the arithmetic inputs

With the data of Chapter 3, 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 18.1 (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 18.2 (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 18.3 (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 $\mathrm{covol}^2_\sigma(y) \le e^V$, then*

$$n(H - 3L) \le V + 2d_{\mathrm{ind}}.$$

Together with the crude bound

$$|H| + m|\mathcal{R}| \le b^2 x^{99}, \tag{18.1}$$

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 18.4 (Crude size comparison)**.** *Along the data in Proposition 3.2, including the chosen second cutoff, inequality* (18.1) *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 (3.2) 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 (18.1). 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. □

18.2 From A and B to an almost-basis of short forms

Choose F as in Theorem 18.2 and use the set S from Theorem 18.1. 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.

18.3 Comparison with the lower bound

Proposition 18.5 (The main arithmetic inequality)**.** *For all sufficiently large x along the admissible family of Proposition 3.2, 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 (18.1) 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:

$$m\mathcal{R} \le 4L + 4x^{99} + 10 \log(1 + b^2)B. \tag{18.2}$$

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 (18.2). 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.

18.4 Extracting the analytic remainder

Here is the short algebra behind (3.9). 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 (3.9), including its constants. In particular, no estimate involving the Chowla–Selberg constant is needed at this final stage.

18.5 Choosing the second cutoff and the power saving

Fix $0 < \varepsilon < 1/100$ and choose admissible data by Proposition 3.2 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 3.2. 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 (3.9) 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,

$$R_{N_0}(x) \le x^{-1/2+8\varepsilon}. \tag{18.3}$$

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 18.6 (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.

18.6 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 (18.3). 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 19

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`; 3
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`; 3
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`; 17.3
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 $r + jt$.	7

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`; 9
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.	7
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`; 4.10
V	Log upper bound on the squared covolume of a lifted family, uniform in its coefficient conjugates.	11
$\lambda_{d,M}(\nu)$	$\mathbf{1}_{\nu \equiv 1 \bmod M}\Lambda(\nu)(1 - \chi_K(\nu))$.	`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`; 6
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`; 17
$Z(\theta)$	Squared volume of the row-chain map at a twist, from its Gram determinant in the given metric.	11
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`; 13
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`; 15
$\star$	Main inequality $m\mathcal{R} \le fa(\log b)^5$.	`star_inequality`; 18

19.1 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 20

Mathematical status and proof dependencies

20.1 What has changed at monograph level

The distinction between a formal proof and a conventional mathematical proof is essential for an argument of this size. The compact development reports that its final theorem is verified with the standard Lean axioms; it does not follow merely from that report that every statement in an expository translation is correct. Conversely, a specialty result whose precise conventional proof has not been supplied here is not an extra axiom asserted to have been used by the formal proof. It is a gap in the reconstruction. The following table records the divisions of mathematical work in this volume. The first column refers to the relevant mathematical chapters. It is not intended as an assertion that the table exhausts the full transitive closure of all the formal results.

Area	Derivations presented in this volume	Further work in the conventional proof
Analytic reduction (2)	Smoothed transform and residues; quadratic-character mean square and zero detector; Landau positivity; high-moment window; fixed-modulus deduction.	Classical nonvanishing, character facts and prime distribution may remain standard references. The derivations of the specialized analytic reductions use no GRH assumption.
Admissible data and coefficient model (3, 4)	Behrend sphere, field and different, prime windows using Siegel–Walfisz, class-number finiteness, free level, class-sum measures, explicit Gram and Fischer identities, weak approximation and local gauges.	The standard integral basis theorem is used; establishment of the common Hermitian frames, Weil data and compatible global normalizations remains necessary.
CM theta kernels (6, 5)	Regularized formal parameter and jet normal forms, null-cone and coefficient calculations, Schur determinant losses, conversion into the local weighted gain, with additional CM normalizations and frame analysis.	Some arithmetic CM and Weil-vector normalization or reciprocity assertions still depend on precise source inputs; their global uniformity must be checked as stated.

Area	Derivations presented in this volume	Further work in the conventional proof
Hecke and seesaw (7, 8)	Root and interpolation arguments, finite integral filtrations, counting/elementary divisors and local repair, exact index and norm cancellation.	Geometric existence of compatible chain data and the genuine iterated seesaw (not merely a formal Fubini identity), and any uniform integral kernel or conjugate-size estimates not established in the two chapters.
Spanning (9, 10)	Finite good-type counting, operator/greedy rank selection, Schur and kernel estimates, local-to-global analytic and sieve reductions.	Quantitative construction of the large-rank Bessel systems, exact global-to-local identifications, as well as other explicitly retained specialized real-place/filling estimates. Ordinary qualitative density does not replace these.
Determinant and periods (11, 12)	Mixed conjugate Gram determinants, height algebra, one determinant and compressed determinants. Finite complementary-frame matching, Fourier/doubling calculations and further level normalizations.	Global complementary-rank period exchange in the genuine model, nonvanishing and convergence, normalized Rallis injectivity/period identities, false-lift and comparison-space assertions not reconstructed.
Selmer and congruences (13, 14)	Integral projection formulation, polynomial and Fitting arguments, auxiliary Kummer/triangle calculations, first-order and integral finite algebra, Tate-string budgets.	Weak attached Galois representations with exact Hecke compatibility; row existence and balance, global duality/extraction, simultaneous twisting and odd-class bounds, and the remaining actual automorphic patching and Ribet inputs.
Crystalline losses (15, 16)	Rank and colength calculations, explicit Kummer-coordinate and exponent arithmetic, safe pairing and window reductions, summation of loss across blocks.	Uniform integral frame comparison with crystalline cocycle modules, period-ring identities, and remaining specialized window/spacing/length theorems at the stated quantitative strength.
Assembly and numerical budgets (18, 17)	Exact prime remainder identity, cancellation of the common leading term, budget sums along the asymptotic family and the final zero-free implication.	No additional arithmetic results beyond the stated inputs to A, B and C, their coherent common data, and the classical analytic facts.

The principal new proofs compared to the shorter specialist document are concentrated in the technical companion chapters included within each part. The statements used as premises in a main exposition sometimes retain their formal-source names to make the logical dependency transparent; when a particular constituent step is proved later in the same volume, the relevant chapter supplies its proof. The entire package should be considered established by conventional arguments only to the extent that all of its preceding stated inputs have been justified.

20.2 Priority for the remaining reconstruction

The highest-priority mathematical obligations are the following groups of exact formal statements. They are not merely recommendations for extra exposition. Each is necessary to recognize that the constructions used in the estimates are the same objects, with the same integral and metric structures, throughout the proof.

1. In A, the overall map in `Data.seesaw` and its control of finite datum support, integral stabilizer and lattice changes: see `New/Seesaw/SeesawTheorem.lean`, `UnfoldingStages.lean` and `GoodDatumExists.lean`. The CM coefficient data and the resulting global estimates must agree with these choices, particularly at nonsplit primes where the arithmetic gain is measured.

2. In B, the normalized global filling datum used in `New/Filling/FillingData.lean`, and the full quantitative interface between the real peak/cut estimates and the global energy sum in `New/Filling/Moments.lean` and `GlobalSieve.lean`. The error b^{-4} in the precise operator bounds is essential. A statement that Bessel coefficients decay rapidly for any fixed parameters does not establish the required uniformity.

3. In C, the actual complementary-rank theta exchange and the comparison of volume ratios, including the conversion of the auxiliary periods to the normalized Hecke-value factors, not merely the finite grid combinatorics. Principal anchors are `New/Exchange/`, `New/Rows/RowVolume.lean`, `LiftVolumeHeight.lean`, and `HeckeValueHeight.lean`. These are the identities that produce the common sign and coefficient H.

4. Also in C, `exists_isFitFamily` in `New/Selmer/FitTwistsExist.lean`. Its constituent results comprise the weak Galois attachment in `Lib/GaloisRep/`, the simultaneous twisting and the actual congruence-to-cohomology argument in `New/Selmer/` and `New/Patching/`, and the exact integral local comparison in `New/Lift/`, `Lib/PadicHodge/` and `Lib/PadicComparison/`. They must give precisely the rank-dependent denominator estimates of the fit definition. Qualitative crystalline or representation-theoretic assertions without quantitative lattice control are not enough.

A draft that contained fewer deep assertions only because they had been collected into larger named theorems would not constitute progress. The reason for the new technical chapters is to separate proofs of elementary algebra and parameter-dependent inequalities from genuinely nonstandard geometric or arithmetic inputs. In a further full-length exposition, each of the groups above should be developed into a sequence of independently checkable theorems, with verified references for the applicable classical results and actual proofs of their specialized refinements. It is possible that more compact mathematical arguments can replace some elaborate formal implementations, but they must be shown to produce exactly the same required bounds.

20.3 The use of published mathematics

The bibliography records classical material expressly used or discussed in this volume, not a complete audited source list of the entire Lean library. Attribution of a theorem to a published paper does not justify a quantitative strengthening that is absent from its hypotheses. For example, the standard algebraic modular form theory gives finiteness and the lattice framework, and the elementary proofs in Chapter 4 give additional details in the present normalizations. This does not establish the quantitative mass and theta-integral formulas at all the special

levels. Similarly the existence and comparison theorems of rational p-adic Hodge theory do not by themselves prove the integral cokernel length bounds in a family of torsion lattices with weights growing with x.

The compact repository's `REFERENCES.md` reports 196 citations compared to varying extents: 57 compared directly and precisely located, 30 locations identified at second hand, 102 works found without a precise location, and 7 not compared. Other files' citations were outside that check. The source list helps a future independent proof identify the intended literature, but is not an audited bibliography of theorem matches. The readable edition's `INDEX.md` and `GUIDE.md` further distinguish background and novel pieces; the distinction between `Lib` and `New` should be checked whenever a statement is replaced by a citation.

The closing inequality of this monograph is elementary given the three arithmetic theorems. The difficult problems concern how to prove those theorems in a form that survives their simultaneous use. With this separation made explicit, the remaining verification is a finite list of substantial mathematical programs rather than an unstructured translation of tens of thousands of lines. This is the intended contribution of the present volume.

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.