Series · Markdown source

III. Heights and the square-root exponent

1. The column and row lattices
2. The slope bound and the main term
   Integrated refinement of the frame bound
3. Reducing the height error to a size bracket
4. From the arithmetic gap to the character sum
5. Division by the normalizer
6. The fixed-level window input
7. The geometric inputs to be established in prose

The arithmetic problem is to bound a weighted prime-power sum by its main term plus a controlled error. The proof compares the slope of the same row lattice with the slopes of a family of column lattices. At a fixed translate, the columns are counted with their actual ranks. The resulting inequalities can then be averaged with equal weights over translates. For the fixed-level extension we use the pointwise version before averaging over the subgroup. All source references are to the release at commit **e7c8972c701d3396947808ab20cc16dc14623d33**. Word and translate means · Mean and pointwise bounds · Fixed-level comparison

In rStation paper 192 v1, span_of_dilates derives the CM balance identity and explains how a hypothesized bound on the arithmetic gap yields the prime estimate at one host. Its Theorem 4.9 invokes `ch_top` for that bound, which appears among the obligations in §4.6. The additional derivation begins in §1, where the finite-place indices give the adjusted column slopes. Section 2 compares their rank-weighted mean with the slope of the common dual row lattice through the integral period maps and the frame rank estimates. The signed height estimate cancels the common divisor term, and §3 uses the finite-prime loss bounds to obtain the same bracket. The comparison is conditional on fitted data at the same cutoff. Sections 4–6 use the shared arithmetic readout and square-root normalization, with a uniform mean over translates also available in the ordinary case under mean loss bounds. Paper 192 v1, Theorems 4.7 and 4.9 and Propositions 4.12–4.13 · Rank filtration · Gain comparison · Pointwise gap bound

The integrated refinement in §2 gives a linear spread term and a centered Gram error. With the general size bounds, §§3, 5 and 6 give a one-sided estimate of order $x^{1/2+\nu\varepsilon}$ in the ordinary and fixed-level cases, where $\nu = \max(3, 2 + e_L, 21/8 + e_D)$. The original comparison also appears throughout.

1. The column and row lattices

Fix a finite nonempty set U and, for each $\alpha \in U$, a finite nonempty word set W_α. A translate depends on α. In the full translate calculation $U = H$, a finite abelian group. In the fixed-level calculation there is a map $\mathrm{tr} : U \to H$ whose image contains the required subgroup. Write $m \geq 1$ for the number of lifts, TT for the set of tops, $n_T = |\mathrm{TT}|$, $a_t \geq 0$ for their sizes and $A = \sum_t a_t$. The common size parameters are $x, \varepsilon, a, b, b', k, L, L_a, B_0, N_c, X''$. Here a is the reference size, b the reference thick weight, b' the thin

scale, k a size parameter for the lattices, and L, L_a logarithmic bounds. Fix $f > 0$ and positive weights $w_{\alpha j\rho}$ for $1 \le j \le m$ and $1 \le \rho \le f$. Scalar data · Lattice specialization · Translate coverage

At a fixed α, the coefficient field K is totally complex, and each lattice is a finitely generated $\mathcal{O}_K$-module spanning its ambient K-space, with a positive Hermitian metric at every infinite place. Put $n_K = [K : \mathbf{Q}]$. In the release's normalization the slope $\mu(M)$ is $2\,\mathrm{slope}(M) - \log 2 + n_K^{-1} \log|\operatorname{disc}(K)|$, where $\mathrm{slope}(M)$ is the negative log of the Euclidean covolume divided by the rational rank. Apply the same affine change to the largest and smallest Harder–Narasimhan slopes. The normalized degree is $\deg M = \mathrm{rk}_K(M)\mu(M)$. For a line $\mathfrak{a}z$ its degree is

$$-\frac{2}{n_K}\log N\mathfrak{a} - \frac{1}{n_K}\sum_{v|\infty}[K_v : \mathbf{R}]\log\|z\|_v^2.$$

Classical: the determinant-line formula for degree (with the displayed normalization) and duality of Harder–Narasimhan slopes give $\mu(R^\vee) = -\mu(R)$ and $\mu_{\min}(R^\vee) = -\mu_{\max}(R)$ with dual metrics. Write $C = R^\vee$ for the common row lattice, realized isometrically in the retained row coordinates. Slope and map-height conventions · Duality and row covolume

For each word $i \in W_\alpha$ there is a full column lattice M_i^{full}, its version M_i after restriction at the type prime T_i, and a sublattice N_i^0 of a fractional-ideal multiple of M_i. The retained rank $d_i = \mathrm{rk}_K M_i = \mathrm{rk}_K N_i^0$ is strictly positive, and $D_i = \mathrm{rk}_K M_i^{\mathrm{full}}$. For each finite place w, set $\ell_w = \log N\mathfrak{p}_w / n_K$, let e_w be the ramification index, and use $l_w = e_w\ell_w$. The finite set S of contributing places contains every place over each rational prime it meets. Given finitely supported integral τ_w, σ_{iw}, set $I_\nu = \prod_w \mathfrak{p}_w^{-\nu_w}$ for either shift and put

$$N_i = I_{\sigma_i}N_i^0, \qquad N_i^0 \subset I_\tau^{-1}M_i, \qquad \lambda_{iw} = \mathrm{length}_{\mathcal{O}_{K,w}}((I_\tau^{-1}M_i/N_i^0)_w).$$

The use of powers with a negative exponent in I_ν fixes the direction of all the slope shifts. The maps B_i are integral from N_i to C, and h_i bounds the height $\sum_{v|\infty}[K_v : \mathbf{R}]\ \log\|B_i\|_v^2/n_K$ for $B_i \neq 0$. **Supplied by the geometric papers:** the lattices, the integral maps, and the restriction and local estimates below are for one and the same family. Family and local hypotheses · Division and map height

Put $\tau_w^* = \tau_w/e_w$, $s_{iw}^* = \sigma_{iw}/e_w$ and $\lambda_{iw}^* = (\lambda_{iw} + d_i\tau_w)/e_w$. Let g_{jw} be the local gain, F_{jw} the common divisor term, and E_{ijw} the type prime exponent term. The family supplies

$$\begin{aligned} S_{iw} &= \sum_j (g_{jw} + F_{jw} - E_{ijw}), \quad s_{iw}^* \le S_{iw}, \\ G &= 2\sum_{w\in S} l_w\Big(\sum_j g_{jw} - \tau_w^*\Big), \qquad F_o = 2\sum_{w\in S} l_w \sum_j F_{jw}, \\ E_i &= 2\sum_{w\in S} l_w \sum_j E_{ijw}, \qquad\qquad \mathcal{R}_i = 2\sum_{w\in S} l_w(S_{iw} - s_{iw}^*). \end{aligned}$$

Here G is the source's Ga, and E_i its Ee. Both G and F_o depend on the translate alone. Write $I_O = \sum_{p\in p(S)} \log p\, \iota(p)$ for the other finite-place charge, with $\iota(p) \ge 0$, and $\mathcal{T} = 2\sum_w l_w\tau_w^*$. The identity $\sum_{w|p} l_w = \log p$ follows from $\sum_{w|p} e_w f_w = n_K$, with f_w the residue degree. Local definitions · Arithmetic identification

Positive-rank comparison (proved here from the local inputs). A full-rank inclusion over a discrete valuation ring has determinant valuation equal to the length of the quotient. Applying this at every finite

place, and using the ideal twist formula for degree, gives

$$\begin{aligned}\mu(N_i) &= \mu(M_i) + S_i^{\mathrm{lat}} - B_i^{\mathrm{len}},\\ \mu_{\max}(N_i) &\le \mu_{\max}(M_i) + S_i^{\mathrm{lat}},\\ S_i^{\mathrm{lat}} &= 2\sum_w l_w(s_{iw}^* - \tau_w^*) = G + F_o - E_i - \mathcal{R}_i,\\ B_i^{\mathrm{len}} &= \frac{2}{d_i}\sum_w l_w(\lambda_{iw}^* - d_i\tau_w^*).\end{aligned}$$

For the maximum slope, each sublattice of N_i lies in $I_{\sigma_i}I_\tau^{-1}M_i$. At a fixed rank its saturated enlargement in that lattice raises degree by the nonnegative index term, while the ideal twist moves every slope by S_i^{lat}. This proves the second line. The positive d_i is also necessary in the first line. For the floor shift $\sigma_{iw} = \lfloor e_w S_{iw}\rfloor$, the contribution to $\mathcal{R}_i$ is exactly $2\ell_w\{e_w S_{iw}\}$, including when $S_{iw} < 0$. Length and slope conversion · Local gain identity · Conversion to place sums · Prior height calculation

Supplied by the geometric papers. For all words, $T_i \ge 4f$, $T_i > 2$, $\mathcal{T} \le b^2$, and $\tau_w^* \ge 0$. At $w \mid T_i$ one has $\tau_w^* = 0$ and $\lambda_{iw}^* \le D_i\lambda_i$, with $\lambda_i \ge 0$. Off these places, $\lambda_{iw}^* \le D_i(\tau_w^* + \iota(p(w)))$. The rank identity and the loss in passage from the full lattice are

$$\frac{d_i}{D_i} = \Big(1 - \frac{1}{T_i - 1}\Big)^f, \quad \mu(M_i) = \mu(M_i^{\mathrm{full}}) - \frac{f\log(T_i-1)}{T_i - 2}, \quad \mu_{\max}(M_i) \le \mu_{\max}(M_i^{\mathrm{full}}).$$

The estimates also give $4fb^2/T_i + f\log(T_i-1)/(T_i-2) \le 1$, $4\lambda_i \le c_{\mathrm{Sl1}} n_T m(a/k + b')$, $\mathcal{R}_i \le c_{\mathrm{Sl1}}X''$, and $c_{\mathrm{Sl1}} \ge 4$. Restrictions and length bounds · Size majorants

Indeed $d_i/D_i = 1 - \xi_i$ with $0 \le \xi_i \le f/(T_i-1) \le 2f/T_i \le 1/2$, by Bernoulli's inequality. The part of B_i^{len} at T_i is at most $4\lambda_i \log T_i$. The other part is at most $2I_O/(1-\xi_i) + \xi_i\mathcal{T}/(1-\xi_i) \le 4I_O + 4fb^2/T_i$. With the preceding loss of slope this proves the following inequalities, where $\mu_{M,i}, \mu_{M,i}^{\max}$ refer to M_i^{full}, and $\mu_{N,i}, \mu_{N,i}^{\max}$ to N_i:

$$\begin{aligned}\mu_{N,i} &\ge \mu_{M,i} + G + F_o - E_i - \mathrm{Lo}_i,\\ \mu_{N,i}^{\max} &\le \mu_{M,i}^{\max} + G + F_o - E_i + c_{\mathrm{Sl1}} m B_0,\\ \mathrm{Lo}_i &= c_{\mathrm{Sl1}} I_O + c_{\mathrm{Sl1}} n_T m(a/k + b')\log T_i + c_{\mathrm{Sl1}}X'' + 1.\end{aligned}$$

The upper inequality holds already with zero in place of $c_{\mathrm{Sl1}} m B_0$ under the local hypotheses just used. We keep the released scalar allowance. Positive-rank scalar comparison · Transfer of the bounds · Word loss

The selected family packages the same calculation as a place-sum input. In that notation the length contribution at T_i is supplied directly with bound $P_{T,i} \le c_{\mathrm{Sl1}} n_T m(a/k + b')\log T_i$, the nonnegative rounding charge is at most $c_{\mathrm{Sl1}}X''$, and the weighted sum of the other local allowances before dividing by the retained share is $2I_O$. The bound $0 \le \xi_i \le 2f/T_i \le 1/2$, the restricted-slope comparison, and the allowance of 1 give exactly the same Lo_i by the preceding argument. The geometric supply for this route can give these place-sum hypotheses directly. Selected place estimates · Selected scalar bounds

2. The slope bound and the main term

The constants $c_{\mathrm{Div}}, c_E, c_H, c_{\mathrm{Count}}, c_{\mathrm{Slope}}, c_{\mathrm{Sl1}}, c_{\mathrm{Loss}}, c_{\mathrm{HH}}, e_{\mathrm{Div}}, e_{\mathrm{Loss}}$ are nonnegative. Write γ for Euler's constant and $\kappa = \operatorname{Re}\ell(\chi_C)$ for the centering constant of the first sign character defined in §4. For positive

v, set $\omega = \gamma + \log(\pi/v) - \kappa$, and set $p(a_t) = a_t \log a_t - a_t$, continuously at zero. Define

$$\begin{aligned} S_{\mathrm{Div}} &= m \sum_t (p(a_t) - a_t \omega), \\ E_{\mathrm{Div}} &= c_{\mathrm{Div}} n_T a L_a + c_{\mathrm{Div}} m (B_0 + N_c) x^{e_{\mathrm{Div}} \varepsilon}, \\ J &= c_{\mathrm{Count}} L_a (1 + m^2/k), \\ \Lambda_\alpha &= \sum_j \frac{1}{f} \sum_\rho \log(w_{\alpha j \rho} v/\pi) + m\, \ell_g - J, \\ P_{\mathrm{Loss}} &= c_{\mathrm{Loss}} n_T m b x^{e_{\mathrm{Loss}} \varepsilon}. \end{aligned}$$

The number ℓ_g is the sum of the logarithms of the two comparison primes g_0, g_1 in the constructed family. These are common across words, whereas T_i is the prime used for the restriction of word i. The live constants set $c_{\mathrm{Div}} = c_E = c_H$, with the values of $c_{\mathrm{Count}}, c_{\mathrm{Slope}}, c_{\mathrm{Sl1}}, c_{\mathrm{Loss}}, c_{\mathrm{HH}}$ furnished by their respective estimates. Main and losses · Prime and count formulas · Choice of constants

Supplied by the geometric papers. At every translate and for every word, the following inequalities hold. In the pairwise inequalities both words belong to the same W_α.

$$\begin{aligned} \mu(C) + h_i - F_o &\le S_{\mathrm{Div}} - A \log T_i + E_{\mathrm{Div}}, \\ |E_i - m A \ell_g| &\le c_E m B_0, \\ |(h_i + A \log T_i) - (h_{i'} + A \log T_{i'})| &\le c_H (n_T a L_a + m B_0), \\ \Lambda_\alpha \le \log T_i \le c_{\mathrm{Count}} m L, &\qquad |\log T_i - \log T_{i'}| \le c_{\mathrm{Count}} \log k, \\ -c_{\mathrm{Slope}} n_T a L_a (1 + m/k) \le \mu_{M,i}, &\qquad \mu_{M,i}^{\max} \le c_{\mathrm{Slope}} n_T a L_a. \end{aligned}$$

The first inequality is a bound for the dual row slope and the map height. In the divisor assembly the upper bound for the base slope contains $+2F$ in the notation of that assembly. The comparison subtracts $2F$ on the left, and the scalar F_o is precisely that $2F$. The same finite sum entered the column comparison with $+F_o$. At the divisor stage the $-\sum_t a_t \log T_i$ term is the type-prime part of the height. The other top contributions are obtained from $p(a_t) - a_t \omega$ once per lift. The geometric papers owe the estimates of the local indices, the Gram determinants, and the archimedean and factorial terms that give these bounds. Divisor assembly · Normalization and transfer · Height and full-lattice slope bounds

The algebra of that first inequality can be seen before the error is absorbed. In the assembly write the height as a type-prime term plus the lift heights, less the archimedean main term and the covariance normalization. The latter normalization cancels the same term in the upper bound for the base dual slope. For each lift and position ρ, the lower estimate of the archimedean term is subtracted from the lift height's factorial term, leaving at most $\mathbf{1}_{\rho \in \mathrm{TT}} p(a_\rho)$ plus its error. The degree multiplying $-\omega$ is $\mathbf{1}_{\rho \in \mathrm{TT}} a_\rho + k_* w_{j\rho} + \Delta_{j\rho}$ in the assembly, where k_* is its degree parameter. The corresponding base slope term is $+\omega\, p_j w_{j\rho}$. Their difference leaves $-\mathbf{1}_{\rho \in \mathrm{TT}} a_\rho \omega$ and costs at most $|\omega|(|p_j - k_*|\,|w_{j\rho}| + |\Delta_{j\rho}|)$. Thus the exact remaining majorant there is the sum of the base-to-row and base errors, the lift-height errors, the archimedean and type-prime errors, the factorial comparison errors, and those two degree discrepancies. The geometric input bounds that sum by E_{Div}. Assembly with each error · Selected height and count inputs · Prior reading of the signed height

Supplied by the geometric papers. At each translate take matrices $A_i : \mathbf{C}^{d_i} \to \mathbf{C}^{D_{\mathrm{all}}}$ and positive scalars c_i. The projections of A_i to the retained rows are the matrices of B_i at a common embedding in

fixed invertible changes of coordinates. Write D_{ret} for the dimension of this retained space, with $D_{\mathrm{ret}} \geq (1-\delta_\alpha)D_{\mathrm{all}} > 0$. For $Q_A = \sum_i c_i A_i A_i^*$ the frame conditions are

$$\sum_i \|c_i A_i^* A_i - I_{d_i}\|_{\mathrm{HS}}^2 \leq e_\alpha \sum_i d_i, \qquad \operatorname{tr}(Q_A^2) \leq (1+\eta_\alpha)\,(\operatorname{tr} Q_A)^2/D_{\mathrm{all}}.$$

The Hilbert–Schmidt norm is the square root of the sum of squared matrix entries in orthonormal bases. The matrix identity is $\Pi A_i = Y^{-1}(B_i)_{\mathbf{C}} Y_i$, with retained projection Π, invertible source changes Y_i, and a single invertible target change Y common to the words. These changes can be nonisometric. In particular $D_{\mathrm{ret}} = \operatorname{rk}_K C$. The parameters $\delta_\alpha, \eta_\alpha, e_\alpha$ lie in $[0,\bar{\delta}]$, $[0,\bar{\eta}]$, $[0,\bar{e}]$, respectively, with $\bar{\delta} < 1$ and $\bar{e} \leq 1/4$. Frame operators · Realization and inequalities

Define the rank mean and the spread

$$\langle Z \rangle_\alpha = \frac{\sum_{i\in W_\alpha} d_i Z_i}{\sum_{i\in W_\alpha} d_i}, \qquad s = \langle \mu_{N,i} - h_i \rangle_\alpha, \qquad u = \max_i(\mu_{N,i}^{\max} - h_i) - s.$$

Then $u \geq 0$. The geometric height input also gives $H_* > 0$ with $s - \mu_{\min}(C) \leq H_* \leq x^{c_{\mathrm{HH}}} b$. For $t \in [0, H_*]$, let $q(t)$ be the fraction of all column slope dimensions with adjusted slope $< s - t$ and $p(t)$ the fraction of retained row slope dimensions with slope $< s - t$. These fractions can be taken in K-dimensions because the same factor n_K converts them to rational ranks. Live mean and spread · Scalar height window

Proved here. The high Harder–Narasimhan step of N_i at $s - t + h_i$ maps into the high step of C at $s - t$. Indeed the minimum slope of the domain step is at least $s - t + h_i$, and the maximum slope of the quotient of C by its high step is strictly less than $s - t$. A nonzero composite would give a positive-rank quotient of the domain step mapping injectively to a saturated sublattice of the target quotient. The two induced metrics make the composite's squared-norm height at most h_i, since the subspace uses the restricted metric and the quotient uses the minimum norm of a lift. The determinant inequality bounds the slope of the domain quotient above by the slope of the saturated image plus h_i, contradicting the two slope bounds. This uses the classical Harder–Narasimhan characterization of extremal slopes by saturated sublattices and quotients. The inclusion holds also when the high step or the composite map is zero. Slope and height conventions · Threshold argument in the release

At the chosen embedding let P_i be the orthogonal projection onto the transported high column step. Put $D_{\mathrm{col}} = \sum_i d_i$, $Q_{\mathrm{high}} = \sum_i c_i A_i P_i A_i^*$, and $T = \operatorname{tr} Q_A$. Cauchy–Schwarz and the Gram bound give $T \geq (1 - \sqrt{e_\alpha}) D_{\mathrm{col}} > 0$. In fact the trace pairing of the Gram deviations with the identities has modulus at most $\sqrt{e_\alpha} D_{\mathrm{col}}$. Pairing instead with $I - P_i$ bounds the discarded trace by $D_{\mathrm{col}}(q(t) + \sqrt{e_\alpha q(t)})$. Thus the discarded trace fraction $r_t = 1 - \operatorname{tr}(Q_{\mathrm{high}})/T$ is at most $2(q(t) + \sqrt{e_\alpha})$. Since $0 \leq Q_{\mathrm{high}} \leq Q_A$, $\operatorname{tr}(Q_{\mathrm{high}}^2) \leq \operatorname{tr}(Q_A^2)$. The trace Cauchy–Schwarz inequality gives $\operatorname{rk}(Q_{\mathrm{high}})/D_{\mathrm{all}} \geq (1-r_t)^2/(1+\eta_\alpha)$. The image of Q_{high} is the span of the high column images. Its codimension cannot increase under projection to the retained rows. The image inclusion just proved then gives the dimension estimate

$$(1-\delta_\alpha)p(t) \leq \eta_\alpha + 4\sqrt{e_\alpha} + 4q(t).$$

Here the bound on the ambient codimension fraction uses $1 - (1-r_t)^2/(1+\eta_\alpha) \leq \eta_\alpha + 2r_t$. Threshold comparison

Proved here. The rank mean of the adjusted slopes is s, and every adjusted slope is at most $s + u$. For $t > 0$ the mean is at most $q(t)(s - t) + (1 - q(t))(s + u)$, so $q(t) \le u/(u + t)$. Since every row deficit below s is at most H_*, $s - \mu(C) \le \int_0^{H_*} p(t)\, dt$. Integrating the dimension bound gives a loss at most

$$\frac{4u \log(1 + H_*/u) + (\eta_\alpha + 4\sqrt{e_\alpha})H_*}{1 - \delta_\alpha}.$$

The first term is zero at $u = 0$, when every adjusted slope is s. This implies the source's scalar comparison, which also allows $\log |W_\alpha|$ as a nonnegative loss. Let $X = x^{c_{\mathrm{HH}}}$, and suppose $u \le U_*$. The inequality $\log z \le z - 1$ applied to $(1 + H_*/u)/(1 + X)$ gives

$$u \log(1 + H_*/u) \le u \log(1 + X) + \frac{H_*}{1 + X} - \frac{uX}{1 + X} \le (U_* + b) \log(1 + X).$$

The last step uses $H_* \le Xb$ and $X/(1 + X) \le \log(1 + X)$. The same bound holds at $u = 0$. Scalar frame comparison · Logarithmic estimate

Integrated refinement of the frame bound

The earlier same-family calculation integrated the actual source deficits, and the frame-energy note centered the Gram operator by its actual trace. The following proof applies both estimates at the filtration thresholds. The high step at a normalized threshold z uses the last grade of slope at least z, with index zero if no grade qualifies. The corresponding Euclidean threshold is $(z - c_K)/2$, with $c_K = -\log 2 + n_K^{-1} \log |\operatorname{disc} K|$. Thus ties belong to the high step and both low counts are strict. Integrated deficits · Centered Gram estimate · Source threshold comparison

Fix $t \in [0, H_*]$. Use $q = q(t)$, $r = r_t$, and $c = \operatorname{codim}(\operatorname{im} Q_{\mathrm{high}})/D_{\mathrm{all}}$. Put $G_i = c_i A_i^* A_i$, $g = T/D_{\mathrm{col}} > 0$, $e = e_\alpha$, and $\eta = \eta_\alpha$. Expanding the square and using $\sum_i \operatorname{tr} G_i = T$ gives

$$\sum_i \|G_i - g I_{d_i}\|_{\mathrm{HS}}^2 = \sum_i \|G_i - I_{d_i}\|_{\mathrm{HS}}^2 - (1 - g)^2 D_{\mathrm{col}} \le (e - (1 - g)^2) D_{\mathrm{col}}.$$

The blocks $(I_{d_i} - P_i) - q I_{d_i}$ have total squared Hilbert–Schmidt norm $D_{\mathrm{col}} q(1 - q)$. Their trace pairing with the centered Gram blocks is $T(r - q)$, since the sum of the centered Gram traces is zero. Divide Cauchy–Schwarz by T and use

$$\frac{e - (1 - g)^2}{g^2} = \frac{e}{1 - e} - \frac{(g - (1 - e))^2}{(1 - e) g^2}$$

to get $r \le q + \sqrt{e' q (1 - q)}$, where $e' = e/(1 - e)$.

Let $J_\perp$ project orthogonally onto the complement of the high frame span. Since Q_{high} is supported on that span and $Q_{\mathrm{high}} \le Q_A$, we have $\operatorname{tr}(Q_A J_\perp) \le rT$. The trace-zero operator $Q_A - (T/D_{\mathrm{all}}) I$ has squared Hilbert–Schmidt norm at most $\eta T^2/D_{\mathrm{all}}$. Pairing it with $J_\perp - cI$, of squared norm $D_{\mathrm{all}} c(1 - c)$, shows that $c - \sqrt{\eta c(1 - c)} \le r$. The high column spaces are transported by Y_i^{-1} before forming their orthogonal projectors. Their images after row projection lie in the high step of C transported by the common Y^{-1}. Scalar extension preserves rank, and row projection reduces the dimension by at most $D_{\mathrm{all}} - D_{\mathrm{ret}}$. Hence $p(t) D_{\mathrm{ret}} \le c D_{\mathrm{all}}$. These dimension comparisons hold for the given invertible basis changes.

Use $2\sqrt{\eta c} \le c + \eta$ and $2\sqrt{e' q} \le q + e'$ to obtain $(1 - \delta_\alpha) p(t) \le c \le 2r + \eta \le 3q + \eta + e'$. An adjusted column slope λ contributes $\min(H_*, (s - \lambda)_+)$ to the integral of the strict low indicator. The rank mean of $(s - \lambda)_+$ equals that of $(\lambda - s)_+$, which is at most u. The integral of $p(t)$ is the mean positive row deficit, since $s - \mu_{\min}(C) \le H_*$. Integrating gives

$$(1 - \delta_\alpha)(s - \mu(C)) \le 3u + \left(\eta_\alpha + \frac{e_\alpha}{1 - e_\alpha} \right) H_*. \tag{F}$$

This includes zero spread and empty steps.

Subtract h_i from the lower column bound. The divisor inequality gives
$F_o - h_i \ge \mu(C) - S_{\mathrm{Div}} + A \log T_i - E_{\mathrm{Div}}$, which cancels the common divisor term. Apply the lower bound for the full-column mean, the upper bound $E_i \le mA\ell_g + c_E m B_0$, and $\log T_i \ge \Lambda_\alpha$, with $A \ge 0$. For the main term use the exact identity

$$\begin{aligned} M_\alpha &= \sum_{t,j} a_t \left(\log a_t - 1 - \gamma + \kappa - \frac{1}{f} \sum_{\rho} \log w_{\alpha j \rho} \right), \\ S_{\mathrm{Div}} + mA\ell_g - A\Lambda_\alpha &= M_\alpha + AJ. \end{aligned}$$

In the second line $mA\ell_g$ cancels the type-prime term, and the two terms involving $\log(\pi/v)$ cancel each other. Averaging over words with the positive weights d_i gives

$$s \ge G + \mu(C) - M_\alpha - K_0 - \langle \mathrm{Lo} \rangle_\alpha, \qquad K_0 = AJ + E_{\mathrm{Div}} + c_E m B_0 + c_{\mathrm{Slope}} n_T a L_a (1 + m/k).$$

Main-term identity and lower mean

To bound u, subtract the lower bound for $\mu_{N,i'} - h_{i'}$ from the upper bound for $\mu_{N,i}^{\max} - h_i$ before taking the mean in i'. Both G and F_o cancel. The variation of $E_i + h_i$ is bounded by

$$V = 2c_E m B_0 + c_H (n_T a L_a + m B_0) + A c_{\mathrm{Count}} \log k,$$

by subtracting and adding the common center $mA\ell_g$ and the two terms $A \log T_i$. Thus

$$\begin{aligned} u &\le K_3 + \langle \mathrm{Lo} \rangle_\alpha, & K_3 &= c_{\mathrm{Slope}} n_T a L_a (2 + m/k) + V + c_{\mathrm{Sl1}} m B_0, \\ \langle \mathrm{Lo} \rangle_\alpha &\le c_{\mathrm{Sl1}} I_O + K_7, & K_7 &= c_{\mathrm{Sl1}} n_T m (a/k + b')(c_{\mathrm{Count}} m L) + c_{\mathrm{Sl1}} X'' + 1. \end{aligned}$$

The maximum over i leaves the estimate unchanged because the right side after the mean in i' is independent of i. Variation and spread · Count bound for the mean loss

Put $D_* = (1 - \bar{\delta})^{-1}$, $J_* = 4 \log(1 + x^{c_{\mathrm{HH}}})$ and $E_* = (\bar{\eta} + 4\sqrt{\bar{e}}) x^{c_{\mathrm{HH}}} b$. Substitute the bound for u into the row comparison and use the lower estimate for s. The coefficient of $\mu(C)$ is one on each side. This gives, for each α,

$$G - M_\alpha \le \log |W_\alpha| + K_0 + \langle \mathrm{Lo} \rangle_\alpha + D_* \big[J_* (K_3 + \langle \mathrm{Lo} \rangle_\alpha + b) + E_* \big].$$

Gain comparison

Let ASK_α be the prime-power sum for a translate, defined in §4. **Supplied by the geometric papers:** the finite-prime loss bounds control $\mathrm{ASK}_\alpha - G$ and I_O each by P_{Loss}, pointwise for the translates used in the fixed-level route. For the mean theorem it suffices that these same bounds hold on their two uniform means

over a nonempty set U. Put $M_{\mathrm{Lo}} = c_{\mathrm{Sl1}} P_{\mathrm{Loss}} + K_7$ and $U_{\mathrm{bar}} = K_3 + M_{\mathrm{Lo}}$. Adding $\mathrm{ASK}_\alpha - G$ and using the nonnegative coefficient $1 + D_* J_*$ of $\langle \mathrm{Lo} \rangle_\alpha$ gives the pointwise inequality

$$\mathrm{ASK}_\alpha - M_\alpha \le \mathrm{Err}_\alpha := K_0 + \log |W_\alpha| + P_{\mathrm{Loss}} + M_{\mathrm{Lo}} + D_* \big[J_*(U_{\mathrm{bar}} + b) + E_* \big].$$

Under mean loss hypotheses the same computation gives $\frac{1}{|U|} \sum_\alpha (\mathrm{ASK}_\alpha - M_\alpha) \le \mathrm{Err}_{\mathrm{mean}}$, where $\mathrm{Err}_{\mathrm{mean}}$ is the displayed expression with $|U|^{-1} \sum_\alpha \log |W_\alpha|$ in place of $\log |W_\alpha|$. This is the uniform translate mean when $U = H$. In the fixed-level route the comparison is pointwise on U, and each subgroup translate is assigned a preimage in U before averaging. Its weight in that average is independent of the number of preimages and the ranks of their words. Loss hypotheses and errors · Uniform mean · Pointwise consumer

3. Reducing the height error to a size bracket

Supplied by the geometric papers. For the same data, the size conditions used here are $x \ge 1$, $0 < \varepsilon \le 1$, $b \ge 1$, $xb \le a$, $n_T \ge 1$, $xm \le k$, $1 \le k \le \min(a, x^{e_K})$, $e_K \ge 0$, $m \le x^{1/2+\varepsilon/2}$, $0 \le L, L_a \le 2x^\varepsilon$, $kb' \le a$, $0 \le X'' \le a$, and $0 \le a_t \le 2a$. Also $B_0, N_c \ge 0$, $\log |W_\alpha| \le (k-3) \log 2$ for each α, $\bar{\delta} \le 1/2$, and

$$E_* \le K_0' x^\varepsilon n_T a.$$

Here $K_0' \ge 0$ dominates the seven coefficients $c_{\mathrm{Div}}, c_E, c_H, c_{\mathrm{Count}}, c_{\mathrm{Slope}}, c_{\mathrm{Sl1}}, c_{\mathrm{Loss}}$. The trace and Gram errors in E_* must be controlled after multiplication by $x^{c_{\mathrm{HH}}} b$. The concrete scales are $b = \exp(x^\varepsilon)$, $a = xb$, $X' = \exp(x^{\varepsilon/3})$, $b' = a/X'$, and $X'' = xnX'$, with $n = |H|$. The estimates below use the displayed size inequalities and permit any scales satisfying them. Error windows · Concrete scales

For the integrated refinement put $E_{\mathrm{int}} = (\bar{\eta} + \bar{e}/(1-\bar{e})) x^{c_{\mathrm{HH}}} b$. It suffices to assume $E_{\mathrm{int}} \le K_0' x^\varepsilon n_T a$ for this route. The original budget implies it since $\bar{e}/(1-\bar{e}) \le 4\sqrt{\bar{e}}$ for $\bar{e} \le 1/4$. From (F) and the lower bound for s, the gap at a label is at most

$$K_0 + (\mathrm{ASK}_\alpha - G) + (1 + 3D_*) \langle \mathrm{Lo} \rangle_\alpha + 3D_* K_3 + D_* E_{\mathrm{int}}.$$

The nonnegative coefficient of $\langle \mathrm{Lo} \rangle_\alpha$ and its bound by $c_{\mathrm{Sl1}} I_O + K_7$ allow the two loss hypotheses to be applied pointwise or after uniform averaging. In both regimes the resulting bound is $K_0 + P_{\mathrm{Loss}} + (1 + 3D_*) M_{\mathrm{Lo}} + 3D_* K_3 + D_* E_{\mathrm{int}}$.

Proved here. In this section write $y = x^\varepsilon$, $d = x^{e_{\mathrm{Div}}\varepsilon}$, $r = x^{e_{\mathrm{Loss}}\varepsilon}$, $h = 1 + c_{\mathrm{HH}}$, $g = 1 + e_K$, $K_* = K_0'$. Take $x \ge 1$ with $\log x \le y$, so $\log k \le gy$ and $D_* \le 2$. Put $Q_0 = n_T a$, $Q_1 = Q_0 m^2/k$, $Q_2 = m(B_0 + N_c)$, and $Q = Q_0 + Q_1 + Q_2$. Then $m \le \min(x, k)$, $mb \le a$, $A \le 2Q_0$, and $Q_0 \ge \max(1, b, X'', \log |W_\alpha|)$. Direct substitution gives

$$\begin{aligned} K_0 &\le 10 K_* y (Q_0 + Q_1) + K_* (d+1) Q_2, \\ P_{\mathrm{Loss}} &\le K_* r Q_0, \\ M_{\mathrm{Lo}} &\le K_*^2 r Q_0 + 4 K_*^2 y Q_1 + (K_* + 1) Q_0, \\ K_3 &\le (8 + 2g) K_* y Q_0 + 4 K_* Q_2. \end{aligned}$$

In the M_{Lo} bound use $a/k + b' \le 2a/k$. In the last bound the contribution of V is bounded by $3K_*Q_2 + (2+2g)K_*yQ_0$. For the integrated refinement, $1 + 3D_* \le 7$, $3D_* \le 6$, and $D_*E_{\mathrm{int}} \le 2K_*yQ_0$. The four coefficients of $yQ_0, rQ_0, yQ_1, (1+d)Q_2$ are bounded respectively by $(67+12g)K_* + 7$, $K_* + 7K_*^2$, $10K_* + 28K_*^2$, and $25K_*$. It follows in both loss regimes that the gap is at most

$$\mathscr{E}_{\mathrm{int}} = C_{\mathrm{gap}}[(y+r)Q_0 + yQ_1 + (1+d)Q_2], \qquad C_{\mathrm{gap}} = (68+12g)(1+K_*)^2. \tag{G}$$

For the original comparison also impose $200(1+K_*)^2hg \le y$, with the original budget E_*. Then $\log(1+x^{c_{\mathrm{HH}}}) \le hy$. The same estimates give $U_{\mathrm{bar}} + b \le 16(1+K_*)^2gyrQ$, and $D_*J_*(U_{\mathrm{bar}} + b) \le 128(1+K_*)^2hg\,y^2rQ$. The remaining terms of the pointwise error are bounded by $22(1+K_*)^2hg\,dry^2Q$, by the four inequalities and $D_*E_* \le 2K_*yQ_0$. Absorbing their coefficients with $200(1+K_*)^2hg \le y$ proves both original error bounds

$$\mathrm{Err}_{\alpha},\ \mathrm{Err}_{\mathrm{mean}} \ \le\ x^{(e_{\mathrm{Div}}+e_{\mathrm{Loss}}+3)\varepsilon}[n_T a(1+m^2/k) + m(B_0+N_c)].$$

The hypotheses of this estimate are uniform over words and translates. The threshold depends on the size constants and ε. Absorption in the release · Explicit coefficient inequality

Let $n_2 = |T_2|$ for the set of thick positions and $\Delta_2 = n_2/n$. **Supplied by the geometric papers:** the upper bounds for B_0, N_c, on the same lattice and prime data, are

$$B_0 \le 2nmbx^{\varepsilon}, \qquad N_c \le \frac{n_2kbx^{\varepsilon}}{n(nx^{-\varepsilon})}.$$

If $C_* \ge e_{\mathrm{Div}} + e_{\mathrm{Loss}} + 3$, an upper bound on either arithmetic gap is consequently

$$\mathscr{E} = x^{C_*\varepsilon}\mathcal{B}, \qquad \mathcal{B} = n_T a(1+m^2/k) + m\left(2nmbx^{\varepsilon} + \frac{n_2kbx^{\varepsilon}}{n(nx^{-\varepsilon})}\right).$$

This substitution uses the identification of tops, weights, centering constants and ASK_{α} with their arithmetic versions. It is valid pointwise on any translates covered by the construction. Tied data and budgets · Bracket and mean identification · Pointwise substitution

4. From the arithmetic gap to the character sum

Let H be finite abelian, $n = |H|$, and $\mathcal{G} = (\mathbf{Z}/2)^2 \times H$. Write $\Phi = \{((0,b_2),v) : b_2 \in \mathbf{Z}/2,\ v \in H\}$, so $f = |\Phi| = 2n$. Put $s_0 = ((1,0),0)$, $c_{\mathrm{conj}} = ((1,1),0)$, and

$$\rho_t(g) = c_{\mathrm{conj}} + g + ((0,0),t), \qquad \chi_C(((b_1,b_2),v)) = (-1)^{b_1}, \qquad \psi_C(((b_1,b_2),v)) = (-1)^{b_2}.$$

Translation by $c_{\mathrm{conj}} + ((0,0),t)$ identifies $\mathcal{G} \setminus \Phi$ with Φ. In Paper II's ordinary coordinates $\chi_C = J_d$ and $\psi_C = J_q$. A position ρ is thick at $z \in H$ if $\rho = ((0,1), s+z)$ for some $s \in T_2$. Coordinates · First sign and partner · Second sign and signed sum

Supplied by Paper II and the host construction. For each prime p there is an inertia subgroup I_p and a lift $\mathrm{Fr}_p \in \mathcal{G}$. The split exceptional set contains $2, q_1$ and excludes the distinguished prime d. For primes in that set both I_p and Fr_p lie in Φ. At d, $I_d = \{0, s_0\}$. At every other prime inertia is trivial. For a unitary character ξ of $\mathcal{G}$, write $\xi_{\mathrm{pr}}(p^i)$ for the mean of ξ over $i\mathrm{Fr}_p + I_p$, extended multiplicatively. For the

principal character $\mathbf{1}$ this function is identically one. Let $\ell(\xi)$ be the primitive logarithmic derivative at 1 for $\xi \neq \mathbf{1}$, and use the centered sums

$$G_\xi(y) = \sum_{n' \le y} \Lambda(n')\,\xi_{\mathrm{pr}}(n')(y/n' - 1) + y\,\ell(\xi) \qquad (\xi \neq \mathbf{1}),$$
$$G_{\mathbf{1}}(y) = G_z(y) = \sum_{n' \le y} \Lambda(n')(y/n' - 1) - y\log y + (1+\gamma)y.$$

Here n' ranges over positive integers, Λ is the von Mangoldt function, and G_z is the ordinary notation used in Paper II, independent of the translate z. Host inertia · Primitive coefficients · Ordinary centered sum

For a top t, $A_t = a_t > 0$, and weight $w : \Phi \to (0,\infty)$, define

$$\mathrm{ASK}(t, A_t, w) = 2 \sum_{\substack{p \le N \\ p \text{ unramified} \\ \mathrm{Fr}_p \notin \Phi}} \log p \sum_{\substack{1 \le i \le N \\ i\mathrm{Fr}_p \notin \Phi}} [A_t/p^i - w(\rho_t(i\mathrm{Fr}_p))]_+$$
$$+ \log d \sum_{i=1}^{N} \sum_{g \in (i\mathrm{Fr}_d + I_d)\setminus\Phi} [A_t/d^i - w(\rho_t(g))]_+,$$
$$\mathrm{MAIN}(A_t, w) = A_t\Big(\log A_t - 1 - \gamma + \ell(\chi_C) - f^{-1}\sum_{\rho\in\Phi}\log w(\rho)\Big).$$

Here $[s]_+ = \max(s, 0)$, and the unramified sum excludes the split exceptional set, whose contribution vanishes. We assume $N \ge A_t/w(\rho)$ for every ρ, so terms with $p^i > N$ vanish. At translate z, the sum over t, j of these ASK's is ASK_α for $\mathrm{tr}(\alpha) = z$, and the sum of the real parts of the corresponding MAIN's is M_α. Denote this gap by $\mathrm{gap}(z)$. Prime-power weights and centering · Identification

Proved here. Character orthogonality gives

$$f^{-1}\sum_{\xi \in \widehat{\mathcal{G}}} \overline{\xi(g)}\,\xi_{\mathrm{pr}}(p^i) = \frac{2}{|I_p|}\,\mathbf{1}_{g \in i\mathrm{Fr}_p + I_p}.$$

The complementary cosets are empty at split primes and have exactly one term at d for each i. This proves the prime weights in ASK from Fourier inversion, including the coefficients $2\log p$ and $\log d$. For $g \notin \Phi$ put $w_g = w(\rho_t(g))$. In the Fourier expansion of ASK use $[A_t/n' - w_g]_+ = w_g[A_t/(w_g n') - 1]_+$ and replace each resulting prime sum by its centered version. The principal character contributes the correction $A_t(\log A_t - \log w_g - 1 - \gamma)$. Each other character contributes $-A_t\ell(\xi)$ before summing over g. The mean of $\overline{\xi(g)}$ on the complement of Φ is -1 for χ_C, 1 for the principal character and 0 for every other character. The centering corrections therefore add up to MAIN, including its $+\ell(\chi_C)$. It follows that

$$\mathrm{ASK}(t, A_t, w) - \mathrm{MAIN}(A_t, w) = f^{-1}\sum_{g \notin \Phi}\sum_{\xi} \overline{\xi(g)}\,w_g G_\xi(A_t/w_g).$$

Prime-power expansion · Centering identity

Supplied by the host and weight construction. For every $t \in \mathrm{TT}$, depth j and translate used in the average, put $y_t = A_t/b$. Assume $x \ge 3,\ |y_t - x| \le 1,\ b, b', X' > 0,\ b'X' \le xb,\ \eta_w, \theta \ge 0$. The weight $w_{zj}(\rho)$ is positive, lies in $[b, b + \eta_w + 1]$ at thick positions, and is at most $2b'$ elsewhere. The conditions also give the cutoff N above, $|G_\xi(y)| \le y\theta$ for all ξ and $y \ge X'/3$, and the two cost bounds

$$2f(x+1)\theta \le 1/x, \qquad 6f(x+1)(\eta_w+1) \le b/x.$$

The weight tolerance η_w and the row frame error η_α occur in different estimates. Node and cutoff hypotheses · Thin and cost bounds · Fixed-level data

Classical input and elementary bounds. Use finite abelian character orthogonality and the identity $\log n' = \sum_{d' \mid n'} \Lambda(d')$. The elementary Chebyshev bound $\Psi(Y) = \sum_{n' \le Y} \Lambda(n') \le 3Y$ follows by comparing $\Psi(2j) - \Psi(j)$ with $\log \binom{2j}{j} \le 2j \log 2$ for integer j and summing at powers of two. Indeed each prime power in $(j, 2j]$ contributes one unit to the valuation of the binomial coefficient and all other contributions are nonnegative. For $Y \ge 1$, summing the divisor identity gives $\log(\lfloor Y \rfloor!) = \sum_{n' \le Y} \Lambda(n') \lfloor Y/n' \rfloor$. Bounding the floors and using $\int_1^{\lfloor Y \rfloor} \log u \, du \le \log(\lfloor Y \rfloor!) \le Y \log Y$ gives

$$\log Y - 2 \le \sum_{n' \le Y} \Lambda(n')/n' \le \log Y + 3.$$

The lower bound uses $\int_{\lfloor Y \rfloor}^{Y} \log u \, du \le \log Y$. Differentiate the continuous, piecewise smooth ramps between their breakpoints. Since $|\xi_{\mathrm{pr}}| \le 1$, one obtains for $Y, Z \ge 1$ and $C_d = 3 + |\gamma|$

$$\begin{aligned} |G_z(Y) - G_z(Z)| &\le C_d |Y - Z|, \\ |G_\xi(Y) - G_\xi(Z)| &\le |Y - Z| \big(\log \max(Y, Z) + C_d + |\ell(\xi)| \big) \quad (\xi \ne \mathbf{1}). \end{aligned}$$

Also $wG_\xi(A_t/w)$ has derivative in w of modulus at most $3A_t/w$. For $\xi \ne \mathbf{1}$ the derivative is the negative of the corresponding twisted $\Psi(A_t/w)$; for the principal character it is $A_t/w - \Psi(A_t/w)$. Integration between two positive weights w, w' gives

$$|wG_\xi(A_t/w) - w'G_\xi(A_t/w')| \le \frac{3A_t}{\min(w, w')} |w - w'|.$$

These bounds give a choice of the absolute constant in the movement estimates. Chebyshev and weight comparison · Movement bounds in the release

Proved here. For every thick partner replace w_g by b, at cost at most $3(x+1)(\eta_w+1)$ per character and position. For a thin partner $A_t/w_g \ge (x-1)b/(2b') \ge X'/3$, so discarding its term costs at most $(x+1)b\theta$. There are f positions and $2f$ characters with a prefactor $1/f$. Thus the modulus of the error in the Fourier sum, for each t, j, z, is at most

$$E_{\mathrm{node}} = 6f(x+1)(\eta_w+1) + 2f(x+1)b\theta \le 2b/x.$$

Keeping the two position counts gives the sharper node bound $2[3n_2(x+1)(\eta_w+1) + (f-n_2)(x+1)b\theta] \le b/x$, as used in Paper II. The coarser allowance E_{node} gives the same released exponent. The thick sum left over is independent of j. Write $e_1(v) = ((1,0), v)$. The thick partners have $g = e_1(s+z-t)$ for $s \in T_2$, so this sum is

$$\mathrm{Th}_{z,t} = \frac{b}{f} \sum_{s \in T_2} \sum_{\xi} \xi(-e_1(s+z-t)) G_\xi(y_t).$$

Single-node comparison and constant

Let $H_0 \le H$ with $T_2 \subset H_0$. Denote the characters trivial on $\{0\} \times H_0$ by Ψ_{H_0}. Summing the character in $\mathrm{Th}_{z,t}$ over $z \in H_0$ gives zero unless $\xi \in \Psi_{H_0}$. For each surviving ξ the factor is $|H_0| \xi(e_1(t))$, since

$s \in H_0$ and the first bit is of order two. Hence

$$\frac{1}{|H_0|}\sum_{z\in H_0}\mathrm{Th}_{z,t} = \frac{b\Delta_2}{2}\sum_{\xi\in\Psi_{H_0}}\xi(e_1(t))G_\xi(y_t).$$

This gives the $1/2$ in the normalizer

$$P = \frac{\Delta_2}{2}\,n_T m b.$$

Whenever $\frac{1}{|H_0|}\sum_{z\in H_0}\mathrm{gap}(z) \le \mathscr{E}$, the modulus estimate for the nodes gives

$$\frac{mb\Delta_2}{2}\sum_{t\in\mathrm{TT}}\mathrm{Re}\sum_{\xi\in\Psi_{H_0}}\xi(e_1(t))G_\xi(y_t) \;\le\; \mathscr{E} + n_T m E_{\mathrm{node}} \;\le\; \mathscr{E} + 2n_T mb/x.$$

The fixed-level hypothesis of pointwise control implies the mean hypothesis in this display. Subgroup orthogonality · Subgroup comparison

In the ordinary case $H_0 = H$, the four survivors are $\mathbf{1}, \psi_C, \chi_C, \chi_C\psi_C$, with signs $+,+,-,-$ at s_0. Suppose the three nonprincipal survivors obey $|G_\xi(x)| \le x^{1/2+\varepsilon}$ and $|\ell(\xi)| \le C_L \log x$ with $C_L \ge 0$, as in the weaker form of Paper II's bounds. The movement from y_t to x contributes at most

$$R_{\mathrm{move}} = C_d + 3(\log(x+1) + C_d + C_L\log x).$$

Before using the upper gap inequality, the Fourier calculation gives the two-sided comparison used for the ordinary normalization in Paper II:

$$\left|G_z(x) - \frac{|H|^{-1}\sum_{z\in H}\mathrm{gap}(z)}{P}\right| \le \sum_{\xi\in\{\psi_C,\chi_C,\chi_C\psi_C\}}|G_\xi(x)| + R_{\mathrm{move}} + \frac{n_T m E_{\mathrm{node}}}{P}.$$

With $\varepsilon > 0$ fixed, $R_{\mathrm{move}} \le x^{1/2+\varepsilon}$ eventually. In particular $P(G_z(x) - 4x^{1/2+\varepsilon}) \le \mathscr{E} + 2n_T mb/x$. Four-character signs and explicit bound · Ordinary translate identification

5. Division by the normalizer

Supplied by the geometric papers. At the chosen cutoff let $a \le 2xb$ and assume the additional size relations

$$n_T \ge \tfrac{1}{2}nx^{-\varepsilon/8}, \qquad x^{-\varepsilon} \le \Delta_2 \le 1, \qquad x^{1/2-\varepsilon} \le m \le x^{1/2+\varepsilon/2},$$
$$xm \le k \le n^2 x^{1/2-\varepsilon}.$$

Then $Q_b = n_T mb > 0$ and $P = Q_b\Delta_2/2 > 0$. **Proved here:** for any real V_x, if

$$P(V_x - 4x^{1/2+\varepsilon}) \le x^{C_*\varepsilon}\mathcal{B} + 2Q_b/x,$$

division gives the exact normalized bound

$$V_x \le 4x^{1/2+\varepsilon} + \frac{2x^{C_*\varepsilon}}{\Delta_2}\frac{\mathcal{B}}{Q_b} + \frac{4}{x\Delta_2}.$$

In the ordinary case put $V_x = G_z(x)$. With the preceding size relations, the four terms of $\mathcal{B}/Q_b$ satisfy

$$\frac{a}{mb} \le \frac{2x}{m} \le 2x^{1/2+\varepsilon}, \qquad \frac{am}{bk} \le \frac{2xm}{k} \le 2,$$
$$\frac{2nmx^{\varepsilon}}{n_T} \le 4x^{1/2+13\varepsilon/8}, \qquad \frac{n_2 k x^{\varepsilon}}{n^2 x^{-\varepsilon} n_T} \le 2x^{1/2+9\varepsilon/8}.$$

The last inequality uses $n_2 \le n$. Thus $\mathcal{B}/Q_b \le 10x^{1/2+13\varepsilon/8}$. Since $2/x \le 1$ and $1/\Delta_2 \le x^{\varepsilon}$, the normalized bound implies

$$V_x \le 26x^{1/2+(C_*+21/8)\varepsilon} \le x^{1/2+(C_*+3)\varepsilon} \quad \text{if } 26 \le x^{3\varepsilon/8}.$$

The two size costs x/m and m enter before density and logarithmic allowances. Their simultaneous balance at m of order $x^{1/2+O(\varepsilon)}$ accounts for the half exponent. Size hypotheses · Size bound

For the integrated refinement substitute $\mathscr{E}_{\mathrm{int}}$ for $x^{C_*\varepsilon}\mathcal{B}$ in the same division. The four bounds on $\mathcal{B}/Q_b$ apply to the two parts of Q_2/Q_b and to $Q_0/Q_b, Q_1/Q_b$ through the budgets in §3. With $\Delta_2^{-1} \le y$, $e_{\mathrm{Div}} \le e_D$, and $e_{\mathrm{Loss}} \le e_L$, this gives

$$\frac{\mathscr{E}_{\mathrm{int}}}{C_{\mathrm{gap}}\Delta_2 Q_b} \le 2x^{1/2+3\varepsilon} + 2x^{1/2+(2+e_L)\varepsilon} + 2x^{2\varepsilon} + 12x^{1/2+(21/8+e_D)\varepsilon}.$$

For $\nu = \max(3, 2+e_L, 21/8+e_D)$ and $x \ge 3$, we obtain $V_x \le (36C_{\mathrm{gap}}+6)x^{1/2+\nu\varepsilon}$, including the other two terms in the division. In the ordinary case this holds for $V_x = G_z(x)$ once $\log x \le x^{\varepsilon}$ and $R_{\mathrm{move}} \le x^{1/2+\varepsilon}$, where R_{move} is defined in §4.

6. The fixed-level window input

The release also needs a bound for one fixed family of primitive characters, even when the Artin group varies with x. Use the subgroup H_0 from §4, of index $o = [H : H_0]$ fixed by the level $\mathcal{L}$. Let $K_{\mathrm{keep}} \subset \Psi_{H_0}$ be the kept characters. **Supplied by the host comparison:** K_{keep} contains the principal character, its characters have $\xi(s_0) = 1$, and $|\Psi_{H_0}| = 4o$. Put $c = (3o+1)^{-1}$ and $a_\xi = 1$ if ξ is trivial on H, $a_\xi = c$ otherwise. For each kept character nontrivial on H, assume

$$\left| \sum_{t \in \mathrm{TT}} \xi((0,0),t) - cn_T \right| \le 4o\, x^{-1+\varepsilon} n_T.$$

The host comparison gives $|\ell(\xi)| \le K_o$ for nonprincipal kept characters, and $|\ell(\xi)| \le C_E \log x$, $|G_\xi(x)| \le x^{1/2+\varepsilon}$ for discarded characters in Ψ_{H_0}. The primitive readout identifies the kept sum with

$$G_{\mathcal{L}}(x) = \sum_{\xi \in K_{\mathrm{keep}}} a_\xi \operatorname{Re} G_\xi(x).$$

As the host changes this is the same function, with the primitive characters and weights fixed by $\mathcal{L}$. In particular $G_{\mathcal{L}} = G_z$ at the conductor-one level of Paper II. Host character and share hypotheses · Reading the fixed-level sum · Bound for discarded characters

Proved here. Let $\Sigma(x) = \sum_{\xi \in \Psi_{H_0} \setminus K_{\mathrm{keep}}} |G_\xi(x)|$. For any nonprincipal character the movement bound of §4 and $G_\xi(1) = \ell(\xi)$ give $|G_\xi(x)| \le x(\log x + C_d + |\ell(\xi)|)$. For a kept character ξ, first move $G_\xi(y_t)$ to $G_\xi(x)$ at each t. The remaining sum of phases is exactly n_T if ξ is trivial on H. In the other case the

deviation from cn_T is bounded by the share hypothesis. For a discarded character the real part is bounded below by $-|G_\xi(x)|$ less the movement error. For the principal character the movement costs at most C_d. Thus the sum on the left of the subgroup comparison in §4, before its factor $mb\Delta_2/2$, is at least

$$n_T\big(G_{\mathcal{L}}(x) - \Sigma(x) - 4oR(x)\big),$$

where one can take

$$R(x) = \max\{C_d,\ \log(x+1) + C_d + K_o + 4o\,x^{\varepsilon}(\log x + C_d + K_o),\ \log(x+1) + C_d + C_E \log x\}.$$

For fixed o, K_o, C_E, ε, $4oR(x) \le 4x^{1/2+\varepsilon}$ eventually. Apply §5 with $V_x = G_{\mathcal{L}}(x) - \Sigma(x)$. For $e_{\mathrm{Div}} \le e_D,\ e_{\mathrm{Loss}} \le e_L$, put $C_* = e_D + e_L + 3$ with $e_D, e_L \ge 0$. Then

$$G_{\mathcal{L}}(x) \le x^{1/2+(e_D+e_L+6)\varepsilon} + 4o\,x^{1/2+\varepsilon} \le x^{1/2+(e_D+e_L+7)\varepsilon},$$

on increasing the threshold so $x^{\varepsilon} \ge 1 + 4o$. In the ordinary case the direct estimate in §5 gives the stronger exponent $1/2 + (e_D + e_L + 6)\varepsilon$ under Paper II's three-character bounds. One-character comparison · Size and discarded-character absorption

For the integrated refinement use the pointwise bound (G) at each subgroup translate in §4. With $x \ge 3,\ \log x \le x^{\varepsilon},\ 4oR(x) \le 4x^{1/2+\varepsilon}$, the integrated estimate of §5 applies to $G_{\mathcal{L}}(x) - \Sigma(x)$. Adding $\Sigma(x) \le 4ox^{1/2+\varepsilon}$ gives $G_{\mathcal{L}}(x) \le (36C_{\mathrm{gap}} + 6 + 4o)x^{1/2+\nu\varepsilon}$. The cutoff conditions hold beyond one threshold for fixed ε and the uniform host constants, since $R(x) = O(x^{\varepsilon}(1 + \log x))$.

Handoff to Paper II and the analytic window argument. In the notation of Paper II, $E(u) = \mathrm{gap}(u)$, $\overline{E} = |H|^{-1}\sum_u E(u)$, and its Q is our P. The a_E in its error bracket is the reference size a. The arithmetic input for its ordinary prime comparison is $\overline{E} \le x^{C_*\varepsilon}\mathcal{B}$ for the host used in that comparison. The pointwise form gives inequality (9) of Paper II. For the release's fixed-level route use $\mathrm{gap}(z) \le x^{C_*\varepsilon}\mathcal{B}$ for every $z \in H_0$. Fix a level $\mathcal{L}$, a window exponent $\lambda > 1$, and e_D, e_L. For every sufficiently small $\varepsilon > 0$ and every requested lower cutoff there must be an interval $[Y, Y^{\lambda}]$ starting above that cutoff. At every real x in that interval one has compatible host and geometric data with bounds uniform in x as above. A different host and lattice family may be selected at each x, with the same $\mathcal{L}$. The release uses $\lambda = 3/2,\ e_D = 2,\ e_L = 0,\ 0 < \varepsilon < 1/52$ along its geometric route, giving the one-sided bound $G_{\mathcal{L}}(x) \le x^{1/2+9\varepsilon}$ on late intervals. For any $1/2 < \beta < 1$ choose $\varepsilon > 0$ so $9\varepsilon \le \beta - 1/2$. Taking the lower cutoff beyond the thresholds in §§3–6 then gives $G_{\mathcal{L}}(x) \le x^{\beta}$ throughout arbitrarily late intervals. That is the hypothesis of the analytic window theorem for the positively weighted family of primitive characters. Its proof and the existence of the windows are supplied by the companion papers. Paper II · Window hypothesis · Host supply · Released exponents · One-point selection · Analytic consumer

7. The geometric inputs to be established in prose

The same window quantifiers in §6 apply to the integrated estimate with either the old frame budget or the allowance E_{int}. The exponent bounds e_D, e_L and the window exponent λ are fixed first. For each sufficiently small fixed $\varepsilon > 0$, the size and host constants are uniform before requesting the lower cutoff Y_0. For either ramp let C_{ramp} be its constant from §5 or §6. For any $\rho > 0$, take Y_0 beyond the relevant movement and size thresholds and $\max(1, C_{\mathrm{ramp}})^{1/(\rho\varepsilon)}$. This gives a unit-coefficient upper bound with exponent $1/2 + (\nu + \rho)\varepsilon$ at every real cutoff in the supplied interval. Given $1/2 < \beta < 1$, first choose ε small enough that $(\nu + \rho)\varepsilon \le \beta - 1/2$. For $e_D = 2, e_L = 0$ the general size bound has $\nu = 37/8$. The

separate fixed-family height calculation obtains $17/16 + \rho$ using the original host ranks, thick density and sharper signed height budgets. Fixed-family height calculation

The implications in §§1–6 are conditional on the following joint existence assertion. For each fixed primitive character and each sufficiently small $\varepsilon > 0$, supply the windows just stated. At every sufficiently large x in those windows, choose compatible data to meet all four requirements below, with the constants of §3 fixed on the family for that ε. The analytic estimates at the host must apply to the same Artin data, weights and tops. Tying the fields and data · Geometric input assembly · Fixed-level interface

1. **Columns, periods and lengths.** Construct the row and column lattices, their metrics and the rational period maps. Prove that the common row is the metric dual of R, that the maps on the shifted columns are integral, that the full and retained ranks are positive with the stated restriction, and that every local length, shift and gain satisfies §1. In particular the finite index loss comes from the same sublattice N_i^0 used to define N_i. The common divisor term used there must be the one in the height inequality. Place data · Period-integral lattice and shifts

2. **Frame, counts and height.** Establish the single-place matrix realization with positive normalizing weights, the trace and Gram estimates with the indicated row and column dimensions, and the height bound for each nonzero map. Prove the quantitative row minimum-slope window, the count and full-column slope estimates, the type-prime compensation and the corrected-height variation in §2. In the divisor estimate use the signed paired local determinant expression of the row covolume and bound its finite and infinite terms to obtain E_{Div} with $e_{\mathrm{Div}} \le 2$. This step requires the common-divisor estimate for the selected family, in addition to the formal cancellation of F_o. Frame hypotheses · Finite determinant expression · Global height input

3. **Finite-prime loss.** For each translate used in the pointwise argument, establish both $\mathrm{ASK}_\alpha - G \le c_{\mathrm{Loss}} n_T mbx^{e_{\mathrm{Loss}}\varepsilon}$ and $I_O \le c_{\mathrm{Loss}} n_T mbx^{e_{\mathrm{Loss}}\varepsilon}$ with $e_{\mathrm{Loss}} \le 0$. This entails the estimates on the orbits served at eligible primes, on the net gain $\sum_j g_{jw} - \tau_w^*$, and on the omitted contributions after weighting by $\log p$. It includes the distinguished ramified prime and the length charge I_O. An averaged pair of bounds would suffice for the ordinary mean route in §2. Local gain and served orbits · Pointwise loss interface · Released loss exponent

4. **Uniform supply and compatibility.** For all cutoffs in the windows, prove the coverage of the required translates, the shared top and weight identifications, the size inequalities of §§3 and 5, and the thick and thin conditions of §4. In the fixed-level case include the share and primitive readout hypotheses of §6. In particular a comparison for one abstract choice of numerical slopes gives the scalar implication, and the needed existence statement requires slopes of the actual lattices built from the chosen host. Pointwise data and coverage · Compatibility equations · Quantified family supply

The geometric statements are inputs here. The height and scalar estimates above show how those statements would yield the ordinary and fixed-level prime comparison used in the series.