Dense minimum-cost flow in $\ddot{O}(n^2 \log n)$ word operations

November 2, 2026

Abstract

We give a randomized word-RAM algorithm that computes an exact minimum-cost flow on a directed graph with n vertices, at most one arc per ordered pair, and polynomially bounded integer capacities, costs, and demands in $\ddot{O}(n^2 \log n)$ word operations, where $\ddot{O}(\cdot)$ hides factors of $(\log \log n)^{O(1)}$. The running time bound holds on every random tape. The output is an exact integral flow with an integer dual certificate, or a verified infeasibility certificate, with probability at least $1 - n^{-c_0}$ for any fixed constant c_0, and FAIL otherwise. Every returned answer is checked exactly before it is released. Dense bipartite minimum-cost perfect matching is a special case. The dense work per attempt is $\ddot{O}(n^2)$ because maintained original-edge cores and packed arithmetic answer the per-frame queries in time proportional to the number of selected labels, and $O(\log n)$ attempts supply both progress and confidence.

This manuscript was generated by Alicia Stepin using GPT6 Astra from a manuscript on short-step interior point methods, and then edited using the Devin harness.

1 Main result

Theorem 1.1. *Let $C, c_0 > 0$ be constants, and let a directed graph on $n \geq 2$ vertices be given with*

- *at most one arc per ordered pair of vertices, and*
- *integer lower capacities, upper capacities, costs, and demands of absolute value at most n^C.*

There is a randomized word-RAM algorithm with $O(\log n)$-bit words and polynomial address space (Section 1.2) that uses at most

$$\ddot{O}\left(n^2 \log n\right)$$

word operations on every random tape, and returns, with probability at least $1 - n^{-c_0}$, either

- *an exact minimum-cost feasible integral flow together with an integer dual certificate, or*
- *a verified infeasibility certificate, which is a bound inconsistency, a nonzero total demand, or a violated capacity cut.*

On the remaining outcomes the algorithm returns FAIL.

Every returned mathematical answer is checked exactly before it is released, so the algorithm never returns an incorrect flow, dual, or cut.

Corollary 1.2. *Let $C, c_0 > 0$ be constants and let $n \geq 2$ be even. Given integer costs $|\mathbf{c}_e| \leq n^C$ on the edges of $K_{n/2,n/2}$, a randomized word-RAM algorithm returns, with probability at least $1 - n^{-c_0}$, an exact minimum-cost perfect matching together with an integer assignment dual certificate, in*

$$\ddot{O}\left(n^2 \log n\right)$$

word operations in the model of Section 1.2, with the same failure convention as Theorem 1.1.

1.1 Logarithmic factors and constants

Throughout, $\ddot{O}(\cdot)$ hides factors of $(\log\log n)^{O(1)}$ and nothing else; every factor of $\log n$ is written out, and $\log m$ is identified with $\log n$ since $m \le n^2$. All constants may depend on the input exponent C and the confidence exponent c_0. Words have $O(\log n)$ bits; the word length is written $\log n$ and its bit length $\log\log n$, up to the integer rounding fixed once here: an integer parameter defined by a real expression is that expression rounded up to the next integer, and a parameter stated as a power of two is rounded up to the next power of two. Every such rounding changes constants only and is not written. The tracking potential exponent is $p = 512$, and the central potential exponent P is the least integer with $P \ge 512$ and

$$P \ge 1 + \log_2\bigl(3 + \log_2(2+n)\bigr),$$

so that $P = \Theta(\log\log n)$ and $P \ge p$. Every bound of the form $O(mP^a(\log\log n)^b)$ with constants a, b is written $\ddot{O}(m)$, and every $O(\log n \log\log n)$ is written $\ddot{O}(\log n)$. The ledger of Section 15 tracks the hidden powers explicitly; the dense work per attempt is $O(mP^{14}(\log\log n)^{19})$, which is $O(m(\log\log n)^{40})$ since $P = O(\log\log n)$. Factors of $\log\log n$ are not optimized or traded against each other. Sparse work is below n^2 by a fixed power of n; its logarithmic factors are not counted and are absorbed in the $o(1)$ of its exponent, as in $n^{15/8+o(1)}$.

1.2 Machine model and notation

The word RAM has indexed memory, comparisons, Boolean operations, shifts, and integer addition, subtraction, multiplication, quotient, and remainder on $O(\log n)$-bit words, together with independent uniform random words. A fixed number of wider products may occupy a fixed number of words. Population count is implemented with bit operations. Square roots, logarithms, powers, and exponentials use certified approximations. Exact geometric sampling may use many words; its actual work is charged before the attempt cap. Polynomial address space may be reserved without clearing it; the initializable-array construction specifies which cells are read.

The augmentation in Section 10 gives a connected working graph with $m = \Theta(n^2)$ labels and rank $r = n - 1 = \Theta(n)$. From that point, n, m, r refer to the working graph. Its dimensions differ from the input dimensions by constant factors. As an undirected graph it has bounded parallel multiplicity, so the structural routines use a fixed number of simple layers.

Appendix B lists the symbols in order of first appearance. The local-step analysis writes the incidence matrix as $\mathbf{B}$ throughout. Products, quotients, powers, and inequalities between edge vectors are coordinatewise. For positive $\mathbf{w}$, $\|\mathbf{x}\|_{\mathbf{w}}^2 = \sum_e \mathbf{w}_e \mathbf{x}_e^2$; matrix inequalities use the positive-semidefinite order. Laplacian inverses act on the quotient by their componentwise constant vectors. The local update defines $\mathbf{H}, \boldsymbol{\tau}, \boldsymbol{\varpi}, \boldsymbol{\delta}, \boldsymbol{\varsigma}$ precisely in Section 2. Logarithmic scales are rounded by integer bit lengths; local integer or dyadic conventions change only constants.

1.3 The algorithm and the two maintenance tasks

The iteration controls stationarity for a weighted logarithmic capacity barrier and tracks regularized leverage scores with adaptive weights $\boldsymbol{\lambda}$. Separate potentials keep the two errors away from their boundaries. A residual test controls approximate flow feasibility during the iteration; one terminal correction restores exact feasibility.

A *frame* is a sparse physical update with reference and metadata service; an *attempt* runs a prescribed number of frames from a checkpoint. An *original label* identifies a working edge, not a sampled copy. A graph *owner* stores labels in cores with certified cut-conductance bounds and in explicit lists.

1. *Initialize.* Augmentation gives an exactly feasible interior midpoint. Initialize tracking and centrality, then save the physical variables and certified references at a checkpoint.

2. *Run an attempt.* Each frame fixes its directions from the old state and draws one independent original-label mask. Only selected flow and tracking weights change. Local drift controls both potentials. Operational failure restores the checkpoint; the next attempt uses fresh tapes, with no local resampling of a failed frame.

3. *Maintain graphs and observations within each frame.* Three graphs, changing only on selected labels, support the ordinary inverse, physical solve, and observer. Cores answer scalar queries; packed computations evaluate fine scores and change sources. Observations change metadata only.

4. *Finish.* Success installs a checkpoint with a constant-factor smaller μ. After enough successes, exact feasibility correction and integral rounding yield the original flow and dual, or the augmentation yields an infeasibility cut.

The owner finds labels with the prescribed sampling probabilities; the observer discovers references needing fine reads, without scanning all edges each frame. Graph updates retain prices frozen before the batch. Autonomous reads spend accumulated score-variation credit; selected reads use frozen list prices. Representation changes preserve outstanding debts and due flags.

1.4 Where the single logarithm comes from

The proof keeps the following costs separate:

- dense expected work of one attempt, $\ddot{O}(m)$;
- additional sparse work of one attempt, $n^{15/8+o(1)} + n^{1011/512+o(1)}$;
- cold initialization, once, $\ddot{O}(m \log n)$;
- exact output, once, $O((m+n)\log n)$.

Each attempt is capped at a fixed multiple of its expected work. The phase analysis gives a positive constant conditional success probability, and $O(\log n)$ total attempts supply both progress and confidence. Since $m = \Theta(n^2)$, the $O(\log n)$ attempts cost $\ddot{O}(n^2 \log n)$ in total, and the sparse exponents' strict gaps below two absorb their logarithmic factors.

Within a stopped attempt, the per-frame scales B_* sum to $O(m)$ (Lemma 11.1). Two constructions convert this budget into the dense work bound. The core owner bounds *moments* of current update prices, including low-capacity records; individual low-density cores may still be expensive. Packed score reads and source assembly process the confidence lanes together, avoiding an additional factor $\log n$ at every dense event.

For the complete argument, read the execution order (Section 13), phase analysis (Section 11), and contract closure (Section 14), followed by the ledger and main proof (Sections 15 and 16).

2 One independent physical step

An independent mask selects the labels whose flow and weight change. Given certified directions and observations, we prove that the step reduces μ, controls feasibility, and contracts the central and leverage-tracking potentials. Section 3 implements this step with finite words; later sections construct its inputs.

2.1 State, directions, and one independent mask

Let $\mathbf{B}$ be an $m \times n$ incidence matrix of rank $r = n - 1$, with $m = \Theta(n^2)$, and put $s = r/m$. The flow is strictly inside its positive capacities, $0 < \mathbf{f}_e < \boldsymbol{\upsilon}_e^{\max}$. Write $\boldsymbol{\lambda}$ for the edge weights and $\mathbf{y}$ for the vertex potential. Define coordinatewise

$$\mathbf{z} = (\mathbf{f}^{-2} + (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-2})^{-1/2}, \quad \overline{\boldsymbol{\Xi}} = \mathbf{z}(\mathbf{f}^{-1} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-1}), \quad \mathbf{D} = \mathbf{z}^3(\mathbf{f}^{-3} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-3}),$$

$$\mathbf{w} = 2s + \boldsymbol{\lambda}, \qquad \boldsymbol{\Xi} = \frac{\mathbf{z}(\mathbf{c} - \mathbf{B}\mathbf{y})}{\mu \mathbf{w}} = \overline{\boldsymbol{\Xi}} + \varepsilon \boldsymbol{\varsigma}, \qquad \boldsymbol{\varrho} = \mathbf{z}/\mathbf{w}, \qquad \mathbf{H} = \mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2)\mathbf{B}.$$

For the logarithmic box barrier, $\mathbf{z}$ is the inverse square root of the second derivative, $\overline{\boldsymbol{\Xi}}$ is minus the first derivative times $\mathbf{z}$, and $\mathbf{D} = dz/df$. Thus e measures centrality in the local flow scale. Here $\boldsymbol{\lambda} \geq s$, $\mu > 0$, and $|\boldsymbol{\varsigma}| < 1$. Work on the rank-$r$ quotient, so that $\mathbf{H}^{-1/2}$ denotes its ordinary inverse there. Set

$$\mathbf{U} = \textsc{Diag}(\boldsymbol{\varrho})\mathbf{B}\mathbf{H}^{-1/2}, \quad \boldsymbol{\Pi} = \mathbf{U}\mathbf{U}^T, \quad \boldsymbol{\tau} = \textsc{Diag}\boldsymbol{\Pi}, \quad \widehat{\boldsymbol{\tau}} = s + \boldsymbol{\tau}, \quad t = \log(\boldsymbol{\lambda}/\widehat{\boldsymbol{\tau}}), \quad \delta = 1 - |t|/\kappa.$$

Assume $|\boldsymbol{\varpi}| < \kappa$ for fixed $\varepsilon, \kappa \leq 10^{-2}$. Then $\sum \boldsymbol{\lambda} \leq 2e^{\kappa} r$ and $\sum \mathbf{w} \leq (2 + 2e^{\kappa})r$. The tracking error $\boldsymbol{\varpi}$ and central error e have potentials

$$\psi_{\varkappa}(x) = (1 - x)^{-\varkappa} + (1 + x)^{-\varkappa}, \quad \Psi_R = \sum_e \psi_p(\boldsymbol{\varpi}_e/\kappa), \quad \Psi_Z = \sum_e \psi_P(\boldsymbol{\varsigma}_e), \quad p = 512,$$

$$10^5 n^{1/4} \leq M \leq 2 \cdot 10^5 n^{1/4},$$
$$\tfrac{1}{2} c_h n^{-3/p} P^{-16} \leq h \leq c_h n^{-3/p} P^{-16}.$$
$$\beta_f = h/M, \quad \beta_s = \gamma \beta_f, \quad B_* = \beta_f m/\sqrt{r}, \qquad B_* \leq \widehat{B}_* \leq 2B_*.$$

Choose $M, h, \gamma, \widehat{B}_*$ dyadic; B_* is analytical only. The constants c_h, γ are fixed below. The direction scale is β_f, the source radius is β_s, and B_* measures dense work per frame.

Tracking feedback. Let $\pi = \boldsymbol{\lambda}(1 + \varepsilon\boldsymbol{\varsigma})$ and $\mathcal{N}(v) = \|v\|_\infty + M\|v\|_\pi$. This mixed norm controls both coordinate jumps and weighted movement. The tracking feedback $\boldsymbol{\varphi}$ obeys

$$\mathcal{N}(\boldsymbol{\varphi}) \leq h, \qquad \mathbf{G}^T\boldsymbol{\varphi} \leq -0.97h\|\mathbf{G}\|_{\mathcal{N}^*} + ChB_R, \qquad \mathbf{G} = \nabla_t \Psi_R, \quad B_R = C_{p,\kappa}\|\mathbf{1}\|_{\mathcal{N}^*}. \tag{1}$$

A lower-wall guard sets negative feedback to zero when $\boldsymbol{\lambda} < (1 + 2a_0)s$, where a_0 upper-bounds every realized coordinate jump. Such guarding loses only ChB_R: a positive gradient removed by it can occur only when $0 < t < \log(1 + 2a_0)$, where the derivative is bounded by the fixed constant included in $C_{p,\kappa}$. Lemma 8.3 constructs this feedback from the maintained joint classes.

Central source. The source points inward from a representative $\overline{\boldsymbol{\varsigma}}_e$ of the central error. Assume

$$|\boldsymbol{\varsigma}_e - \overline{\boldsymbol{\varsigma}}_e| \leq \eta(1 - |\overline{\boldsymbol{\varsigma}}_e|)/P^4, \qquad d_{Z,e} = 1 - |\overline{\boldsymbol{\varsigma}}_e|, \qquad G_{Z,e} = |\psi_P'(\overline{\boldsymbol{\varsigma}}_e)|, \qquad J_e = \psi_P''(\overline{\boldsymbol{\varsigma}}_e), \tag{2}$$

where η is a sufficiently small fixed constant. Let $\widetilde{\mathbf{w}}_e/\mathbf{w}_e \in [1 - \eta, 1 + \eta]$ be constant within the classes used for the source. Use certified rational class values

$$G_{Z,e} \leq \mathbf{G}^c_{Z,e} \leq (1 + \eta)G_{Z,e}, \qquad J_e \leq J^c_e \leq (1 + \eta)J_e.$$

The zero-gradient class uses $\mathbf{G}^c_{Z,e} = 0$ exactly. Define the capped waterfill

$$\mathbf{F}_e = -\textsc{Sgn}(\overline{\boldsymbol{\varsigma}}_e) \min\{c_Z d_{Z,e}/P^2,\ \gamma h, \vartheta \mathbf{G}^c_{Z,e}/\widetilde{\mathbf{w}}_e\}. \tag{3}$$

Choose rational ϑ so that $(1-\eta)\beta_s \le \|\mathbf{F}\|_{\mathbf{w}} \le \beta_s$ when possible. If the fully capped source has norm at most β_s, it may instead be returned with $\vartheta^{-1} = 0$, whether or not its norm lies in the filled range. Use this fully capped branch when the filled range cannot be reached.

The nonzero source-cap set consists of coordinates with $\mathbf{F}_e \neq 0$ at which either fixed cap $c_Z d_{Z,e}/P^2$ or γh attains the minimum, including ties. Let $\mathcal{J}$ be its complement. A zero-gradient coordinate has $\mathbf{F}_e = 0$ and is never forced by the source; it stays in $\mathcal{J}$, where its positive Hessian contributes to the curvature normalizer. Define

$$A_Z = \sum_e G_{Z,e}|\mathbf{F}_e|, \qquad V = \widehat{B}_* + \beta_s^2/\vartheta.$$

Freeze both inputs, $\boldsymbol{\varphi}, \mathbf{F}$, and the source-cap set before the mask.

Solving for the flow direction. Choose the actual next ratio $\theta = \mu'/\mu$ with $\beta_s/(2\sqrt{r}) \le 1-\theta \le \beta_s/\sqrt{r}$ and put

$$\mathbf{g}_0 = (\theta-1)\boldsymbol{\Xi} + \theta\varepsilon\mathbf{F}, \quad \mathbf{g} = \mathbf{g}_0 + \boldsymbol{\Xi}\boldsymbol{\varphi}/3, \quad \mathbf{Y} = 3\boldsymbol{\lambda}/\mathbf{w}, \quad \mathbf{k} = DC + \mathbf{Y}(1-DP_0),$$
$$\mathbf{L} = \mathbf{B}^T \textsc{Diag}\left(\frac{3\boldsymbol{\lambda}\mathbf{z}^2}{\mu\mathbf{w}^2\mathbf{k}}\right)\mathbf{B}.$$

The two metrics have different roles: the solve uses $\mathbf{L}$ to control feasibility, while the tracked leverage scores belong to $\mathbf{H}$. One has $1-\varepsilon \le \mathbf{k} \le 3+\varepsilon$. Use a direction $\mathbf{x}^{\mathrm{dir}}$ satisfying

$$\mathbf{L}\mathbf{x}^{\mathrm{dir}} = \rho + \mathbf{B}^T \textsc{Diag}(\mathbf{Y}\mathbf{z}/\mathbf{k})\mathbf{g} - r_{\mathrm{sol}}, \quad \rho = \mathbf{B}^T\mathbf{f} - b_{\mathrm{dem}}, \quad \|r_{\mathrm{sol}}\|_{\mathbf{L}^{-1}} \le c_{\mathrm{sol}}\beta_s\sqrt{\mu}, \tag{4}$$

and the old residual has norm at most $\beta_s\sqrt{\mu}$. Define

$$\mathbf{q} = \mathbf{z}\mathbf{B}\mathbf{x}^{\mathrm{dir}}/(\mu\mathbf{w}), \quad \boldsymbol{\Upsilon} = (\mathbf{g}-\mathbf{q})/\mathbf{k}, \quad \mathbf{d} = |\boldsymbol{\Upsilon}| + |\boldsymbol{\varphi}|, \quad \boldsymbol{\nu} = (\mathbf{g}_0 - \mathbf{q})/(\varepsilon\theta).$$

Here $\mathbf{d}$ is the movement scale before the mask is applied, and $\boldsymbol{\nu}$ is the coefficient of the central noise. The bounds below give $U_{\boldsymbol{\nu}} = O(\beta_f)$ with $\|\boldsymbol{\nu}\|_{\mathbf{w}} \le U_{\boldsymbol{\nu}}$.

The sampling rates. Let $a_{\mathrm{jump}} > 0$ be a sufficiently small fixed scalar jump allowance, distinct from the row-weight vector $\boldsymbol{\varrho} = \mathbf{z}/\mathbf{w}$. Let $\overline{\ell}_e$ be a supplied upper bound for the actual leverage of row e in $\mathbf{L}$. Let $\overline{\boldsymbol{\delta}}_e$ be the stored certified tracking-margin scale, with a fixed known constant $C_\delta \ge 1$ such that $C_\delta^{-1}\boldsymbol{\delta}_e \le \overline{\boldsymbol{\delta}}_e \le C_\delta\boldsymbol{\delta}_e$. Freeze all inputs, and write $\mathcal{F}$ for this history before the fresh mask coins. Form a fixed number of rational component probabilities, upward within fixed factors of the following clipped rates:

$$K\log n\boldsymbol{\lambda}_e, \quad K\log n\overline{\ell}_e, \quad \frac{KLr\mathbf{w}_e d_e^2}{\beta_s^2}, \quad \frac{C_\delta\mathbf{d}_e}{a_{\mathrm{jump}}\overline{\boldsymbol{\delta}}_e}. \tag{5}$$

The implemented margin target uses only $\overline{\boldsymbol{\delta}}_e$ and dominates the analytical floor $\mathbf{d}_e/(a_{\mathrm{jump}}\boldsymbol{\delta}_e)$. Every local estimate below continues to use that analytical lower bound; evaluating a proposal probability requires no fresh fine-score read. On $\mathcal{J}$, also include

$$\frac{k_P|\boldsymbol{\nu}_e|}{d_{Z,e}}, \qquad \frac{k_P B_{\boldsymbol{\nu}}|\boldsymbol{\nu}_e|\sqrt{J_e^c}}{V}, \qquad B_{\boldsymbol{\nu}} \ge U_{\boldsymbol{\nu}}\sqrt{\sum_{e\in\mathcal{J}} J_e^c/\mathbf{w}_e}, \qquad k_P = CP^{12}. \tag{6}$$

The leverage and quadratic rates give the matrix and residual concentration bounds below. The tracking and central wall rates limit individual jumps; the remaining central rate controls second-order potential cost. Lemma 2.7 chooses its $B_{\boldsymbol{\nu}}$ from common scalars without a reciprocal scan. Force the nonzero source-cap set. These are all the physical-mask components; (13) supplies the gradient bound without another sampling rate.

The selected update and its commit. Draw the components independently between original labels. Their union indicators I_e are independent, with actual union probabilities p_e. Use $\mathbf{T}_e = I_e/p_e$ with the zero-target no-op convention; the denominator is the union probability, not the sum of component probabilities. Set $v = Tu$, $\zeta = \mathbf{T}\boldsymbol{\varphi}$, and on selected labels define

$$\boldsymbol{\lambda}' = s + [\mathbf{w}(1 + \zeta/3) - 3s](1 + Dv)^2, \qquad \mathbf{w}' = 2s + \boldsymbol{\lambda}'. \tag{7}$$

Set $\mathbf{y}' = \mathbf{y} - \mathbf{x}^{\mathrm{dir}}$ and $\mu' = \theta\mu$. Choose the new flow by the following scalar root equation, which will make the central noise exactly $(\mathbf{T} - 1)\boldsymbol{\nu}$ in (9). The root near $\mathbf{f}$ solves

$$\frac{\mathbf{z}(\mathbf{f}')(\mathbf{c} - \mathbf{B}\mathbf{y}')}{\theta\mu\mathbf{w}'} - \overline{\boldsymbol{\Xi}}(\mathbf{f}') = \frac{C + \mathbf{q}}{\theta} - \overline{\boldsymbol{\Xi}} + \frac{kv - C\zeta/3}{\theta}. \tag{8}$$

Unselected labels retain their stored $\mathbf{f}, \boldsymbol{\lambda}$ and structural coefficients. Each frame has one candidate: a failed residual or source-normalizer check aborts the whole clean attempt, with no local retry. No moment bound conditions on these checks or eventual phase success.

Buffer candidate points and their residual/source-normalizer data. Retain the old $\mathbf{H}$ and ordinary owner through all conserved-source products and old-observer batches; then commit and service queued reads at frozen old prices, in Steps 6–8.

2.2 Scalar identities and direction bounds

Lemma 2.1. *Put $\mathbf{X}_e = T_e d_e$ and $\sigma_e = |\mathbf{q}_e| + |1 - \theta|$. For $\mathbf{X}_e \le 10^{-4}$ and $\sigma_e \le 10^{-4}$, the root in $|\mathbf{f}'_e - \mathbf{f}_e| \le \mathbf{z}_e/40$ exists uniquely and stays in the box. The guard gives $\boldsymbol{\lambda}'_e \ge s$. Before numerical rounding,*

$$\boldsymbol{\varsigma}' = \boldsymbol{\varsigma} + \mathbf{F} + (\mathbf{T} - 1)\boldsymbol{\nu}, \tag{9}$$

$$\rho' = r_{\mathrm{sol}} + \mathbf{B}^T \mathrm{DIAG}(\mathbf{Y}\mathbf{z})(\boldsymbol{\Lambda} - \boldsymbol{\Upsilon}) + \mathbf{B}^T \mathrm{DIAG}(\mathbf{z}) r_f, \tag{10}$$

$$(\mathbf{f}' - \mathbf{f})/\mathbf{z} = \mathbf{Y}\boldsymbol{\Lambda} + r_f, \qquad |r_f| \le C(X^2 + \sigma X), \tag{11}$$

$$(a'/a) - 1 = Dv - \zeta/3 + r_a, \qquad |r_a| \le C(X^2 + \sigma X). \tag{12}$$

All constants are uniform as $\mathbf{f}$ approaches either capacity wall.

Proof. Suppress the edge index. Write $\alpha = \mathbf{z}/\mathbf{f}$ and $\beta = \mathbf{z}/(\boldsymbol{\upsilon}^{\max} - \mathbf{f})$, so $\alpha^2 + \beta^2 = 1$, $\overline{\boldsymbol{\Xi}} = \alpha - \beta$, and $\mathbf{D} = \alpha^3 - \beta^3$. Hence $\mathbf{D} = (1 + \alpha\beta)\overline{\boldsymbol{\Xi}}$, $|\mathbf{D}|, |\overline{\boldsymbol{\Xi}}| \le 1$, and $1 \le 1 + \alpha\beta \le 3/2$. Direct differentiation gives $dz/df = \mathbf{D}$ and $dP_0/df = (DP_0 - 1)/\mathbf{z}$.

For normalized displacement $x = (\mathbf{f}' - \mathbf{f})/\mathbf{z}$, put $\mathbf{R}(x) = \mathbf{z}(\mathbf{f} + zx)/\mathbf{z}$ and $\mathbf{K}_0 = (C + \mathbf{q})/(\theta(1 + b_{\mathbf{w}}))$, where $b_{\mathbf{w}} = \mathbf{w}'/\mathbf{w} - 1$. Let $\Phi(\varkappa) = \mathbf{R}(\varkappa)\mathbf{K}_0 - \overline{\boldsymbol{\Xi}}(\mathbf{f} + \mathbf{z}\varkappa)$. At the displaced point, with its own D_x, α_x, β_x, differentiation gives

$$\mathbf{R}' = D_x, \qquad \Phi' = \frac{1 + D_x\Phi}{\mathbf{R}}, \qquad D'_x = -\frac{3\alpha_x^2\beta_x^2(1 + 2\alpha_x\beta_x)}{\mathbf{R}}, \qquad \Phi'' = D'_x\Phi/\mathbf{R}.$$

For $|x| \le 1/40$ the hypotheses imply $0.975 \le \mathbf{R} \le 1.025$, $|\Phi| < 0.062$, $\Phi' > 0.915$, and $|\Phi''| < 0.1$. These bounds follow directly from $|D_x| \le 1$, $|D'_x| \le 3/(2\mathbf{R})$, $|\boldsymbol{\Xi} - \overline{\boldsymbol{\Xi}}| \le 0.01$, and the stated jump bounds.

Let $l = (\boldsymbol{\lambda} - s)/\mathbf{w}$. Expansion of (7) gives

$$b_{\mathbf{w}} = \zeta/3 + 2l\mathbf{D}\boldsymbol{\Lambda} + r_{\mathbf{w}}, \qquad |r_{\mathbf{w}}| \le 2\mathbf{X}^2, \qquad \mathbf{Y} = 1 + 2l.$$

If T_0 is the target in (8), direct subtraction gives

$$\theta\{T_0 - \Phi(0) - \mathbf{Y}\mathbf{\Lambda}\Phi'(0)\} = \mathbf{q}\zeta/3 + [-Dq + \mathbf{Y}(1-\theta)(1-DP_0)]v + (C+\mathbf{q})r_{\mathbf{w}} \\ - \frac{(C+\mathbf{q})b_{\mathbf{w}}(\zeta/3 - Dv + r_{\mathbf{w}})}{1+b_{\mathbf{w}}}.$$

Its magnitude is at most $5X^2 + 4\sigma X$. Also $0.989 < \Phi'(0) < 1.011$ and $|T_0 - \Phi(0)| < 3.034X$. Monotonicity and the endpoint displacements give the unique root, initially $|x| < 4X$, and Taylor's formula then gives $|x - \mathbf{Y}\mathbf{\Lambda}| \le 6X^2 + 5\sigma X$. Since $\mathbf{R}(x) = 1 + Dx + O(x^2)$ and $\boldsymbol{\varrho}'/\boldsymbol{\varrho} = \mathbf{R}(x)/(1+b_{\mathbf{w}})$, the row estimate follows. The entire window lies in the box because $\mathbf{z} \le \min\{\mathbf{f}, \boldsymbol{\upsilon}^{\max} - \mathbf{f}\}$.

If the guard is active, $\zeta \ge 0$. Otherwise $\mathbf{w}(1+\zeta/3) \ge s(3+2a_0)(1-a_0/3) \ge 3s$, proving the lower weight bound. Substituting $ku = \mathbf{g} - \mathbf{q}$ and $\mathbf{g} = \mathbf{g}_0 + \mathbf{\Xi}\boldsymbol{\varphi}/3$ proves (9). The solve equation gives $\mathbf{B}^T \textsc{Diag}(\mathbf{Y}\mathbf{z})\boldsymbol{\Upsilon} = -\rho + r_{\mathrm{sol}}$, which proves (10). For $I_e = 0$, the old coordinate is exactly the root. □

The next lemma analyzes the unsampled first-order step, obtained by setting every $\mathbf{T}_e = 1$. Under the independent mask its derivative becomes the mean linear increment, before any restriction to a safety event.

Lemma 2.2. *Let $\mathbf{L}_0 = \mu\mathbf{L}$ and let ℓ_e be its row leverage. Define*

$$\mathbf{Q} = \textsc{Diag}(\boldsymbol{\varrho})\mathbf{B}\mathbf{L}_0^{-1}\mathbf{B}^T\textsc{Diag}(3\boldsymbol{\lambda}\boldsymbol{\varrho}/\mathbf{k}), \quad \mathbf{K}_0 = \textsc{Diag}(\mathbf{k})^{-1}(\mathbf{I} - \mathbf{Q}),$$

$$A_{\boldsymbol{\tau}} = \textsc{Diag}(\widehat{\boldsymbol{\tau}})^{-1}(\mathbf{\Pi} \circ \mathbf{\Pi}), \qquad \boldsymbol{\iota} = 1 - s/\boldsymbol{\lambda} - \boldsymbol{\tau}/\widehat{\boldsymbol{\tau}}.$$

Then

$$\ell_e \le \min\{1, 4\boldsymbol{\lambda}_e^2/s\}, \quad \ell_e/\boldsymbol{\lambda}_e \le 2/\sqrt{s}, \quad \|\mathbf{Q}\mathbf{x}\|_\infty \le 3s^{-1/4}\|\mathbf{x}\|_{\boldsymbol{\lambda}}.$$

The feedback part of the first-order increment of $\boldsymbol{\varpi}$ is $(\mathbf{I} - \mathbf{R})\boldsymbol{\varphi}$, where

$$\mathcal{R} = \frac{2}{3}(A_{\boldsymbol{\tau}} + \textsc{Diag}\boldsymbol{\iota})[\mathbf{I} - \textsc{Diag}(\mathbf{D})\mathbf{K}_0\textsc{Diag}(\mathbf{\Xi})], \qquad \|\mathcal{R}\|_{\mathcal{N}\to\mathcal{N}} < 0.86.$$

The other first-order terms have $\mathcal{N}$ norm at most $C\gamma h$. Moreover

$$\|\mathbf{d}\|_{\mathbf{w}} + \|\boldsymbol{\nu}\|_{\mathbf{w}} \le C\beta_f, \qquad \|\mathbf{d}\|_\infty + \|\mathbf{q}\|_\infty + |1-\theta| \le Ch, \qquad \mathcal{N}(\mathbf{d}) \le C_{\mathrm{amp}}h.$$

Proof. Since $\mathbf{L}_0 \succeq 3sH/(3+\varepsilon)$ and $\boldsymbol{\tau}_e \le e^\kappa \boldsymbol{\lambda}_e$, inverse comparison proves $\ell_e \le 4\boldsymbol{\lambda}_e^2/s$. Every leverage is at most one; taking the smaller bound proves $\ell_e/\boldsymbol{\lambda}_e \le 2/\sqrt{s}$. The matrix $\mathbf{Q}$ is an orthogonal projection in the $\boldsymbol{\lambda}/\mathbf{k}$ inner product. Row Cauchy–Schwarz proves its displayed infinity bound. In a simple augmented dense graph, $s^{-1/4} \le Cn^{1/4}$; the fixed augmentation constant can be absorbed in M if necessary.

The matrix $\mathbf{K}_0$ is self-adjoint positive semidefinite in the $\boldsymbol{\lambda}$ inner product and $0 \preceq \mathbf{K}_0 \preceq \textsc{Diag}(1/\mathbf{k})$ there. With $\alpha = \mathbf{z}/\mathbf{f}$ and $\beta = \mathbf{z}/(\boldsymbol{\upsilon}^{\max} - \mathbf{f})$ as in the scalar calculation, put $\chi = 1 + \alpha\beta \in [1, 3/2]$, so $\mathbf{D} = \chi\overline{\mathbf{\Xi}}$. The matrix $\textsc{Diag}(\mathbf{D})\mathbf{K}_0\textsc{Diag}(\overline{\mathbf{\Xi}})$ is positive semidefinite in the $\boldsymbol{\lambda}/\chi$ inner product and bounded by $(1-\varepsilon)^{-1}\mathbf{I}$. Consequently

$$\|\mathbf{I} - \textsc{Diag}(\mathbf{D})\mathbf{K}_0\textsc{Diag}(\mathbf{\Xi})\|_{\boldsymbol{\lambda}\to\boldsymbol{\lambda}} \le b_0 := \sqrt{3/2} + \varepsilon/(1-\varepsilon).$$

The perturbation from $\mathbf{\Xi} - \overline{\mathbf{\Xi}} = \varepsilon\boldsymbol{\varsigma}$ contributes the second term. The projection identity $\mathbf{\Pi}^2 = \mathbf{\Pi}$ and weighted Cauchy–Schwarz give

$$\|A_{\boldsymbol{\tau}}\|_{\boldsymbol{\lambda}\to\boldsymbol{\lambda}} \le e^\kappa, \quad \|A_{\boldsymbol{\tau}}\mathbf{x}\|_\infty \le e^{\kappa/2}\|\mathbf{x}\|_{\boldsymbol{\lambda}}, \quad \|\boldsymbol{\iota}\|_\infty \le J_0 := e^\kappa - 1.$$

Combining these weighted bounds in the mixed norm gives at most

$$\frac{2}{3}b_0\left[e^{\kappa}\sqrt{\frac{1+\varepsilon}{1-\varepsilon}}+\frac{e^{\kappa/2}}{10^5\sqrt{1-\varepsilon}}\right]+\frac{2}{3}J_0\max\left\{\frac{2}{1-\varepsilon}+\frac{20(1+\varepsilon)}{10^5(1-\varepsilon)^{3/2}},\, b_0\sqrt{\frac{1+\varepsilon}{1-\varepsilon}}\right\}<0.854.$$

Differentiating the local weight and row formulas gives

$$\Delta\boldsymbol{\varpi}^{(1)}=\boldsymbol{\varphi}+2(A_{\boldsymbol{\tau}}+\textsc{Diag}\boldsymbol{\iota})(Du-\boldsymbol{\varphi}/3).$$

The feedback contribution to $\boldsymbol{\Upsilon}$ is $\mathbf{K}_0\textsc{Diag}(\boldsymbol{\Xi})\boldsymbol{\varphi}/3$, giving the asserted response. For the other terms, write $\mathbf{q}=Qg+q_{\mathrm{res}}$, where q_{res} is the contribution of $\rho-r_{\mathrm{sol}}$ in (4). We have $\|\mathbf{g}_0\|_{\mathbf{w}}\le C\beta_s$, the projection $\mathbf{Q}$ is bounded by an absolute constant in $\mathbf{w}$, and

$$\|q_{\mathrm{res}}\|_{\mathbf{w}}\le\sqrt{1+\varepsilon}\|\rho-r_{\mathrm{sol}}\|_{\mathbf{L}^{-1}}/\sqrt{\mu}\le C\beta_s.$$

Its infinity norm and that of Qg_0 are at most $Cn^{1/4}\beta_s$. Inserting these inequalities in the displayed derivative and the bounds for $A_{\boldsymbol{\tau}},\boldsymbol{\iota}$ gives $C\gamma h$. Applying the same estimates to $\boldsymbol{\Xi}\boldsymbol{\varphi}/3$ gives all remaining direction bounds. Constants may depend on the fixed ε, because $\boldsymbol{\nu}$ divides by it. □

2.3 Independent moments and high-probability safety

The direction bounds and jump floor put every mask outcome in the scalar window of Lemma 2.1. For independent mean-zero Hilbert vectors Z_i with $\|Z_i\|\le R_{\mathrm{B}}$ and $\sum_i\mathbb{E}\|Z_i\|^2\le\sigma^2$, use

$$\Pr\{\|\sum_i Z_i\|\ge t\}\le 2\exp\{-c_\star t^2/(\sigma^2+R_{\mathrm{B}}t)\}.$$

This dimension-free bound applies to matrices with the Frobenius inner product.

This is the Hilbert-space case of Pinelis' martingale inequality [Pin94, Theorem 3.4]: its Bennett exponent implies the displayed bound because $(1+x)\log(1+x)-x\ge x^2/(2(1+x))$ for $x\ge 0$. The same statement applies conditionally on a frozen pre-state; zero variance or zero packet bound gives the deterministic zero sum.

In matrix-change and observer formulas, $\boldsymbol{\sigma}_e$ denotes a relative coefficient change; centrality formulas use it for the stored representative.

Lemma 2.3. *Suppose an actual structural coefficient change is $\boldsymbol{\sigma}_e=I_e\boldsymbol{\sigma}_e^{\mathrm{sel}}$ with $|\boldsymbol{\sigma}_e^{\mathrm{sel}}|\le C\mathbf{d}_e/p_e$. This includes the squared row change from Lemma 2.1, and its absolute value. Let $\mathbf{E}=\mathbf{U}^T\textsc{Diag}(\boldsymbol{\sigma})\mathbf{U}$ and $\mathbf{E}_{\mathrm{abs}}=\mathbf{U}^T\textsc{Diag}(|\boldsymbol{\sigma}|)\mathbf{U}$. For a sufficiently large K, at a fixed valid pre-state:*

$$\|\mathbb{E}\mathbf{E}\|_{\mathrm{op}},\ \|\mathbb{E}\mathbf{E}_{\mathrm{abs}}\|_{\mathrm{op}}\le Ch,$$
$$\|\mathbb{E}\mathbf{E}\|_F,\ \|\mathbb{E}\mathbf{E}_{\mathrm{abs}}\|_F\le C\beta_f,$$
$$\mathbb{E}\|\mathbf{E}-\mathbb{E}\mathbf{E}\|_F^2,\ \mathbb{E}\|\mathbf{E}_{\mathrm{abs}}-\mathbb{E}\mathbf{E}_{\mathrm{abs}}\|_F^2\le C\beta_s^2/(K\log n).$$

With any prescribed polynomial confidence,

$$\|\mathbf{E}-\mathbb{E}\mathbf{E}\|_F\le c_1\beta_s,\quad \|\mathbf{E}\|_{\mathrm{op}}\le Ch,\quad \mathbf{E}_{\mathrm{abs}}\preceq ChI,\quad \|\mathbf{E}\|_F\le C\beta_f.$$

Also, if $\Gamma\le C_\Gamma\sum_e\boldsymbol{\lambda}_eI_ed_e/p_e$ up to a negligible deterministic numerical reserve, then

$$\mathbb{E}\Gamma\le C\beta_f\sqrt{r},\qquad \Gamma\le C_{\Gamma,\mathrm{cap}}\beta_f\sqrt{r}$$

with any prescribed polynomial confidence after increasing fixed constants and handling the finitely many bounded inputs separately.

Proof. The mean row coefficient has magnitude at most Cd_e. Since $\mathbf{U}^T\mathbf{U} = \mathbf{I}$, its operator norm is at most $C\|\mathbf{d}\|_\infty$. Put $S = \mathbf{\Pi} \circ \mathbf{\Pi}$. The elementary Schur-product identity $\textsc{Diag}(\boldsymbol{\tau}) - S = \mathbf{\Pi} \circ (\mathbf{I} - \mathbf{\Pi}) \succeq 0$ gives $\|\mathbb{E}\mathbf{E}\|_F^2 \le C\sum_e \boldsymbol{\tau}_e d_e^2 \le C\|\mathbf{d}\|_{\mathbf{w}}^2$; the same holds for absolute coefficients.

Saturated coordinates are deterministic and omitted from the variance calculation. For every nonsaturated coordinate,

$$\sum_e \textsc{Var}(\boldsymbol{\sigma}_e)\boldsymbol{\tau}_e^2 \le C \sum_{e:p_e<1} \frac{\mathbf{d}_e^2\boldsymbol{\tau}_e^2}{p_e} \le \frac{C\beta_s^2}{KLr}\sum_e \frac{\boldsymbol{\tau}_e^2}{\mathbf{w}_e} \le \frac{C\beta_s^2}{K\log n}.$$

The geometric mean of the ordinary and quadratic floors gives

$$|\boldsymbol{\sigma}_e^{\mathrm{sel}}|\boldsymbol{\tau}_e \le \frac{C\beta_s}{K\log n\sqrt{r}}\sqrt{\boldsymbol{\tau}_e/\mathbf{w}_e} \le \frac{C\beta_s}{K\log n\sqrt{r}}.$$

Independence makes cross terms in the Frobenius second moment zero. Hilbert Bernstein at $c_1\beta_s$ now has exponent $\Omega(K\log n)$, after K is chosen for the desired confidence. Apply it to both signed and absolute coefficients. The absolute matrix is positive semidefinite on every outcome, so its operator bound is the asserted order bound. A fixed number of such fields on the same ordinary graph is handled by a fixed union bound. A different graph requires its own leverage-domination premise before reusing this estimate.

For Γ, its expectation is bounded by $C\sum \boldsymbol{\lambda}_e d_e \le C\beta_f\sqrt{r}$. Each nonsaturated summand has size at most $Ch/(K\log n)$, by the ordinary floor. Its variance sum is at most $(Ch/(K\log n))C\beta_f\sqrt{r}$. Scalar Bernstein at a sufficiently large fixed multiple of $\beta_f\sqrt{r}$ has exponent

$$\Omega\left(K\log n\frac{\beta_f\sqrt{r}}{h}\right) = \Omega(K\log n\sqrt{r}/M) = \Omega(K\log nn^{1/4}/10^5).$$

Saturated contributions are deterministic and already included in the mean bound. This proves the source-normalizer statement for the original unconditioned mask. □

Lemma 2.4. *Choose $a_{\mathrm{jump}}, c_{\mathrm{sol}}$ small compared with the fixed γ, then K large. With any prescribed polynomial confidence the new residual is at most a fixed small fraction of $\beta_s\sqrt{\mu}$ in the old metric. Uniform comparison of the new and old physical metrics then gives $\|\rho'\|_{(\mathbf{L}')^{-1}} \le \beta_s\sqrt{\mu'}$. A finite norm test with a constant gap may abort the attempt on failure. Its failure probability is inverse polynomial; it need not be used for local rejection sampling.*

Proof. The selected normalized displacement is a fixed old-state value $\boldsymbol{\xi}_e^{\mathrm{sel}}$ with $|\boldsymbol{\xi}_e^{\mathrm{sel}}| \le Cd_e/p_e$. Its nonlinear mean differs from $Y_e u_e$ by at most $C(a_{\mathrm{jump}} + h)\mathbf{d}_e$, since $\mathbb{E}\mathbf{X}_e = \mathbf{d}_e$ and $\mathbb{E}\mathbf{X}_e^2 = \mathbf{d}_e^2/p_e \le a_{\mathrm{jump}}\delta_e d_e$. The physical incidence contraction is

$$\|\mathbf{B}^T\textsc{Diag}(\mathbf{z})x\|_{\mathbf{L}^{-1}} \le C\sqrt{\mu}\,\|\mathbf{x}\|_{\mathbf{w}}.$$

Hence the mean residual has norm at most $c_{\mathrm{sol}}\beta_s\sqrt{\mu} + C(a_{\mathrm{jump}} + h)\beta_f\sqrt{\mu}$.

The centered *actual displacement* is an independent vector sum. Its coordinate packet norm is at most $C\sqrt{\mu\mathbf{w}_e\ell_e}\,\mathbf{d}_e/p_e$. The quadratic floor gives total second moment $C\mu\beta_s^2/(K\log n)$, using $\sum \ell_e = r$. The geometric mean of the physical-leverage and quadratic floors gives maximum packet $C\beta_s\sqrt{\mu}/(K\log n\sqrt{r})$. Hilbert Bernstein therefore bounds this full nonlinear noise by any fixed small multiple of $\beta_s\sqrt{\mu}$ with exponent $\Omega(K\log n)$.

The scalar window gives constant comparisons of $\mathbf{z}', \mathbf{w}', \boldsymbol{\lambda}'$ with their old values. Both $\mathbf{k}, \mathbf{k}'$ lie in $[1-\varepsilon, 3+\varepsilon]$ and θ is close to one, so the new physical metric is uniformly comparable to the old. Choose the old-metric threshold with the corresponding fixed reserve. Lemma 6.4 implements the test between separated thresholds using certified inverse estimates, and includes its cost in sparse preparation. □

2.4 Expected nonlinear tracking drift

Theorem 2.5. *At a valid pre-state with $\Psi_R \le C_{R,\mathrm{out}} m$, use the law above. On the simultaneous safety event of Lemma 2.3, the tracking coordinate stays inside its tube. On a primitive safety or arithmetic first failure, assign the analytical candidate potential zero before evaluating a barrier outside its domain. The resulting nonnegative potential satisfies*

$$\mathbb{E}[\Psi_R' \mid \mathcal{F}] \le (1 - a_R)\Psi_R + C a_R m + \epsilon_{\mathrm{num}}, \qquad a_R = c_\star \beta_f / \sqrt{r}.$$

The horizon sum of ϵ_{num} can be made arbitrarily small. This is an unconditional pre-mask expectation, not an expectation conditional on passing a later phase test.

Proof. Separate each edge's own update from the inverse-Laplacian change. Write $\mathcal{G}$ for the safety event used in the Taylor calculation; whenever a nonlinear tracking bound uses $\mathcal{G}$, its expectation includes that indicator. Its nonnegative majorant is averaged under the original mask law. The signed linear term remains defined on all masks and is treated in the final paragraph.

Direction moments and the inverse remainder. All gradients below are with respect to the raw coordinate $\boldsymbol{\varpi}$. Set $H_e^R = |\mathbf{G}_e| + C_{p,\kappa}$ and $\mathbf{G}_* = \|\mathbf{G}\|_{\mathcal{N}^*}$. Coordinatewise sign invariance of the mixed norm and Lemma 2.2 give

$$\alpha_* := \sum_e H_e^R d_e \le C_{\mathrm{amp}} h(\mathbf{G}_* + B_R). \tag{13}$$

Let $\boldsymbol{\sigma}_e = (\boldsymbol{\varrho}_e'/\boldsymbol{\varrho}_e)^2 - 1$ and $\mathbf{E} = \mathbf{U}^T \textsc{Diag}(\boldsymbol{\sigma})\mathbf{U}$. Lemma 2.1 gives

$$|\boldsymbol{\sigma}_e| \le C X_e, \qquad \mathbb{E}\boldsymbol{\sigma}_e = 2D_e u_e - 2\boldsymbol{\varphi}_e/3 + r_{b,e}, \qquad |r_{b,e}| \le C(a_{\mathrm{jump}} + h)\mathbf{d}_e.$$

For a symmetric $\mathbf{T}$, define $\mathscr{D}(\mathbf{T})_e = \mathbf{U}_e T \mathbf{U}_e^T / \widehat{\boldsymbol{\tau}}_e$. Weighted Cauchy–Schwarz proves

$$\|\mathscr{D}(\mathbf{T})\|_\infty \le \|\mathbf{T}\|_{\mathrm{op}}, \qquad \|\mathscr{D}(\mathbf{T})\|_{\boldsymbol{\lambda}} \le e^{\kappa/2}\|\mathbf{T}\|_F.$$

Indeed $(\mathbf{U}_e T \mathbf{U}_e^T)^2 \le \boldsymbol{\tau}_e \mathbf{U}_e T^2 \mathbf{U}_e^T$, and summing uses $\boldsymbol{\lambda}_e/\widehat{\boldsymbol{\tau}}_e \le e^\kappa$. The moment lemma implies

$$\mathbb{E}\|\mathscr{D}(\mathbf{E})\|_{\mathcal{N}} \le Ch. \tag{14}$$

On its good event, $\|\mathbf{E}\|_{\mathrm{op}} \le Ch$ and $\|\mathbf{E}^2(\mathbf{I}+\mathbf{E})^{-1}\|_F \le Ch\|\mathbf{E}\|_F$. Thus the expected mixed norm of the resolvent remainder is $O(h^2)$, using $M\beta_f = h$.

Separating the two changes. To separate the coordinate update from the inverse change, put

$$V_e = \mathscr{D}(\mathbf{E})_e, \quad \mathrm{Rem}_e = \mathscr{D}(\mathbf{E}^2(\mathbf{I}+\mathbf{E})^{-1})_e, \quad \mathbf{c}_e = \frac{1+\boldsymbol{\sigma}_e}{1 + (\boldsymbol{\tau}_e/\widehat{\boldsymbol{\tau}}_e)\boldsymbol{\sigma}_e}.$$

The inverse identity gives

$$\frac{s + \boldsymbol{\tau}_e'}{s + \boldsymbol{\tau}_e} = \left(1 + \frac{\boldsymbol{\tau}_e}{\widehat{\boldsymbol{\tau}}_e}\boldsymbol{\sigma}_e\right)[1 - \mathbf{c}_e(V_e - \mathrm{Rem}_e)].$$

Consequently $\Delta\boldsymbol{\varpi} = \log n + S$, where $\log n, S$ are edge vectors used only in this proof. The local term is

$$L_e = \log(\boldsymbol{\lambda}_e'/\boldsymbol{\lambda}_e) - \log(1 + (\boldsymbol{\tau}_e/\widehat{\boldsymbol{\tau}}_e)\boldsymbol{\sigma}_e),$$

and the inverse term is $S_e = -\log[1 - \mathbf{c}_e(V_e - \mathrm{Rem}_e)]$. Expansion of the local scalar formulas gives

$$L_e = \mathbf{T}_e[(1 - 2\boldsymbol{\iota}_e/3)\boldsymbol{\varphi}_e + 2\boldsymbol{\iota}_e D_e u_e] + O(\mathbf{X}_e^2 + hX_e).$$

Also $|\mathbf{c}_e - 1| \le Ca_{\mathrm{jump}}$, $\|S\|_\infty \le Ch$ on the good event, and, by (14),

$$\mathbb{E}\|S - \mathscr{D}(\mathbf{E})\|_{\mathcal{N}} \le C(a_{\mathrm{jump}} + h)h.$$

The mean linear increment is therefore precisely the derivative in Lemma 2.2, up to a mixed error $C(a_{\mathrm{jump}} + h)h$. The local linear noise and the matrix linear noise have zero expectation, because $\mathbb{E}\mathbf{T}_e = 1$.

Curvature of the two terms. For fixed p, κ, direct differentiation of ψ_p gives

$$\partial^2_{\boldsymbol{\varpi}_e} \Psi_R \le C_{p,\kappa} H_e^R / \boldsymbol{\delta}_e.$$

The same bound, within a fixed factor, holds on any segment whose relative coordinate displacement is a sufficiently small fixed fraction of $\boldsymbol{\delta}_e$. The local jump floor gives $\mathbf{X}_e \le a_{\mathrm{jump}}\boldsymbol{\delta}_e$ on nonsaturated coordinates. On saturated coordinates, $\mathbf{X}_e = \mathbf{d}_e \le Ch \le a_{\mathrm{jump}}\boldsymbol{\delta}_e$ after the step-size choice. Therefore

$$\mathbb{E}\sum_e H_e^R \mathbf{X}_e^2/\boldsymbol{\delta}_e \le a_{\mathrm{jump}} \sum_e H_e^R \mathbb{E}\mathbf{X}_e = a_{\mathrm{jump}}\alpha_*.$$

The same calculation bounds the local scalar remainder by $C(a_{\mathrm{jump}} + h)\alpha_*$.

For the inverse term S, duality and the good-event infinity bound give

$$\mathbb{E}\sum_e \frac{H_e^R S_e^2}{\boldsymbol{\delta}_e} \le \frac{Ch}{\delta_{\min}}(\mathbf{G}_* + B_R)\mathbb{E}\|S\|_{\mathcal{N}} \le \frac{Ch^2}{\delta_{\min}}(\mathbf{G}_* + B_R).$$

The outer cap implies $\delta_{\min} \ge (C_{R,\mathrm{out}}m)^{-1/p}$. Thus $h/\delta_{\min} \le Cc_h n^{-1/p} P^{-16}$, which is uniformly small after choosing c_h.

Descent and affine contraction. Taylor's formula now gives

$$\mathbb{E}\Delta\Psi_R \le -(0.97 - 0.86 - C\gamma)hG_* + ChB_R + C(a_{\mathrm{jump}} + h + h/\delta_{\min})h(\mathbf{G}_* + B_R) + \epsilon_{\mathrm{num}}.$$

Choose γ small, then a_{jump} small, then K large for safety, and finally c_h small. The result is $-c_\star hG_* + ChB_R + \epsilon_{\mathrm{num}}$.

Write $W_\pi = \sum_e \pi_e$. Since $\mathcal{N}(\mathbf{x}) \le (1 + M\sqrt{W_\pi})\|\mathbf{x}\|_\infty$, the mixed unit ball contains the infinity ball of radius $(1 + M\sqrt{W_\pi})^{-1}$. Hence

$$\mathbf{G}_* \ge \frac{\|\mathbf{G}\|_1}{1 + M\sqrt{W_\pi}} \ge c_\star(\Psi_R - 2m)/(M\sqrt{r}).$$

Here $W_\pi \le Cr$, and $|\psi_p'(x)| \ge p(\psi_p(x) - 2)$ follows directly from the two powers; the raw-$\boldsymbol{\varpi}$ derivative also contains the fixed factor $1/\kappa$. Separately, $\mathcal{N}(\mathbf{x}) \ge M\|\mathbf{x}\|_\pi$ and weighted Cauchy–Schwarz give

$$\|\mathbf{1}\|_{\mathcal{N}^*} \le \frac{1}{M}\left(\sum_e \pi_e^{-1}\right)^{1/2} \le \frac{m}{M\sqrt{(1-\varepsilon)r}},$$

because $\pi_e \ge (1 - \varepsilon)s$ and $s = r/m$. Thus $B_R \le Cm/(M\sqrt{r})$, proving the affine drift.

The same-frame exceptional event. To account for $\mathcal{G}^c$, specify the signed linear term on every mask outcome:

$$\boldsymbol{\sigma}_e^{(1)} = 2D_eT_eu_e - 2\mathbf{T}_e\boldsymbol{\varphi}_e/3, \qquad \mathbf{E}_1 = \mathbf{U}^T\textsc{Diag}(\boldsymbol{\sigma}^{(1)})\mathbf{U},$$
$$L_{1,e} = \mathbf{T}_e[(1-2\boldsymbol{\iota}_e/3)\boldsymbol{\varphi}_e + 2\boldsymbol{\iota}_eD_eu_e], \qquad Z = \mathbf{G}^T[L_1 + \mathscr{D}(\mathbf{E}_1)].$$

Its all-mask expectation is exactly the coupled first-order derivative. The nonlinear row remainder is not centered: instead use the nonnegative majorant $|\mathbf{G}|^T A_{\boldsymbol{\tau}}|\boldsymbol{\sigma} - \boldsymbol{\sigma}^{(1)}|$ for its signed contribution on $\mathcal{G}$. Since $\mathbb{E}|\boldsymbol{\sigma} - \boldsymbol{\sigma}^{(1)}| \le C(a_{\text{jump}} + h)\mathbf{d}$ and $A_{\boldsymbol{\tau}}$ is bounded in the mixed norm, its full expectation is at most $C(a_{\text{jump}} + h)hG_*$. This is the previously reserved remainder. The old gradient is bounded by the outer cap and $\mathbf{X}_e \le a_{\text{jump}}\boldsymbol{\delta}_e$, so $\sup|Z| \le Ca_{\text{jump}}\|\mathbf{G}\|_1 \le C_{p,\kappa}(C_{R,\text{out}}m)^{1+1/p}$ on every mask outcome. Explicitly use

$$\mathbb{E}[Z\mathbf{1}_{\mathcal{G}}] = \mathbb{E}Z - \mathbb{E}[Z\mathbf{1}_{\mathcal{G}^c}], \qquad |\mathbb{E}[Z\mathbf{1}_{\mathcal{G}^c}]| \le \sup|Z|\Pr(\mathcal{G}^c).$$

Thus zero mean is used before good-event restriction, never after it. On the bad event the analytical new potential is zero; omitting a nonnegative remainder only helps. Choose the polynomial confidence exponent after the all-outcome polynomial bound so the displayed correction is included in ϵ_{num}. No barrier is evaluated outside its domain and no future-success conditioning is used. Any further operational or numerical kill only lowers this nonnegative matrix-killed candidate; no new signed cancellation is required. For the finite implementation of Section 3, use $\boldsymbol{\Upsilon}_c, \mathbf{d}_c$ in the all-mask expressions above. The deterministic direction bias $\boldsymbol{\Upsilon}_c - \boldsymbol{\Upsilon}$ and the final endpoint-storage errors are charged separately to ϵ_{num} by that section; neither is included in a purported centered random remainder. □

2.5 Expected central drift and its dependence on P

Theorem 2.6. *For the same single independent mask,*

$$\mathbb{E}[\Psi'_Z \mid \mathcal{F}] \le \Psi_Z - cA_Z + C_\gamma P^2 B_* + \epsilon_{\text{num}},$$

and

$$A_Z \ge c_\star \frac{P\beta_s}{\sqrt{r}}(\Psi_Z - Cm).$$

Consequently the affine contraction rate is $a_Z = cP\beta_s/\sqrt{r}$, with equilibrium $O_\gamma(Pm)$. Every coordinate remains inside its central domain before numerical rounding, on every mask outcome. No high-probability central potential inequality is required.

Proof. Use the scalar identity (9).

Curvature and waterfill bounds. The elementary explicit derivative bounds, valid for $P \ge 2$, are

$$J_e \le CPG_{Z,e}/d_{Z,e} + CP^2, \qquad G_{Z,e}d_{Z,e} \ge P(\psi_P(\overline{\boldsymbol{\varsigma}}_e) - 2).$$

For the first, split $|\overline{\boldsymbol{\varsigma}}| \le 1/P$, where both denominators give $J = \ddot{O}(1)$, from $|\overline{\boldsymbol{\varsigma}}| > 1/P$, where the larger derivative term dominates the smaller by a fixed factor. The second follows by subtracting the two explicit powers. Moreover $|(\log\psi''_P)'(x)| \le (P+2)/(1-|x|)$. The representative error and the coordinate displacement bounds therefore keep true curvatures within a fixed factor of J_e.

For finite ϑ, the literal rational waterfill and the derivative and weight comparisons give $G_{Z,e}|\mathbf{F}_e| \geq c_\star \mathbf{w}_e F_e^2/\vartheta$ on every label. The uncapped formula and actual filled radius then imply

$$\sum_{e\in\mathcal{J}} G_{Z,e}^2/\mathbf{w}_e \leq C\beta_s^2/\vartheta^2, \qquad A_Z \geq c_\star \beta_s^2/\vartheta, \qquad V \leq C(A_Z + B_*).$$

For the fully capped branch, all uncapped $G_{Z,e}$ are zero and the assertions use $\vartheta^{-1} = 0$.

Noise and one-step drift. Put $\boldsymbol{\xi} = (\mathbf{T} - 1)\boldsymbol{\nu}$. Forced and saturated coordinates have $\boldsymbol{\xi}_e = 0$. On all other coordinates the wall floor gives $|\boldsymbol{\xi}_e| \leq d_{Z,e}/k_P$. Thus $|\mathbf{F}_e + \boldsymbol{\xi}_e| \leq C d_{Z,e}/P^2$, which stays inside the true domain by (2). For each coordinate, $\mathbb{E}\boldsymbol{\xi}_e = 0$ and $\mathbb{E}\boldsymbol{\xi}_e^2 = (1 - p_e)\boldsymbol{\nu}_e^2/p_e$. The curvature floor yields

$$\mathbb{E}\sum_{e\in\mathcal{J}} J_e \boldsymbol{\xi}_e^2 \leq \frac{V}{k_P B_{\boldsymbol{\nu}}} \sum_{e\in\mathcal{J}} |\boldsymbol{\nu}_e|\sqrt{J_e} \leq V/k_P.$$

The linear-noise expectation is exactly zero: $\mathbb{E}\sum \psi_P'(\boldsymbol{\varsigma}_e)\boldsymbol{\xi}_e = 0$.

The representative reserve and the derivative bound give

$$\sum_e \psi_P'(\boldsymbol{\varsigma}_e)\mathbf{F}_e \leq -(1 - o(1))A_Z + CP^2 B_*.$$

For example the error per coordinate is at most $C\eta(G_{Z,e}/P^3 + 1/P^2)|\mathbf{F}_e|$; sum it using $\sum |\mathbf{F}_e| \leq \beta_s\sqrt{m/(3s)} = O_\gamma(B_*)$. The deterministic source curvature is at most

$$C\sum J_e F_e^2 \leq CP \max_e(|\mathbf{F}_e|/d_{Z,e})A_Z + CP^2\beta_s^2/(3s) \leq (Cc_Z/P)A_Z + CP^2\beta_s^2/(3s).$$

Since $(\beta_s^2/s)/B_* = \gamma^2\beta_f/\sqrt{r}$, the last term is at most CB_* for the chosen step size. Taylor's theorem, $(\mathbf{F} + \boldsymbol{\xi})^2 \leq 2\mathbf{F}^2 + 2\boldsymbol{\xi}^2$, and $V \leq C(A_Z + B_*)$ prove the drift when c_Z is small and k_P is large. The stated $P^2 B_*$ allowance is conservative.

Progress of the inward source. For the lower bound, put $W = \sum \mathbf{w}$ and $b_0 = \beta_s/\sqrt{W}$. The step size ensures $b_0 \leq \gamma h$ and $b_0 \leq c_Z/P^2$. The feasible comparison source has inward magnitude $\min\{c_Z d_{Z,e}/P^2, b_0\} \geq b_0 d_{Z,e}$. The waterfill maximizes the G_Z^c linear objective for the $\widetilde{\mathbf{w}}$ ellipsoid at its attained radius. Its objective is within a factor $1+\eta$ of the true objective. Scaling any $\mathbf{w}$-feasible comparison by $(1-\eta)\sqrt{(1-\eta)/(1+\eta)}$ makes it feasible for that ellipsoid and its possibly reduced filled radius. Thus its objective is at least a fixed factor of the comparison objective; the fully capped branch is immediate. Consequently

$$A_Z \geq c_\star b_0 \sum G_{Z,e} d_{Z,e} \geq cP\beta_s(\Psi_Z - Cm)/\sqrt{r}.$$

The last comparison uses the multiplicative bound between $\psi_P(\boldsymbol{\varsigma}_e)$ and $\psi_P(\overline{\boldsymbol{\varsigma}}_e)$ implied by the representative reserve. This proves the claimed affine rate and equilibrium.

Stopping on failures. The wall/source caps make the all-mask candidate's potential defined and nonnegative. Killing it at any additional matrix, field, operational, or arithmetic failure only decreases it, so the same inequality holds for the killed candidate without conditioning the noise on a good event. □

Lemma 2.7. *Assume the outer central cap $\Psi_Z \le C_{Z,\mathrm{out}} P^2 m$. For a sufficiently large constant depending only on that outer allowance and fixed parameters,*

$$B_{\boldsymbol{\nu}} = C_* P \widehat{B}_* \sqrt{1 + \beta_f^2/(\vartheta \widehat{B}_*)}$$

upper-bounds the normalizer required in (6), uniformly over $P \ge 512$. It may be enclosed upward by a fixed factor.

Proof. The central cap and power-mean inequality imply $\sum d_Z^{-2} \le Cm(C_{Z,\mathrm{out}}P^2)^{2/P} = O(m)$, uniformly in P. With $\omega = 3s$,

$$\begin{aligned}\sum_{\mathcal{J}} J_e^c/\mathbf{w}_e &\le CP \left(\sum_{\mathcal{J}} G_{Z,e}^2/\mathbf{w}_e\right)^{1/2} \left(\sum_{\mathcal{J}} 1/(w_e d_{Z,e}^2)\right)^{1/2} + CP^2 \sum_{\mathcal{J}} 1/\mathbf{w}_e \\ &\le CP\frac{\beta_s}{\vartheta}\sqrt{m/\omega} + CP^2 m/\omega.\end{aligned}$$

Multiply by $U_{\boldsymbol{\nu}}^2 = O(\beta_f^2)$, use $\sqrt{m/\omega} = \Theta(m/\sqrt{r})$ and $B_* = \beta_f m/\sqrt{r}$, and take a square root. Increasing C_* only increases probabilities, improving the moment bounds without changing the preceding drift constants. Thus outer allowances may be chosen after those constants. □

2.6 Finite arithmetic and stopping rules

Fix certified deterministic rounding before the mask. A selected candidate uses only frozen data and its actual union probability; evaluation may wait until selection. Its displacement and exact rational coefficient change are two-point functions of its own I_e, with a literal no-op at $I_e = 0$, so change packets remain independent.

Solving the scalar update. On a declared polynomial state range, the derivative bounds of Lemma 2.1 permit certified Newton iteration. Let x_* be its root in the normalized coordinate x, and start at the old point $x_0 = 0$. The initial error is at most $4X \le 4 \cdot 10^{-4}$. Evaluate the residual and its derivative with outward absolute error at most ϵ_N, divide by a derivative enclosure bounded below by 0.9, and round the iterate to the same fixed grid. Throughout the stated window the mean-value identity gives

$$|x_{j+1} - x_*| \le C|x_j - x_*|^2 + C\epsilon_N$$

with an absolute C; the sharper derivative and second-derivative bounds above make the initial error lie strictly in its quadratic convergence region. Choose ϵ_N a sufficiently small fixed fraction of the requested normalized accuracy, and perform a predetermined $O(\log\log n)$ number of stages. The final outward residual enclosure, divided by 0.9, certifies the error. Window and positive-denominator guards precede every evaluation; they do not fire under the stated scalar-window and range hypotheses. All stages use one fixed precision, so rational denominators do not accumulate. For square roots use the integer-Newton enclosure in the proof of Lemma 3.2; scalar powers use the bounded-power arithmetic of Section 12. The scalar work is $\ddot{O}(1)$ per selected label.

Precision and actual coefficients. Section 10.3 and 12 bound physical quantities, inverse margins, and every fixed-order derivative by fixed powers of n, uniformly for $P = O(\log\log n)$. Choose precision after those bounds and the horizon so the summed residual, normalizer, and drift defects are at most n^{-c}. Central identities then have this additive storage defect.

The stored metric uses exact rational $\boldsymbol{\varrho}_e^2 = \mathbf{z}_e^2/\mathbf{w}_e^2$. Polynomial sensitivity and a relative coefficient error ϵ give Loewner error ϵ and regularized-log-score error $O(\epsilon)$. Section 3 charges these errors and proves two-point moments for the actual stored-state changes, including arbitrarily small ideal directions.

Operational checks and proof-only stopping. Check the residual, Γ, ledger, and work operationally; kill the reference at the first hidden matrix-norm or observer-projection failure, under the finite arithmetic convention. The deterministic row cap gives $\|\mathbf{E}\|_{\mathrm{op}} = O(a_{\mathrm{jump}})$ and $\mathbf{I}+\mathbf{E} \succ 0$; Lemma 2.3 supplies the stronger $O(h)$ bound required by tiny tracking margins.

3 A finite implementation of the physical step

Fix the direction mesh Δ_D after the polynomial physical ranges and horizon, before the finer state and source/filter grids. Every nonzero computational direction size $d_{c,e} = |\boldsymbol{\Upsilon}_{c,e}| + |\boldsymbol{\varphi}_e|$ is at least Δ_D; when it is zero, the label's stored $(\mathbf{f}_e, \boldsymbol{\lambda}_e)$ is unchanged. Thus point-storage error can be made smaller than every nonzero $d_{c,e}$, even near cancellation, without refining Δ_D.

3.1 Finite global scalars

Use the integer P and dyadic $M, h, \gamma, \widehat{B}_*$ of Section 2. The scale B_* remains analytical: use $\widehat{B}_* \in [B_*, 2B_*]$ in V and the curvature normalizer, with an upward enclosure of the latter's square root.

Store μ on one fixed mesh. Let $\mu_{\mathrm{low}} > 0$ be its supplied lower bound throughout the expanded attempt range, and choose $\Delta_\mu \le c_\star \mu_{\mathrm{low}} \beta_s/\sqrt{r}$. At each frame choose a grid value μ' such that its actual rational ratio $\theta = \mu'/\mu$ satisfies

$$1 - \beta_s/\sqrt{r} \le \theta \le 1 - \beta_s/(2\sqrt{r}).$$

The interval length is at least $\mu_{\mathrm{low}}\beta_s/(2\sqrt{r})$, so a rounded certified midpoint has reserve. Use this μ' and its actual ratio θ throughout; both have $O(\log n)$-bit numerators and denominators.

3.2 The computational direction and its exact identities

First store the class feedback $\boldsymbol{\varphi}$ on the grid $\Delta_D\mathbb{Z}$, rounding toward zero and reserving its uniform objective error in the feedback computation. This cannot increase its mixed norm or violate a lower-wall guard. Use this actual stored $\boldsymbol{\varphi}$ in $\mathbf{g}$ and in the physical solve. The source $\mathbf{F}$, scalars μ, θ, and returned vertex direction are also their actual represented values; the solve's defect is measured against those values.

For each original label, replace the ideal direction $\boldsymbol{\Upsilon} = (\mathbf{g} - \mathbf{q})/\mathbf{k}$ by a deterministic certified scalar evaluation and grid rounding:

$$\boldsymbol{\Upsilon}_{c,e} \in \Delta_D\mathbb{Z}, \qquad |\boldsymbol{\Upsilon}_{c,e} - \boldsymbol{\upsilon}_e| \le 2\Delta_D.$$

Enclose $\boldsymbol{\upsilon}_e$ to error $\Delta_D/100$ and round the deterministic midpoint. This uses only local flow/weight data and the frozen vertex direction's endpoints, so it defines a fixed old-state function even when evaluated lazily, without a leverage-score read.

Set

$$\mathbf{d}_c = |\boldsymbol{\Upsilon}_c| + |\boldsymbol{\varphi}|, \qquad \boldsymbol{\nu}_c = \frac{\mathbf{k}\boldsymbol{\Upsilon}_c - \boldsymbol{\Xi}\boldsymbol{\varphi}/3}{\varepsilon\theta}, \qquad \delta\boldsymbol{\Upsilon} = \boldsymbol{\Upsilon}_c - \boldsymbol{\Upsilon}, \qquad \delta\boldsymbol{\nu} = \frac{\mathbf{k}\delta\boldsymbol{\Upsilon}}{\varepsilon\theta}, \qquad \mathbf{F}_c = \mathbf{F} + \delta\boldsymbol{\nu}.$$

The direction rounding introduces deterministic biases $\delta\boldsymbol{\Upsilon}, \delta\boldsymbol{\nu}$. Use $\mathbf{d}_c$ for physical jumps, quadratic rates, owner bounds, and observer bounds; use $\boldsymbol{\nu}_c$ for central wall rates, curvature rates, and scheduler bounds. Source waterfill headers and forced caps continue to use the class source $\mathbf{F}$; $\mathbf{F}_c$ is only its analytical, uniformly small correction.

Draw the original-label independent mask with these actual computational targets. The selected point map uses $v = \mathbf{T}\boldsymbol{\Upsilon}_c$ and $\zeta = \mathbf{T}\boldsymbol{\varphi}$. If $d_{c,e} = 0$, return the old stored $\mathbf{f}_e, \boldsymbol{\lambda}_e$ and their exact rational coefficients literally, regardless of any unrelated floor selecting it.

Lemma 3.1. *Before point storage rounding, the computational map satisfies*

$$\begin{aligned}\boldsymbol{\varsigma}' &= \boldsymbol{\varsigma} + \mathbf{F}_c + (\mathbf{T}-1)\boldsymbol{\nu}_c,\\ \rho' &= r_{\mathrm{sol}} + \mathbf{B}^T\mathrm{DIAG}(\mathbf{Y}\mathbf{z})(\mathbf{T}-1)\boldsymbol{\Upsilon}_c + \mathbf{B}^T\mathrm{DIAG}(\mathbf{Y}\mathbf{z})\delta\boldsymbol{\Upsilon} + \mathbf{B}^T\mathrm{DIAG}(\mathbf{z})r_f,\end{aligned}$$

where $|r_f| \le C(X_c^2 + hX_c)$ *and* $X_c = \mathbf{T}\mathbf{d}_c$.

Proof. Subtraction in the root equation gives

$$\boldsymbol{\varsigma}' - \boldsymbol{\varsigma} = \mathbf{F} - \boldsymbol{\nu} + \mathbf{T}\frac{\mathbf{k}\boldsymbol{\Upsilon}_c - \boldsymbol{\Xi}\boldsymbol{\varphi}/3}{\varepsilon\theta} = \mathbf{F} + \delta\boldsymbol{\nu} + (\mathbf{T}-1)\boldsymbol{\nu}_c.$$

The original solve still gives $\mathbf{B}^T\mathrm{DIAG}(\mathbf{Y}\mathbf{z})\boldsymbol{\Upsilon} = -\rho + r_{\mathrm{sol}}$. Insert $\mathbf{T}\boldsymbol{\Upsilon}_c - \boldsymbol{\Upsilon} = (\mathbf{T}-1)\boldsymbol{\Upsilon}_c + \delta\boldsymbol{\Upsilon}$ to get the residual identity. The scalar root proof uses the actual point arguments $\boldsymbol{\Upsilon}_c, \boldsymbol{\varphi}$ and hence gives the same remainder in terms of X_c. □

A local lower bound for the central radii. Fix a dyadic range parameter $d_{Z,\min} > 0$ that lower-bounds $1 - |\overline{\boldsymbol{\varsigma}}_e|$ throughout the declared one-step expanded range of occupied central classes. Fix the positive rational constant c_{cen} used by the central scheduler, and choose a dyadic $\rho_{\min}$ with

$$\frac{c_{\mathrm{cen}}}{4}P^{-8}\min\{d_{Z,\min}, P^{-1}\} \le \rho_{\min} \le \frac{c_{\mathrm{cen}}}{2}P^{-8}\min\{d_{Z,\min}, P^{-1}\}.$$

Every scheduler radius obtained by downward dyadic rounding within a factor two of

$$c_{\mathrm{cen}}P^{-8}\min\{1 - |\overline{\boldsymbol{\varsigma}}_e|, P^{-1}\}$$

is then at least $\rho_{\min}$ in the direction-mesh and point-interval accuracy conditions.

3.3 The deterministic perturbation budget

Since $\sum w \le Cr$ and $1 - \varepsilon \le \mathbf{k} \le 3 + \varepsilon$,

$$\|\delta\boldsymbol{\Upsilon}\|_{\mathbf{w}} \le C\Delta_D\sqrt{r}, \qquad \|\delta\boldsymbol{\nu}\|_{\mathbf{w}} \le C_\varepsilon\Delta_D\sqrt{r}, \qquad \|\delta\boldsymbol{\nu}\|_\infty \le C_\varepsilon\Delta_D.$$

Choose $\Delta_D\sqrt{r} \ll \beta_s$ and $\Delta_D \ll \rho_{\min}$ on the fixed expanded range. Let $C_{\boldsymbol{\nu}}$ be a fixed constant for the ideal bound $\|\boldsymbol{\nu}\|_{\mathbf{w}} \le C_{\boldsymbol{\nu}}\beta_f$ in Lemma 2.2. After a fixed increase of constants, $\mathbf{d}_c, \boldsymbol{\nu}_c$ retain the ideal norm and infinity bounds. Let $B_{\boldsymbol{\nu},c}$ be the normalizer of Lemma 2.7, with C_* enlarged to dominate $2C_{\boldsymbol{\nu}}\beta_f\sqrt{\sum_{\mathcal{J}} J_e^c/\mathbf{w}_e}$. Its proof uses the original class waterfill and only the noise norm, so no new dense aggregate is needed.

The mean residual defect from rounding has norm at most $C\Delta_D\sqrt{r\mu}$. The mean tracking derivative changes by $2(A_{\boldsymbol{\tau}} + \mathrm{DIAG}\boldsymbol{\iota})\mathbf{D}\delta\boldsymbol{\Upsilon}$, whose mixed norm is at most $C\Delta_D(1 + M\sqrt{r})$. The extra central deterministic linear term is bounded by

$$\left|\sum_e \psi_P'(\boldsymbol{\varsigma}_e)\delta\boldsymbol{\nu}_e\right| \le C_\varepsilon\Delta_D P d_{Z,\min}^{-1}\Psi_Z.$$

All associated quadratic terms obey the same polynomial range bounds. If the declared number of frames is n^{A_T} and these old-state derivative and norm bounds are at most n^{A_E}, choose $\Delta_D \le n^{-A_T-A_E-c_\star-10}$. Their total perturbation is then at most n^{-c}. These bounds are fixed before sampling, independently of whether the attempt later succeeds.

For the owner energy bound, add to the direct class occurrence the envelope $C_\varepsilon \Delta_D \sqrt{\mathbf{w}_e}$. Its weighted squared envelope mass is at most $Cr\Delta_D^2$, so the physical-envelope energy remains $O(\beta_f^2)$. This is the only change needed in that owner's point envelope for $\mathbf{d}_c$ and $\boldsymbol{\nu}_c$.

3.4 Stored points and exact coefficient changes

After fixing Δ_D, choose the state mesh fine enough that every selected point/root storage error, and the consequent change in its metric coefficient, contributes at most $c_\star \Delta_D$ in the normalized scalar inequalities. Polynomial bounds on $\mathbf{f}, \mathbf{z}, \mathbf{w}, \mu$ and their required derivatives permit this using a fixed number of words. The ordinary coefficient is always the exact rational function of the stored state,

$$\boldsymbol{\varrho}_e^2 = \frac{\mathbf{z}_e^2}{\mathbf{w}_e^2} = \frac{\mathbf{f}_e^2(\boldsymbol{\upsilon}_e^{\max} - \mathbf{f}_e)^2}{\mathbf{w}_e^2\big(\mathbf{f}_e^2 + (\boldsymbol{\upsilon}_e^{\max} - \mathbf{f}_e)^2\big)}.$$

Thus $\mathbf{H} = \mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2)\mathbf{B}$ retains its original meaning without materializing a dense rational sum. Every source and observer uses the exact rational change $\boldsymbol{\sigma}_e = \boldsymbol{\varrho}_e'^2/\boldsymbol{\varrho}_e^2 - 1$ between stored states.

Preserving the weight floor. Write $\Delta_{\text{state}} = 2^{-B_{\text{state}}}$ and store the weight on the shifted grid $\boldsymbol{\lambda}_e = s + \mathbf{q}_e^{\text{grid}}\Delta_{\text{state}}$ with integer $\mathbf{q}_e^{\text{grid}} \ge 0$. Round the ideal nonnegative offset upward, or round and clamp that integer offset at zero. The error is at most one mesh unit, the wall $\boldsymbol{\lambda}_e \ge s$ holds exactly, and the common denominator is $m2^{B_{\text{state}}}$. No-ops retain their stored offset. Since $\Delta_D \le d_{c,e} \le d_{c,e}/p_e$ whenever $d_{c,e} > 0$, the actual coefficient change obeys the two-point bound

$$\boldsymbol{\sigma}_e = I_e \boldsymbol{\sigma}_e^{\text{sel}}, \qquad |\boldsymbol{\sigma}_e^{\text{sel}}| \le C d_{c,e}/p_e, \qquad \mathbb{E}|\boldsymbol{\sigma}_e| \le C d_{c,e}.$$

These are the observer's two-point bounds. The central and residual point-rounding defects enter the numerical reserve fixed before the mask.

Summing the source normalizer. To compute Γ in (19), we sum enclosures of the changes while retaining the exact coefficients. Stored point operands have $W = O(\log n)$ bits. Each coefficient, old–new difference, and $\boldsymbol{\lambda}_e|\boldsymbol{\sigma}_e|$ is a fixed-depth rational expression, with numerator and denominator lengths at most $C_{\text{rat}}W$. Thus every nonzero difference or summand is at least $2^{-C_\Gamma W}$ for a fixed enlarged constant. After this bound is fixed, sum upward enclosures on a common grid of width at most $\eta_\Gamma 2^{-C_\Gamma W}$. Exact rational comparison recognizes zero; upward division gives fixed relative error using a fixed number of words per summand. Retain the exact coefficients and identity $\boldsymbol{\varrho}_e^2|\boldsymbol{\sigma}_e| = |\boldsymbol{\varrho}_e'^2 - \boldsymbol{\varrho}_e^2|$. The exact coefficients need not lie on the enclosure grid. Sum only the enclosures of $\boldsymbol{\lambda}_e|\boldsymbol{\sigma}_e|$, avoiding a dense exact rational sum.

3.5 Rational sampling rates

Sampling uses a rational enclosure of $|\boldsymbol{\nu}_c|$, whereas the point map keeps its exact noise coefficient $\boldsymbol{\nu}_c$. The next lemma supplies relative accuracy and an exact zero test without changing either physical grid.

Lemma 3.2. *Fix the direction and state grids first, and let $W = O(\log(2+n))$ bound the logical numerator and denominator lengths of their stored operands, the represented vertex direction, and the rational design constants. In particular take $\varepsilon > 0$ rational. For any fixed rational $0 < \eta_{\mathrm{rate}} < 1$, there is a deterministic point routine returning a rational $\boldsymbol{\nu}_e^+$ such that*

$$|\nu_{c,e}| \le \boldsymbol{\nu}_e^+ \le (1+\eta_{\mathrm{rate}})|\nu_{c,e}|, \qquad \boldsymbol{\nu}_e^+ = 0 \quad \Longleftrightarrow \quad \nu_{c,e} = 0.$$

It also determines the exact sign of $\nu_{c,e}$, uses $O(\log W) = \ddot{O}(1)$ word operations on a fixed number of words, and has no numerical background rate. Its precision is chosen after W and does not refine either physical storage grid.

Proof. Suppress the label and set

$$\mathbf{Z} = \mathbf{z}^2, \qquad \mathbf{R} = (\mathbf{c} - \mathbf{B}\mathbf{y})/(\mu \mathbf{w}), \qquad \mathbf{S} = \mathbf{f}^{-3} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-3}, \qquad \mathbf{R}_0 = \mathbf{f}^{-1} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-1}.$$

All four quantities are rational fixed-depth expressions in the stored operands, and $Z > 0$. The shared positive square root $\mathbf{z} = \sqrt{\mathbf{Z}}$ gives

$$\boldsymbol{\Xi} = \mathbf{R}\sqrt{\mathbf{Z}}, \quad \overline{\boldsymbol{\Xi}} = \mathbf{R}_0\sqrt{\mathbf{Z}}, \quad \mathbf{D} = \mathbf{Z}\mathbf{S}\sqrt{\mathbf{Z}}, \quad \mathbf{k} = \mathbf{Z}^2\mathbf{S}\mathbf{R} + \mathbf{Y}(1 - \mathbf{Z}^2\mathbf{S}\mathbf{R}_0).$$

Thus $\mathbf{k}$ is evaluated by its exact rational expression, and

$$\boldsymbol{\nu}_c = \alpha_{\boldsymbol{\nu}} + \beta_{\boldsymbol{\nu}}\sqrt{\mathbf{Z}}, \qquad \alpha_{\boldsymbol{\nu}} = \frac{k u_c}{\varepsilon\theta}, \qquad \beta_{\boldsymbol{\nu}} = -\frac{\mathbf{R}\boldsymbol{\varphi}}{3\varepsilon\theta}.$$

Choose $H_{\mathrm{bit}} = C_{\mathrm{bits}} W \ge 1$ so that the numerators and positive denominators of $\alpha_{\boldsymbol{\nu}}, \beta_{\boldsymbol{\nu}}, Z$ have absolute value at most 2^H. The constant C_{bits} is fixed: these expressions have fixed depth, and the reduced cost uses only the two endpoint coordinates of $\mathbf{y}$.

Exact zero and sign. If $\alpha_{\boldsymbol{\nu}} = \beta_{\boldsymbol{\nu}} = 0$, return zero. If exactly one vanishes, use the sign of the other, since $Z > 0$. If both have the same nonzero sign, that is the sign of their sum. Otherwise compare the rational numbers $\alpha_{\boldsymbol{\nu}}^2$ and $\beta_{\boldsymbol{\nu}}^2\mathbf{Z}$ by exact cross multiplication. Equality means $\boldsymbol{\nu}_c = 0$; otherwise its sign is that of $\alpha_{\boldsymbol{\nu}}$ when $\alpha_{\boldsymbol{\nu}}^2 > \beta_{\boldsymbol{\nu}}^2\mathbf{Z}$, and that of $\beta_{\boldsymbol{\nu}}$ when $\alpha_{\boldsymbol{\nu}}^2 < \beta_{\boldsymbol{\nu}}^2\mathbf{Z}$. The sign test is necessary: equality of the squares alone does not imply cancellation when the signs agree.

Separation from zero. There is a uniform nonzero separation bound

$$\boldsymbol{\nu}_c \ne 0 \quad \Longrightarrow \quad |\boldsymbol{\nu}_c| \ge 2^{-8H_{\mathrm{bit}}}.$$

For terms of one sign this follows from their rational lower bounds and $\sqrt{Z} \ge 2^{-H_{\mathrm{bit}}/2}$. In the opposite-sign nonzero case, $\alpha_{\boldsymbol{\nu}}^2 - \beta_{\boldsymbol{\nu}}^2\mathbf{Z}$ is a nonzero rational with denominator at most $2^{5H_{\mathrm{bit}}}$, and

$$|\boldsymbol{\nu}_c| = \frac{|\alpha_{\boldsymbol{\nu}}^2 - \beta_{\boldsymbol{\nu}}^2\mathbf{Z}|}{|\alpha_{\boldsymbol{\nu}} - \beta_{\boldsymbol{\nu}}\sqrt{\mathbf{Z}}|} \ge \frac{2^{-5H_{\mathrm{bit}}}}{2^H + 2^{3H_{\mathrm{bit}}/2}} \ge 2^{-8H_{\mathrm{bit}}}.$$

A certified square-root enclosure. Write $Z = N_1/N_2$ with positive integers N_1, N_2, and take $R_* = 9H_{\mathrm{bit}} + \log_2(16/\eta_{\mathrm{rate}})$. Form

$$\mathbf{Q} = \left\lfloor \frac{2^{2R_*} N_1}{N_2} \right\rfloor, \qquad j = \lfloor \sqrt{\mathbf{Q}} \rfloor.$$

Then $[j2^{-R_*}, (j+1)2^{-R_*}]$ encloses $\sqrt{\mathbf{Z}}$. All these integers have $O(H_{\mathrm{bit}})$ bits. Binary search on bit positions finds an initial power-of-two upper bound for $\sqrt{\mathbf{Q}}$ in $O(\log H_{\mathrm{bit}})$ word operations. Integer Newton iteration $x \leftarrow \lfloor (x + \lfloor \mathbf{Q}/x \rfloor)/2 \rfloor$, started at that upper bound and stopped when the next value does not decrease, returns j in $O(\log H_{\mathrm{bit}})$ iterations. The usual Newton error-squaring bound, followed by the final integer check, proves this count. It does not use a unit-cost square-root or most-significant-bit instruction.

Multiply the enclosure by the signed rational $\beta_{\boldsymbol{\nu}}$ and add $\alpha_{\boldsymbol{\nu}}$ exactly. Its width is at most $2^H 2^{-R_*} \le \eta_{\mathrm{rate}} 2^{-8H_{\mathrm{bit}}}/16$. For a nonzero value, take the upper endpoint of its absolute value using the already determined exact sign. It is positive, rational, and at most $(1 + \eta_{\mathrm{rate}}/16)|\boldsymbol{\nu}_c|$. If a common dyadic output grid is desired, round upward on a grid of width at most $\eta_{\mathrm{rate}} 2^{-8H_{\mathrm{bit}}}/16$; the displayed lemma still holds. For the exact-zero case return zero before the square-root calculation. All rational expressions have $O(H_{\mathrm{bit}})$-bit numerators and denominators, so a fixed enlargement of the eventual arithmetic word suffices. □

Using the enclosure. Use $\boldsymbol{\nu}_e^+$ only in sampling rates, with constant-factor rational upper enclosures of $B_{\boldsymbol{\nu},c}$ and $\sqrt{J_e^c}$. The resulting rational probabilities bound the clipped wall and curvature targets upward within fixed factors. The union of these fixed many components has exact rational probability p_e, determined by frozen data even when evaluated lazily after a proposal.

The point map still uses $\mathbf{T} = \mathbf{I}/p$, $v = Tu_c$, and $\zeta = \mathbf{T}\boldsymbol{\varphi}$; its noise coefficient remains $\boldsymbol{\nu}_c$ in Lemma 3.1. Increasing the floors preserves safety and variance bounds, and the fixed-factor increase preserves weighted-rate work bounds.

Choose the rate grid after the logical state and direction bit bounds, before the exact-stream input exponent and final word width. It refines only rational probabilities; the state grid and $\boldsymbol{\Upsilon}_c$ stay fixed.

The direction also has the form $(\alpha_{\boldsymbol{\nu}}\sqrt{\mathbf{Z}} + \beta_{\boldsymbol{\nu}})/\mathbf{k}$ with new rational scalar coefficients $\alpha_{\boldsymbol{\nu}}, \beta_{\boldsymbol{\nu}}$. Its nonzero separation depends on the stored-flow mesh and may be a higher power of that mesh. Fixing Δ_D first avoids circular refinement; its perturbation was charged above.

4 Original-edge cores and adaptive pruning

To avoid full original-edge scans at every recursive level, a degree guard supplies known references and a game on a fresh sparse sample finds a cut or a near-expander. Trimming and degree-bias repair produce an original induced core; a final pass assigns original labels. The repair uses weighted pruning charged by stored sample atoms. Installed cores use unweighted pruning charged per genuine deletion at a density-dependent price. Both pruners keep their birth references.

Original graphs are simple and undirected. An original label is an edge, stored with two direct incidence handles. A constant number of simple families may be treated separately, for example the two orientations of a simple directed graph. A weighted sparse atom is stored once regardless of its weight. Loops count once in degree and zero in cuts. References are positive integers fixed at core birth. A *ghost arc* is a deleted edge retained only in the pruner's private flow state. For a reference r, write $r(X) = \sum_{u \in X} r_u$, $D_r = \textsc{Diag}(r)$, $\|\mathbf{x}\|_r^2 = x^T D_r x$, and $\|\mathbf{x}\|_{r,1} = \sum_u r_u |x_u|$. For a row embedding X the squared norm sums over its columns. Operator Hilbert–Schmidt norm is $\|\mathbf{T}\|_{\mathrm{HS},r} = \|D_r^{1/2} T D_r^{-1/2}\|_F$.

An induced h-core has $\partial_{G[A]}(U) \ge h \min\{r(U), r(A \setminus U)\}$. A near-expander in a larger graph $\mathbf{G}$ has instead $\partial_G(U) \ge h r(U)$ for $U \subseteq A$ with $r(U) \le r(A)/2$. The game returns near-expanders; the subsequent conversion must certify expansion inside the retained set. Original-label access counts row scans, probes, and all selected sample positions, including rejected ones; arithmetic, sparse

processing, and indexing are charged separately. All words have $O(\log n)$ bits, which includes a fixed polynomial call horizon; the desired overall failure probability is $\delta = n^{-O(1)}$. Constants in $n^{o(1)}$ are absolute and can be large.

The only graph-algorithm inputs are the sparse approximate Cholesky construction of Kyng–Sachdeva, including its low-degree elimination variant [KS16]; standard dynamic-tree link, cut, path-minimum, path-add, and edge-value operations [ST83]; and deterministic exact flow for integral data in $m^{1+o(1)}$ time with logarithmic capacity and cost factors [vdBCP+23, Theorem 1.1]. Flow is used only on explicit sparse networks. We also use elementary independent scalar and matrix concentration. Appendix A.5 supplies finite arithmetic, the sparse flow reduction, and work caps.

For $\mathbf{q} > 0$, define the density deficit $D(\mathbf{q}) = 1 + \max\{0, \log_2(n^2/\mathbf{q})\}$. Each residual phase processes the preceding phase's uncovered edges, starting with the full input. Write h_j for the reference-conductance guarantee of its phase-j cores.

Theorem 4.1. *Let $\mathbf{G}$ be a frozen simple graph on at most n vertices with $\mathbf{q}$ labels. There is a construction with polynomial confidence that partitions its labels into cores and an explicit residual of size at most $\boldsymbol{\tau} = n^{9/8}$. Put $D_j = D(q_j)$ for the size $q_j > \boldsymbol{\tau}$ of each residual phase. For a fixed absolute A, put $S_j = AD_j^2$. Every phase-j core has a known reference r, retained from its local construction, such that*

$$h_j \geq cD_j^{-5}, \qquad \min r_u \geq q_j/(4nS_j), \qquad \deg_{\mathrm{core}}(\boldsymbol{\Upsilon}) \leq r_u.$$

The construction uses

$$O(\mathbf{q} + n\log(2 + \mathbf{q}/\boldsymbol{\tau})) \quad \textit{original-label work}, \qquad n^{1+o(1)} \quad \textit{additional work}.$$

Total core vertex population is $n^{1+o(1)}$. On installing a core, use its actual original degrees as immutable birth reference and round ϕ down within a factor two of h_j. There is a deterministic structure for any adaptive sequence of genuine deletions, exporting $O(U/\phi)$ live labels after U such deletions, with total work

$$O(q_{\mathrm{birth}} + v_{\mathrm{birth}} + H_{\mathrm{dep}} + UD_j/\phi^2), \qquad H_{\mathrm{dep}} = O(D_j/\phi).$$

Every retained live cut has conductance at least $\phi/4$ in that birth reference. Exports are not fed back as genuine deletions.

Here $q_{\mathrm{birth}}, v_{\mathrm{birth}}$ are birth sizes and H_{dep} is pruning height. Cores are induced in their creating phase's residual. For $\mathbf{q} = 0$ return nothing. Construction is Monte Carlo with a deterministic work cap; pruning is deterministic conditional on a valid birth certificate. Appendix A proves these two parts in order.

5 Maintaining original edges within an attempt

An original label's update cost depends on its current core. We bound current price moments, price fresh proposals, and charge historical copying under an admission cap. This includes small records with expensive deletions. All prices are frozen before the mask.

Put $m = \Theta(n^2)$, $r = \Theta(n)$, $s = r/m$. All implicit constants are independent of n. Bounded parallel-label multiplicity permits a fixed number of simple layers before coefficient bucketing. In each of the fixed number of metrics, every label has one live analytical owner: its current core or explicit route.

Core guarantees used below. Use the static cover and fixed-birth pruner of Theorem 4.1, with residual cutoff $\tau_0 \asymp n^{9/8}$ and its stated birth-work and population bounds. The cover-phase sizes q_j decrease by at least a factor two (Lemma A.11). Write $D_j = D(q_j)$; then $\phi_j^{-1} \le CD_j^5$. Per genuine deletion, charge $CD_j\phi_j^{-2}$ for pruning and $C\phi_j^{-1}$ exported live labels, through retirement. Use the theorem's immutable actual-birth-degree references and current cut certificates.

Store coefficients in dyadic amplitude buckets. Choose the fixed-reference weighted median g_C from the stored field coordinates, breaking ties by vertex ID, and share it at both endpoints. This needs no gauge rounding. Apply any common scalar normalization to the median too; constant fields then have exactly zero centered differences. Let F_{eo} be the coefficient-scaled centered field value at endpoint occurrence o of label e; an explicit route instead uses the scaled edge difference. Apply Cheeger's inequality to the core completed to its birth degrees by loops, centered at the birth-reference median, to get

$$\sum_{e,o} \phi_{eo}^2 F_{eo}^2 \le C\|v\|_K^2 \tag{15}$$

for the appropriate coefficient graph K and vertex field v. If the core containing e was created in cover phase j, write $D_e = D_j$. An explicit route uses its actual edge difference, with $D_e = 1$ and $\phi = 1$ in all price and moment bounds. Header subdivisions do not increase the fixed number of analytical occurrences per label.

5.1 Current prices in the mixed-radix counter

Use powers of two $\tau_0 \asymp n^{9/8}$, $R \asymp n^{1/16}$ and $T_0 \asymp m/\log^{1024} n$. The cutoff limits the population of expensive small records; the lower radix limits the number of carries paid by each token. The finitely many n for which $T_0 < \tau_0$ are handled by the bounded-input branch. Historical admission tokens remain as tombstones after deletion. A genuine admission or pruner export enters the pending list of its current amplitude bucket and receives one token; carries only regroup existing tokens. In each bucket, full τ_0-token blocks enter a counter. Below T_0 its radix is R, except for one final power-of-two radix in $[2, R]$; above T_0 it is binary. A carry rebuilds on the children's current live labels and discards their private flow state. A record's final explicit residual stays in that record until a real carry.

Bulk initialization. In a bucket with q_b labels, write $\lfloor q_b/\tau_0 \rfloor$ in the mixed radices. Partition the labels into records of the indicated capacities and a pending list of fewer than τ_0 labels. Build each record's cover once, with birth clocks and handles from the current coefficients. Since all labels are live, this directly constructs the final sequential-admission state.

The initial record count is $O(R + \log n)$ per amplitude bucket. Original label scans sum to $O(m)$, and sparse cover work is $n^{17/16+o(1)}$; the explicit routes and proxy/header preparation have the larger $n^{19/16+o(1)}$ allowance. Grouping labels by bucket and installing scalar/deadline metadata adds $\ddot{O}(m)$. A clean checkpoint refreshes references before zeroing and saving debts. A retry restores them literally, without random fine reads, and may rebuild the owner while preserving the saved inner certificates. Within an attempt, counters persist; the initialization and historical-token charges are separate. Let N count all admission tokens so far, including initial tokens and tombstones.

Lemma 5.1. *Under any predeclared polynomial admission stop $N \le n^A$, for every fixed $0 \le j \le 512$,*

$$\sum_{e \text{ current}} D_e^j \le C_{A,j} m (\log\log n)^j.$$

Give a deletion from an explicit route the price one, and replace every core deletion price by its maximum with one. If W_e sums these old deletion prices over all affected metric owners, then

$$\sum W_e^{12} = \ddot{O}(m), \qquad \sum_{e,o} \phi_{eo}^{-12} = \ddot{O}(m), \qquad \max W_e + \max \phi_{eo}^{-1} = n^{o(1)}.$$

These are pathwise statements about the frozen state before a mask. They do not require an event count or a bound on historical copying.

Proof. For consecutive lower capacities c_j and radices r_j,

$$\sum_{c_j < T_0} (r_j - 1)c_j = T_0 - \tau_0.$$

Adding the unfinished pending block shows that all lower records have total capacity less than T_0 per amplitude bucket. There are a constant number of lower levels: the ratio from τ_0 to m has exponent $7/8$, whereas one full radix has exponent $1/16$. Thus their record count is $O(R)$ per bucket. There are $O(\log n)$ amplitude buckets and $O_A(\log n)$ upper records per bucket under the polynomial stop.

Every live label in a cover phase with $q_j \geq T_0$ has $D_j = \ddot{O}(1)$. After an upper record's first cover phase below T_0, the remaining input sizes sum to less than $2T_0$. The lower-capacity population and these upper-record tails therefore contain at most $C_A \log^2 nT_0$ labels. Their deficits are at most $C \log n$. It follows that

$$\sum D_e^j \leq C_j m(\log\log n)^j + C_{A,j} \log^2 nT_0 \log^j n \leq C_j m(\log\log n)^j + C_{A,j} m \log^{j+2-1024} n.$$

Each deletion price is at most CD^{11}; constant-layer sums and the fixed occurrence count give the remaining assertions without independence between metrics. □

5.2 Pricing proposals, updates, and reads

Price each raw proposal for its own work, structural callbacks, and an old-margin fine read. The valid-state bounds of Section 2 give

$$\mathbf{w}_e \geq c_\star s, \quad \sum \overline{\boldsymbol{\delta}}_e^{-p} \leq Cm, \quad \sum d_{Z,e}^{-P} \leq CP^2 m, \quad p = 512, \quad P \geq 512,$$

and $\|\mathbf{d}\|_{\mathbf{w}} + \|\boldsymbol{\nu}\|_w \leq C\beta_f$. Use the stored certified tracking scale with $C_\delta^{-1}\delta \leq \overline{\boldsymbol{\delta}} \leq C_\delta \delta$ for the fixed known comparison constant C_δ introduced there. For every fixed order $j \leq 36$, the central premise implies $\sum d_Z^{-j} \leq m(CP^2)^{j/P} = O_j(m)$. Retain the source multiplier ϑ, central derivative magnitude $G_{Z,e}$, certified curvature upper bound J_e^c, and uncapped set $\mathcal{J}$ from the physical step. For a filled source the finite waterfill gives

$$\|G_{Z,\mathcal{J}}/\sqrt{w}\|_2 \leq C\beta_s/\vartheta, \qquad J_e^c \leq CPG_{Z,e}/d_{Z,e} + CP^2.$$

In the fully capped branch set $1/\vartheta = 0$; $G_{Z,e} = 0$ on $\mathcal{J}$, even when the source norm lies in the filled range. Put

$$B_* := \beta_f\sqrt{m/s} = \beta_f m/\sqrt{r}, \quad \zeta_* := \beta_f^2/(\vartheta B_*), \quad V := B_* + \beta_s^2/\vartheta, \qquad \beta_s = \gamma\beta_f.$$

The implemented upward normalizer can be chosen within a fixed factor of $PB_*\sqrt{1+\zeta_*}$, so that

$$B_{\boldsymbol{\nu}} \geq C\beta_f \sqrt{\sum_{\mathcal{J}} J_e^c/\mathbf{w}_e}, \qquad B_{\boldsymbol{\nu}}/V \leq C_\gamma P(1+\zeta_*)^{-1/2}. \tag{16}$$

This is Lemma 2.7 and its finite-direction extension in Section 3: use $\widehat{B}_* \asymp B_*$ and $V = B_*(1+\gamma^2\zeta_*)$.

The physical base owner has coefficient $\mathbf{z}^2/\mathbf{w}$. Its two endpoint occurrences, together with one direct source-class occurrence, give nonnegative F_{eo} with

$$\sqrt{\mathbf{w}_e}\max\{d_e, |\boldsymbol{\nu}_e|\} \le C\sum_o F_{eo}, \qquad \sum_{e,o} \phi_{eo}^2 F_{eo}^2 \le C\beta_f^2. \tag{17}$$

For the endpoint terms take $\mathbf{z}_e|v_y(u) - g_C|/(\mu\sqrt{\mathbf{w}_e})$. For the direct term take a fixed multiple of $\sqrt{\mathbf{w}_e}(|1-\theta| + |F_e| + |\boldsymbol{\varphi}_e|)$. The direction estimates and (15) prove both claims, with fixed envelope factors for coefficient buckets.

For the implemented direction, replace $d, \boldsymbol{\nu}$ by $d_c, \boldsymbol{\nu}_c$ and add $C_\varepsilon\Delta_D\sqrt{\mathbf{w}_e}$ to this direct occurrence. Section 3 gives $|u_c - u| \le 2\Delta_D$ and $|\boldsymbol{\nu}_c - \boldsymbol{\nu}| \le C_\varepsilon\Delta_D$. This adds squared energy $O(\Delta_D^2 \sum w) = O(r\Delta_D^2)$; the direction mesh makes its stopped weighted charge negligible. It adds no graph occurrence or score read. Below use these actual directions and the enlarged direct envelope.

The combined point price R_e below is distinct from the counter radix R.

Lemma 5.2. *Freeze all old core identities, margins, classes and prices, and put*

$$R_e = 1 + W_e + \overline{\boldsymbol{\delta}}_e^{-3}.$$

The three dense component families

$$\min\{1, C_\delta d_e/(a_{\mathrm{jump}}\overline{\boldsymbol{\delta}}_e)\}, \quad \min\{1, k_P|\boldsymbol{\nu}_e|/d_{Z,e}\}, \quad \min\{1, k_P B_{\boldsymbol{\nu}}|\boldsymbol{\nu}_e|\sqrt{J_e^c}/V\}, \qquad k_P = CP^{12},$$

where the last two families are restricted to the uncapped source set $\mathcal{J}$, have total conditional R-priced raw proposal expectation at most $\ddot{O}(B_*)$. The same bound pays the forced source-wall cap list. The constants may depend on the fixed $a_{\mathrm{jump}}, \gamma, C_\delta$. Every actual selected label's W-priced structural callbacks and its old-margin fine read are paid by these prices or by the sparse component rows below.

Proof. The current moment lemma and the tracking cap give $\sum R^{12} = \ddot{O}(m)$. Hölder, with exponents $3, 3, 3$, yields

$$\sum R^4\phi^{-4}d_Z^{-2} = \ddot{O}(m), \qquad \sum R^2\phi^{-2} = \ddot{O}(m).$$

All sums include the fixed number of direct/endpoint occurrences. Using the derivative bound and then the uncapped source estimate,

$$\begin{aligned} S_R &:= \sum_{e\in\mathcal{J}}\sum_o R_e^2\phi_{eo}^{-2}J_e^c/\mathbf{w}_e \\ &\le CP\|G_{Z,\mathcal{J}}/\sqrt{w}\|_2\left(\sum R^4\phi^{-4}/(wd_Z^2)\right)^{1/2} + CP^2\sum R^2\phi^{-2}/w \\ &\le C_\gamma P^2(\log\log n)^{32}(m/s)(1+\zeta_*). \end{aligned}$$

Equations (17) and (16) therefore bound the raw curvature envelope by

$$Ck_P(B_{\boldsymbol{\nu}}/V)\beta_f\sqrt{S_R} \le C_\gamma k_P P^2(\log\log n)^{16}B_*.$$

The margin families instead use $\sum R^2\phi^{-2}\overline{\boldsymbol{\delta}}^{-2}$ and $\sum R^2\phi^{-2}d_Z^{-2}$. Their respective sixth moments and Hölder bound both sums by $\ddot{O}(m)$; Cauchy–Schwarz in the owner energy bound completes their price bounds.

If a nonzero source is capped at $c_Z d_Z/P^2$, its actual rounded magnitude is within a fixed factor of that cap. Consequently $1 \le CP^2|F_e|/d_Z$. Weighted Cauchy gives the forced price

$$CP^2 \sum R_e|F_e|/d_{Z,e} \le CP^2\beta_s s^{-1/2}(\sum R^4)^{1/4}(\sum d_Z^{-4})^{1/4} \le CP^2(\log\log n)^{11}B_*.$$

The flat γh cap is charged in the sparse rows.

Use homogeneous incidence lists and exact original-label thinning. The untruncated envelopes pay the frozen R_e price of every raw success, including rejected and duplicate proposals. Each label has only a fixed number of list occurrences. Coalescing the component masks by original label can only reduce cost. Each retained label gets one final packed read at cost $\ddot{O}(\bar{\delta}_{e,\mathrm{old}}^{-3})$, by the supplied one-frame margin comparison. Freeze the whole batch's old owners before its coins and detach all genuine changes before pruning or carrying, so that later bits never acquire a newly chosen price. □

The sparse component rows. The ordinary floor $K\log n\boldsymbol{\lambda}$ has raw mass $O(r\log n)$ using the actual-$\boldsymbol{\lambda}$ bins. The quadratic floor $KLrwd^2/\beta_s^2$ has raw mass $rn^{o(1)}$ by the squared version of (17) and $\max\phi^{-2} = n^{o(1)}$. Lemma 6.2 supplies a physical-leverage envelope of mass $n^{19/16+o(1)}$, including cores and explicit labels, without a dense leverage table. The number of nonzero flat source caps is at most $C\beta_s^2/(s\gamma^2h^2) = O(n^{1/2})$.

All these rows may pay the worst point price $R_{\max} \le n^{6/p+o(1)}$. For at most $T = n^{3/4+3/p+o(1)}$ frames their largest exponent is $3/4 + 3/p + 19/16 + 6/p = 1001/512 < 2$. List starts, medians, class-header scans and global alias construction cost $[n(R+\log n) + \tau_0(R+\log n)]n^{o(1)}$ per field, not per output. Additional inverse, source, and observer fields use their preparation bounds below. Numerical background rates have summable total weighted mass at most one.

5.3 Admissions and historical copying

Lemma 5.3. *Suppose a clean attempt has at most the preceding T frames and its predictable physical budget is $\sum B_* \le Cm$. Then its structural, selected-list and metadata work has expected cost $\ddot{O}(m)$, with an additional constant attempt-abort probability as small as desired. The statement uses a deterministic sharp token cap and actual expected scalar costs.*

Proof. Expected genuine-change cost. First impose the coarse stop $N \le n^A$, with A fixed large. This is enough for the predictable current moments and hence the weighted raw-rate theorem. Tower expectation at the predictable event that a frame begins gives

$$\mathbb{E}Z \le \ddot{O}(m), \qquad Z = \sum_{\text{genuine}} W_e.$$

The sparse rows are smaller. Normalize core deletion prices to dominate $\max\{1, D_C\phi_C^{-2}\}$ and give explicit routes price one. There are $N_{\mathrm{init}} = O(m)$ initial tokens. A core receiving U genuine deletions exports at most $16U/\phi_C$ labels in pruning and $64U/\phi_C$ at retirement (Theorem A.13). Carries mint no tokens, and genuine reinsertions number at most the genuine metric deletions. Summing gives the cumulative certificate

$$N \le N_{\mathrm{init}} + 81Z. \tag{OA.1}$$

A retirement may spend deletion credit accumulated in earlier batches.

A deterministic admission cap. Choose a fixed polynomial $B(P, \log\log n)$ and a constant C_0 so that $N_{\mathrm{sharp}} = C_0 m B(P, \log\log n)$ exceeds an arbitrarily large constant multiple of the displayed expectation. Before structural commit, use the batch's prices frozen before its mask to compute $Z_{\mathrm{pros}} = Z + \sum_{e\in\mathrm{batch}} W_e$. Abort if $N_{\mathrm{init}} + 81 Z_{\mathrm{pros}} > N_{\mathrm{sharp}}$. Otherwise (OA.1) bounds every commit prefix by N_{sharp}. Compute the test from the enumerated genuine batch in constant work per affected metric occurrence. Also guard every insertion and resulting carry by the literal token count, on every outcome.

Markov's inequality on the coarse-stopped reference, including the proposed crossing batch, bounds this guard's failure probability by the desired small constant without conditioning masks on survival. For large n, the sharp cap is below the coarse cap.

Copying and sparse births. Each historical token participates in constantly many lower carries and at most $1+\log_2(1+N_{\mathrm{sharp}}/T_0) = \ddot{O}(1)$ upper carries. Thus literal live-label copies and original cover passes are pathwise at most $CN\log\log n$. The number of births is at most $2N/\tau_0$. Their sparse construction work is bounded using the sharp cap by

$$(N_{\mathrm{sharp}}/\tau_0) n^{1+o(1)} = n^{15/8+o(1)} = o(m).$$

Retiring arrays and private state is charged to their birth sizes. Central, source and tracking class moves use direct handles, paid by displacement or an already-paid read, and never enter a pruner. Owner migrations preserve signed clock-minus-baseline debts and due flags.

Scalar costs and the final work cap. List and signed-deadline updates cost $\ddot{O}(1)$ per touched label; header scans are paid at field preparation. Exact stream enumeration and rational thinning cost conditional expected $\ddot{O}(1)$ per start and raw proposal. Combine these with $\mathbb{E}Z$, then impose the attempt work cap. Markov gives a small constant exhaustion probability; exhaustion aborts the entire attempt. □

Sections 3 and 8 implement the scalar operations and encoded update. This section pays selected reads and admissions; the observer pays autonomous reads. Lemma 14.1 combines their accounts.

6 Finite sparse inverses and physical solves

Fixed sparse inverses support physical directions and later source fields. We implement low-degree approximate Cholesky [KS16, Theorem B.1], sample the physical right-hand side, and test candidates through an exact residual cache. Numerical ranges use the ambient n, including for small child graphs.

6.1 Constructing and applying a fixed sparse inverse

Lemma 6.1. *Let P be a connected graph Laplacian with q stored edge copies, positive common-dyadic-grid weights in $[n^{-A}, n^A]$, and at most n vertices. Assume also $q \le n^A$ and that its common input mesh is $2^{-b_{\mathrm{in}}}$ with $0 \le b_{\mathrm{in}} \le A\log_2 n$; enlarging the fixed A covers all three bounds. For fixed $A, c_\star > 0$, one can construct, except with probability n^{-c}, a fixed symmetric operator T_P satisfying*

$$(1-\epsilon)P^\dagger \preceq T_P \preceq P^\dagger \quad \text{on } \mathbf{1}^\perp.$$

For constant ϵ, construction and application use $qn^{o(1)}$ word work; for $\epsilon = n^{-O(1)}$ they still use this bound with another fixed logarithmic exponent. A finite application can have any prescribed

inverse-polynomial absolute energy error on a supplied polynomial-range common-grid vector using $O(\log n)$-bit words. The analytical T_P is linear and symmetric; its rounded evaluation need not be linear.

Proof. Elimination and exact sampling. Write v for the number of vertices of P. The one-vertex case returns the zero map in constant work; the stated sparse-work bounds concern the nontrivial connected case $q \geq 1$. Split each input edge into $\rho = \Theta_c(\log^2 n)$ equal copies, choosing ρ to be a power of two. A split copy has leverage at most $1/\rho$ relative to P. At each elimination choose uniformly among live vertices of multiedge degree at most twice the average. Rejection from uniform live-vertex draws implements this rule with success probability at least $1/2$.

Use the published clique sampler with its incident list and weighted distribution frozen for the whole elimination bag. If it emits a harmonic weight $h = w_1 w_2/(w_1 + w_2)$, replace it by $\delta\lfloor h/\delta \rfloor$, discarding zero copies. The fine dyadic mesh δ divides the input mesh after splitting. All weighted choices use the actual current integer grid weights; integer aliases and uniform-integer rejection give exact choices.

Concentration with downward rounding. In original vertex coordinates, let Z_j be the sum of the emitted rank-one factor terms and the current rounded Schur graph. The clique expectation identity on that current graph gives

$$Z_j = P + \mathcal{M}_j - \mathcal{D}_j, \qquad \mathcal{D}_j \succeq 0,$$

where $\mathcal{M}_j$ sums the centered unrounded clique samples and $\mathcal{D}_j$ sums the downward edge errors in original coordinates. Later triangular factors do not conjugate these errors. If the two incident copies have leverage at most $1/\rho$, resistance triangle inequality gives

$$hR_P(u,v) \leq \frac{w_1 w_2}{w_1 + w_2}\left(\frac{1}{\rho w_1} + \frac{1}{\rho w_2}\right) = 1/\rho.$$

Rounding preserves this leverage invariant. The variance calculation below follows the two random choices in an elimination: first the pivot, then the independent samples in its clique. At an earlier prefix with $\mathcal{M}_i \preceq \epsilon_0 P$, the negative drift gives $Z_i \preceq (1+\epsilon_0)P$, as required by the upper-stopped argument. Set $\epsilon_0 = 1/4$ and $\overline{X} = P^{\dagger/2} X P^{\dagger/2}$. Let A_i indicate that $\mathcal{M}_j \preceq \epsilon_0 P$ for every $j < i$. In bag i, condition first on the past and then on the chosen pivot u. Write S for the current rounded Schur graph, S_u for its incident-star Laplacian at u, and $C_u(S)$ for the clique Laplacian obtained by eliminating u from that star. For a centered unrounded roll $X = \mathbf{Y} - \mathbb{E}\mathbf{Y}$, $\|\overline{X}\| \leq 1/\rho$ and

$$B_i := A_i \sum_{\text{rolls}} \mathbb{E}\overline{X}^2 \preceq \frac{A_i}{\rho}\overline{C_u(S)}, \qquad \|B_i\| \leq \frac{5}{4\rho}.$$

Here $S \preceq Z_{i-1}$ and $C_u(S) \preceq S$. With a live vertices the eligible set has at least $a/2$ members, and $\sum_u C_u(S) \preceq \sum_u S_u = 2S$. Thus, averaging only over the uniform eligible pivot, still conditional on the past,

$$\mathbb{E}_u B_i \preceq \frac{5}{\rho a}\mathbf{I}.$$

The bags-of-dice bound [KS16, Theorems 4.3 and 4.5] permits $\sigma_1 = 1/\rho$, $\sigma_2^2 = 5/(4\rho)$, and deterministic $\Omega_i = 5\mathbf{I}/(\rho(v-i+1))$, whose sum has norm at most $5\log v/\rho$. The predictable indicator A_i is fixed throughout the bag, preserving conditional independence. Upper stopping with $\epsilon_0 = 1/4$ and a sufficiently large constant in ρ gives

$$-P/4 \preceq \mathcal{M}_{v-1} \preceq P/4$$

with any fixed polynomial confidence. The concentration applies to the centered unrounded samples.

Rounding drift and zero-pivot failures. There are at most ρq live copies. Low-degree elimination therefore emits at most

$$J \le 4\rho q \sum_{a=2}^{n} a^{-1} = O(q \log^3 n)$$

samples. If $w_{\min}$ is the minimum input weight, a path in P gives $R_P(u,v) \le n/w_{\min}$. Hence

$$\|P^{\dagger/2}\mathcal{D}_{v-1}P^{\dagger/2}\| \le J\delta n/w_{\min}.$$

Choose $\delta \le w_{\min}/(16nJ)$ using the deterministic upper bound for J, with δ a dyadic divisor of the split input mesh. The bounds on q and the input mesh make δ inverse polynomial. The emitted factor Z then satisfies $11P/16 \preceq Z \preceq 5P/4$ on the concentration event. A zero pivot while at least two vertices remain is a reported failure. For its analytical continuation, freeze the previous factor sum F, the current graph S, and both $\mathcal{M}, \mathcal{D}$ at that instant, and append zero bags for the remaining steps. The displayed original-coordinate identity is then unchanged, and future conditional variances are zero. If $\mathbf{k}$ vertices have already been eliminated and $a = v - \mathbf{k} \ge 2$ remain, the isolated live vertex implies $\textsc{Rank}S \le a - 2$, while $\textsc{Rank}F \le \mathbf{k}$. Thus $Z = F + S$ has rank at most $v - 2$ and an extra kernel direction, contradicting the lower spectral event. The frozen zero-bag continuation therefore accounts for early reporting within the same failure bound.

Stable finite factor applications. Ground the last eliminated vertex. The unit triangular factor is $\log n = \mathbf{I} - T$, with $T \ge 0$ strictly lower triangular and column sums at most one. An entry of $\log^{-1} n = \sum_{j=0}^{v-2} T^j$ is the visit probability of a vertex in an acyclic substochastic chain, and is at most one. Thus both triangular solves have only polynomial error amplification. Positive pivots are at least δ; pivots and coefficients are polynomially bounded. Store each coefficient as a word-sized rational. After each rational product or division round back to a fixed common vector grid. Denominators do not multiply along a solve. With s_Z factor entries, a sufficiently fine fixed grid makes the total absolute error as small as prescribed in $O(s_Z)$ operations. Subtracting the output mean supplies the ungrounded action for zero-sum inputs.

Refining the inverse. For t stages define the exact fixed map

$$T_P = \sum_{j=0}^{t-1} (\mathbf{I} - \tfrac{1}{2} Z^\dagger P)^j \tfrac{1}{2} Z^\dagger.$$

Conjugation by $P^{1/2}$ makes it $\mathbf{I} - (\mathbf{I} - \frac{1}{2} P^{1/2} Z^\dagger P^{1/2})^t$. The parenthesized matrix has spectrum in a fixed subinterval of $[0, 1)$, proving symmetry and the claimed inverse comparison with $t = O(1 + \log(1/\epsilon))$. Sparse products and rounded factor applications implement these stages to the prescribed accuracy. On the spectral event, contraction and polynomial conditioning bound all intermediates. The pre-operation guards reject operands outside these fixed ranges on every tape; the chosen reserves ensure they pass on the spectral event.

A cap for random work. Each exact uniform-integer rejection trial succeeds with probability at least $1/2$; combining the range and eligible-vertex tests gives success probability at least $1/4$ for a pivot trial. There are at most $K = O(J + v)$ logical random choices, each using a constant number of such rejection calls. For their conditional trial counts G_j and $\boldsymbol{\lambda} = \log(8/7)$, $\mathbb{E}[e^{\boldsymbol{\lambda} G_j} \mid \text{previous calls}] \le 2$. Padding to a deterministic upper bound $\mathbf{K}_0 = O(J + v)$ and iterating this inequality gives

$$\Pr\left[\sum_j G_j > T\right] \le 2^{K_0} e^{-\boldsymbol{\lambda} T}.$$

The predetermined cap $T = C_c(J + v + \log n)$ therefore fails with probability at most n^{-c} after adjusting its constant. No independence between different rejection counts is assumed. □

Grounded blocks. Adjoin one ground vertex to handle a grounded Laplacian block. For a right-hand side b, lift it to $(b, -\mathbf{1}^T b)$. Applying the exact Laplacian pseudoinverse and subtracting the ground coordinate gives the exact inverse of the grounded block. The same lifting and subtraction transfer the preceding inverse approximations and finite evaluation guarantees. Independent connected components are processed separately.

6.2 Sparse proxies from current cores

Current core certificates supply the leverage bounds for proxies and the physical right-hand-side sampler, without a prior inverse.

Lemma 6.2. *Suppose a maintained owner partitions its original labels into explicit labels and cores of current fixed-reference conductance at least $\phi_C/4$, with current degrees at most the birth reference r_C. Suppose all target coefficients in one bucket are within fixed factors of its representative. There are dyadic upper leverage scores*

$$U_e = \min\{1, C\phi_C^{-2}(r_C(u)^{-1} + r_C(v)^{-1})\}$$

on a core label $e = uv$, with each summand rounded upward within a factor two; use $U_e = 1$ on explicit labels. They dominate the leverage in the full target graph and satisfy

$$\sum_e U_e \le [n(R + \log n) + \tau_0(R + \log n)]n^{o(1)} = n^{19/16 + o(1)}.$$

They can be sampled through the owner's incidence lists with this same bound on total raw envelope mass. No dense leverage table or previous inverse is required.

Proof. Fixed-reference Cheeger gives, on a core of bucket coefficient $\mathbf{g}$,

$$K \succeq c_\star \mathbf{g} L_C \succeq c' \mathbf{g} \phi_C^2 (D_{r_C} - r_C r_C^T / r_C(C)).$$

The inverse quadratic form of the last matrix at $e_v - e_u$ is $r_C(u)^{-1} + r_C(v)^{-1}$. Energy monotonicity for this zero-sum load gives the claimed full-graph resistance upper bound. Multiplication by the actual target edge coefficient changes only fixed constants. Summing endpoint terms uses $\deg_C^{\text{live}}(u)/r_C(u) \le 1$. The owner's vertex and explicit populations give the stated mass. Endpoint list length times its constant endpoint rate gives the same raw, rather than only retained, bound. All rounded scores have one common dyadic grid. □

Finite categorical proxy law. For target graph $K = \mathbf{B}^T \textsc{Diag}(c_\star)\mathbf{B}$, set $\rho_e = C\epsilon^{-2}\log nU_e$ and fix a public polynomial integer B_{pkt} larger than every permitted total proposal mass. At each of its independent positions retain e with probability ρ_e/B_{pkt}, or return zero. A retained packet is $c_e B_e^T B_e/\rho_e$. Its expected sum is exactly K, and its K-normalized norm is at most $\epsilon^2/(C\log n)$. Matrix Chernoff gives a $(1 \pm \epsilon/2)$ proxy with the desired conditional polynomial confidence. Its expected copy count is $O(\epsilon^{-2}\log n \sum U_e)$ and its corresponding count tail permits the public sparse-work cap.

An actual private envelope of total V implements this law by activating a position with probability V/B_{pkt}, drawing from its alias, and thinning to the target rate. The resulting edge/null probability is independent of the valid envelope. Theorem 7.1 enumerates the active positions by exact skipping. The trial count never changes with private V. Round each retained positive weight to a common grid with relative error at most $\epsilon/10$; summing positive semidefinite errors reserves the remaining spectral tolerance. Every target coefficient may be a fixed-depth word-sized rational before that packet rounding.

6.3 Sampling and solving the physical right-hand side

The forcing vector may have a dense edge representation. We sample it before applying the proxy inverse and bound the error in the true physical metric.

Lemma 6.3. *Use the state, classes and physical metric* $\mathbf{L}$ *of the local step. Assume the old residual obeys* $\|\rho\|_{\mathbf{L}^\dagger} \le \beta_s\sqrt{\mu}$. *Fix any positive constant* $c_{\mathrm{sol}} \le 10^{-2}$ *before the construction; its work constants may depend on* c_{sol}. *Using the physical-base owner of coefficient* $\mathbf{z}^2/w$, *one can construct a common-grid direction* $\mathbf{x}^{\mathrm{dir}}$ *with*

$$\|\rho + \mathbf{B}^T \textsc{Diag}(\mathbf{Yz}/\mathbf{k})\mathbf{g} - \mathbf{L}\mathbf{x}^{\mathrm{dir}}\|_{\mathbf{L}^\dagger} \le c_{\mathrm{sol}}\beta_s\sqrt{\mu}$$

with any fixed polynomial confidence, in $n^{19/16+o(1)}$ *sparse preparation work. There is no dense residual multiplication in the algorithm.*

Proof. The forcing vector and its metric. Write $b = \rho + \mathbf{B}^T \textsc{Diag}(\mathbf{Yz}/\mathbf{k})\mathbf{g}$. The class-constant envelope

$$H_j = 1.01(|\theta - 1| + |\boldsymbol{\varphi}_j|/3) + \varepsilon|F_j|$$

dominates $|\mathbf{g}_e|$ for $e \in j$. Writing $H_e = H_j$ on each class, the direction/source radii give $\|\mathbf{H}\|_{\mathbf{w}} \le 4\beta_f$. Incidence contraction gives

$$\|\mathbf{B}^T \textsc{Diag}(\mathbf{Yz}/\mathbf{k})\mathbf{g}\|_{\mathbf{L}^\dagger} \le \sqrt{\mu}\,\|\sqrt{\mathbf{wY}/\mathbf{k}}\,\mathbf{g}\|_2 \le C\beta_f\sqrt{\mu}.$$

The actual metric coefficient is $(\mathbf{z}^2/w)\mathbf{Y}/(\mu\mathbf{k})$, with $\mathbf{Y}/\mathbf{k}$ bounded above and below by fixed constants. The physical owner therefore supplies the preceding upper scores for this metric, even though $\mathbf{k}$ and μ are not structural bucket keys. Construct a $(1 \pm c_\star\gamma)$ proxy P for $\mathbf{L}$ by the public categorical proxy law, including its positive weight rounding reserve.

Sampling the right-hand side. Put $\boldsymbol{\varrho}_e^{\mathrm{rhs}} = w_e H_j^2/(16\beta_f^2)$, so that $\sum \boldsymbol{\varrho}_e^{\mathrm{rhs}} \le 1$, and use the rates $\rho_e^{\mathrm{rhs}} = Q(U_e + \boldsymbol{\varrho}_e^{\mathrm{rhs}})$ with $Q = C\gamma^{-2}\log n$. Use these rational rates, or their upward constant-factor encodings, in both the sampling probabilities and denominators. At each of a public polynomial number of positions return label e with probability $\rho_e^{\mathrm{rhs}}/B_{\mathrm{rhs}}$ and vector packet

$$(Y_e z_e/\mathbf{k}_e)\mathbf{g}_e B_e^T/\rho_e^{\mathrm{rhs}}.$$

Add these packets to ρ. The analytical sum has expectation b. If $\ell_e \le U_e$ is the actual physical leverage, the undivided packet norm is

$$|(Y_e z_e/\mathbf{k}_e)\mathbf{g}_e| \|B_e^T\|_{\mathbf{L}^\dagger} = \sqrt{\mu \mathbf{w}_e(\mathbf{Y}_e/\mathbf{k}_e)\ell_e}\,|\mathbf{g}_e|.$$

The inequality $\sqrt{aU}/(a+U) \le 1/2$ bounds a divided packet by $C\beta_f\sqrt{\mu}/Q$. Its total variance trace is at most $C\mu\beta_f^2/Q$. Centering the independent positions only decreases covariance and changes the maximum packet bound by at most two. The Hilbert-space Bernstein estimate therefore gives error at most $(c_{\mathrm{sol}}/16)\beta_s\sqrt{\mu}$ after increasing the fixed constant in Q as a function of c_{sol}.

Conserved packet rounding. Enclose the scalar packet, round to a common grid, and add opposite endpoint integers. The stored sum is zero-sum and its deterministic error is at most

$$B_{\mathrm{rhs}}\epsilon_{\mathrm{pkt}} \max_e \|B_e^T\|_{\mathbf{L}^\dagger}.$$

Polynomial ranges and the public count make this less than $(c_{\mathrm{sol}}/16)\beta_s\sqrt{\mu}$ with fixed precision on every outcome. Write $\widehat{b}$ for the stored vector.

Solving and grounding the returned direction. Refine the fixed inverse of P to relative error $\eta = c_\star\gamma$. If δ is the proxy error, then on the quotient

$$\|\mathbf{I} - \mathbf{L}^{1/2}T_P\mathbf{L}^{1/2}\| \le \eta + 2\delta.$$

Choose their constants so that this norm times $\|\widehat{b}\|_{\mathbf{L}^\dagger} \le C_b\beta_f\sqrt{\mu}$ is at most $(c_{\mathrm{sol}}/4)\beta_s\sqrt{\mu}$. Evaluate and round $T_P\widehat{b}$ with energy error at most $(c_{\mathrm{sol}}/16)\beta_s\sqrt{\mu}$. Triangle inequality proves the claimed true-metric residual bound. Round onto the fixed vertex-potential grid within the reserve, then subtract the ground coordinate. This exact $O(n)$ operation preserves edge differences and makes $\mathbf{y}' = \mathbf{y} - \mathbf{x}^{\mathrm{dir}}$ exact on the same grounded grid.

Implementing the class stream. Refine classes into narrow w bins and use the class sampler, with an upward common-dyadic enclosure $E_j \asymp H_j^2 w_j^+$ that preserves zero, and member thinning probability $H_j^2\mathbf{w}_e/E_j$. Header preparation is sparse work; swap deletion and cache updates cost $\ddot{O}(1)$ per touched label. Core endpoint lists supply the upper-score stream. Total rate mass is $n^{19/16+o(1)}$ and all inverses are sparse. Fix the direction before the physical mask, charging preparation failure separately. □

6.4 Checking the candidate and preparing later inverses

Lemma 6.4. *The residual and source-normalizer tests in the local step can be implemented in the same sparse preparation row, without a dense candidate scan. Their passing thresholds certify the required returned-state residual bound. On a valid old state their abort probabilities are inverse polynomial, after the local constants are chosen with the test reserves.*

Proof. Exact candidate residual. Maintain the vertex vector $\rho = \mathbf{B}^T\mathbf{f} - b_{\mathrm{dem}}$ exactly on the common primal grid. Initially it is zero. Each selected stored change $\mathbf{f}'_e - \mathbf{f}_e$ updates only its two endpoints in a candidate copy of this vector. The copy has n entries, and its incidence updates take constant work per selected label. Install the copy only on commit. Thus ρ' is the stored candidate's exact residual, including root rounding.

A certified residual test. Retain the old physical proxy and its fixed inverse from the direction computation. Let its analytical map T_{phys} satisfy $\frac{1}{2}\mathbf{L}^{\dagger} \preceq T_{\mathrm{phys}} \preceq 2\mathbf{L}^{\dagger}$. The scalar-window comparisons give a fixed $C_{\mathrm{cmp}} \geq 1$ such that $(\mathbf{L}')^{\dagger} \preceq C_{\mathrm{cmp}}\mathbf{L}^{\dagger}$ on every candidate for which the scalar premises hold. Fix $\boldsymbol{\tau}^2 = \beta_s^2\mu/(16C_{\mathrm{cmp}})$. Compute $q = \rho'^{T}T_{\mathrm{phys}}\rho'$ to absolute error at most $\boldsymbol{\tau}^2/100$, and accept this test only when the computed value is at most $\boldsymbol{\tau}^2/4$. Passing implies $\|\rho'\|^2_{\mathbf{L}^{\dagger}} < \boldsymbol{\tau}^2$ and hence $\|\rho'\|_{(\mathbf{L}')^{\dagger}} < \beta_s\sqrt{\mu'}$, since $\mu' \geq \mu/2$. Conversely an old-metric residual at most $\boldsymbol{\tau}/4$ passes. Lemma 2.4 supplies this stronger bound at inverse-polynomial failure probability by choosing $a_{\mathrm{jump}}, c_{\mathrm{sol}}$ smaller and K larger. No mask is conditioned on this test passing; a failure aborts the attempt.

One fixed-inverse application and dot product have the stated sparse cost. Polynomial ranges, including the lower bound on μ, permit the error reserve with fixed word width; the threshold $\boldsymbol{\tau}^2$ is rational.

The source-normalizer test. On the selected list, compute the common-grid upward sum of actual $\boldsymbol{\lambda}_e|\boldsymbol{\sigma}_e|$. Its zero branch is exact. A test of $\Gamma^2 \leq C^2_{\Gamma,\mathrm{cap}}\beta_f^2 r$, with a fixed sufficiently large threshold constant, is exact rational word arithmetic and certifies $\Gamma \leq C_{\Gamma,\mathrm{cap}}\beta_f\sqrt{r}$. The independent normalizer-safety bound includes the upward enclosure factor and makes its abort probability inverse polynomial. Both tests finish while the old owners and inverse are still installed. □

Source inverse on the old graph. Let $\mathbf{H}$ be the old pre-mask ordinary graph. Prepare or retain an ordinary proxy P with $3\mathbf{H}/4 \preceq P \preceq 5\mathbf{H}/4$ and $3P^{\dagger}/4 \preceq T_P \preceq P^{\dagger}$, then use $T_H = T_P/2$. This scaling gives $0.3\mathbf{H}^{\dagger} \preceq T_H \preceq (2/3)\mathbf{H}^{\dagger}$, including the upper bound required by the raw source moment proof. Fix this inverse before the physical mask and retain it until all change-source replicas and old-view clock increments finish. The candidate coefficients define b while these operations still use the old $\mathbf{H}$.

Fine fields after the commit. Only after that old-graph work is complete, install the candidate metric and prepare its fresh fine fields. Update actual-$\boldsymbol{\lambda}$ dyadic bins on the already enumerated physical commit, before any new fine read. Their total upper mass is $O(r)$. Sampling at rates $C\log n\epsilon_f^{-2}\boldsymbol{\lambda}_e$ gives a fine ordinary proxy, since $\boldsymbol{\tau}_e \leq C\boldsymbol{\lambda}_e$. With ϵ_f a fixed fraction of $\delta_{\min}$ and $\delta_{\min} \geq c_\star n^{-2/p}$, it has $n^{1+4/p+o(1)}$ copies. All $\Theta(\log n\delta_{\min}^{-2})$ fine lanes and inverse actions cost $n^{1+8/p+o(1)}$ per frame. These bins do not depend on joint classes that are awaiting fine observations.

7 Exact independent Bernoulli streams on a word RAM

Geometric gaps enumerate successes without testing every list position. We use a published exact sampler inside attempts. The one-time endpoint routines use a finite implementation coupled to the same law, with a separately budgeted failure probability.

Write $N = 2 + n$. Fix a constant A_{st} before the word length. A call supplies integers $0 \leq N_1 \leq N_2 \leq N^{A_{\mathrm{st}}}$, $N_2 \geq 1$, the rational $p = N_1/N_2$, and a remaining length $0 \leq d_{\mathrm{st}} \leq N^{A_{\mathrm{st}}}$. The declared reference horizon has at most $N^{A_{\mathrm{st}}}$ calls. For $\bar{p} = 1 - p$, the truncated failure gap has law

$$\Pr(G = j) = p\bar{p}^j \quad (0 \leq j < d_{\mathrm{st}}), \qquad \Pr(G = d_{\mathrm{st}}) = \bar{p}^d.$$

Handle $d_{\mathrm{st}} = 0$ and $p = 0, 1$ directly. All other calls use fresh tapes.

7.1 Exact gaps and adaptive streams

Bringmann–Friedrich [BF13, Theorem 5] sample this truncated law exactly in expected time

$$O(1 + \log(\min\{1/p, d_{\mathrm{st}}\})/w)$$

on a word RAM with $w = \Omega(\log\log(1/p))$. For a rational p whose numerator and denominator occupy a constant number of words, the rational implementation has polynomial worst-case space in d_{st} and $1/p$.

Its input convention also supplies magnitude exponents for p, d_{st} and approximations to the normalized mantissa of p. Binary search among bit positions finds both exponents in $\ddot{O}(1)$ work; rational division supplies every requested precision, with multiword operations charged by their actual cost. Thus a sufficiently large fixed $w = \Theta(\log n)$ gives polynomial address space and conditional expected gap cost $\ddot{O}(1)$, including setup. Finite expected cost also gives almost sure termination.

Theorem 7.1. *Under these input bounds, enumerating a length-d_{st} independent Bernoulli(p) stream costs $\ddot{O}((\mathbb{E}Z + 1))$ in expectation, where Z is its number of successes. More generally, an adaptive computation making C_{gap} calls costs $\ddot{O}(\mathbb{E}C_{\mathrm{gap}})$ in expectation. Each call has the stated conditional kernel given the history fixing its inputs. For fixed list contents and rates, the returned indicators are independent Bernoulli variables; disjoint list tapes give independent vectors. Rates chosen from earlier outputs need not give unconditionally independent outputs.*

Proof. A stream needs at most $Z + 1$ gap calls, including its exhausted suffix. The tower rule with predictable call-start indicators sums the uniform conditional cost bound. We retain the looser $\ddot{O}(1)$ price used by the scalar ledger.

For a fixed indicator vector with z successes, let $g_1, \dots, g_z$ be the failure gaps preceding them and g_{z+1} the final failure suffix. Since $\sum_i g_i = d_{\mathrm{st}} - z$, fresh gap kernels give probability

$$\left(\prod_{i=1}^{z} \bar{p}^{g_i} p\right) \bar{p}^{g_{z+1}} = p^z q^{d-z}.$$

This proves the product law, including $p = 0, 1$ and empty suffixes. For adaptive lists, condition at each start before its fresh tape. □

Finite endpoint gaps. The cold and terminal routines need an all-tape work cap, without repeating their dense scans. Use the following separate implementation there. Read $b_0 = B \log n$ bits specifying a uniform dyadic cell of width 2^{-b_0}, and couple it to a continuous uniform U in that cell. The ideal gap is determined by the ordered thresholds $\bar{p}^j$, $1 \le j \le d_{\mathrm{st}}$. Binary-search these thresholds using outward-rounded repeated squaring on a dyadic mesh ζ. Enclose $\bar{p}$ on this mesh first, and clamp all power intervals to $[0, 1]$. Each enclosure has width $O(d_{\mathrm{st}}\zeta)$: multiplication on $[0, 1]$ adds operand errors, and the repeated-squaring expression for $\bar{p}^j$ has at most $O(j)$ leaf and multiplication errors when unrolled. The actual evaluation uses only $O(\log n)$ multiplications.

If the uniform cell overlaps a queried power interval, report failure; otherwise the comparison is exact. Every completed call returns the ideal gap on this coupling, and its work is deterministically $O(\log^2 n)$. All possible ambiguous cells lie within $2^{-b_0} + O(d_{\mathrm{st}}\zeta)$ of the at most d_{st} thresholds. Thus

$$\Pr(\text{failure} \mid \text{call inputs}) \le C\big(d_{\mathrm{st}} 2^{-b_0} + d_{\mathrm{st}}^2 \zeta\big).$$

Choose B and $\zeta = 2^{-D \log n}$ after A_{st} and the desired confidence. A union bound over the declared horizon makes the total coupling failure arbitrarily small inverse polynomial. No failed gap is redrawn. These $O(\log^2 n)$ costs occur only in sparse endpoint preparation; restartable dense attempts use the exact sampler above.

7.2 Thinning incidence proposals to original labels

Thin each original label once, at its first incidence proposal.

Envelope domination and the first proposal. Suppose one target component has a fixed number k of rational rate contributions $r_{ej} \geq 0$ and desired probability $p_e = \min\{1, \sum_{j=1}^{k} r_{ej}\}$. Suppose each rate has at most N_2 nonnegative incidence envelopes, where N_2 is fixed, and $r_{ej} \leq \sum_{x=1}^{b} a_{ej,x}$. Zero envelopes may pad a shorter list. The value $N_2 = 3$ permits the two endpoint occurrences and a direct class occurrence without identifying them. An incidence occurrence is put in a homogeneous random-access list with a constant dyadic proposal probability in

$$[\min\{1, bka_{ej,x}\},\ \min\{1, 2bka_{ej,x}\}].$$

Freeze the lists, incidence probabilities, and targets before their independent proposal tapes. At label e's first proposal, compute its incidence-union probability q_e, use an independent exact p_e/q_e coin, and stamp it processed even on rejection. At least one envelope is at least $(\sum_j r_{ej})/(bk)$, so $q_e \geq p_e$, including the saturated case. Thus the retained label has probability exactly p_e. The product-vector identity makes incidence indicators independent. Different labels use disjoint groups of them and independent retention tapes, so their retained indicators are independent. With all component inputs frozen, apply this construction separately to each prescribed component, with fresh stamps and tapes. Their union is still independent between labels and has probability $1 - \prod_j (1 - p_{e,j})$.

Exact coins and raw work. For probability N_1/N_2, draw a uniform integer in $\{0, \ldots, 2^{\lceil \log_2 N_2 \rceil} - 1\}$, reject draws at least N_2, and accept the coin if the retained draw is less than N_1. Its conditional expected trial count is less than two. All denominators here contain a fixed number of word-sized factors.

Charge every list start and raw proposal, including rejections, at its scalar and label-evaluation prices. List maintenance and weighted envelope bounds are supplied separately. Indexed stamps cost $O(1)$ and reset only at a paid rebuild or the polynomial horizon.

Use with an attempt work cap. Couple an attempt to the uncapped same-tape computation until its work cap. The exact sampler has polynomial memory; only the capped implementation stores an actual-work counter. Apply Markov's inequality to the unconditional stopped expected cost. A cap restarts the whole attempt from its checkpoint, never an individual slow gap. The law is not conditioned on eventual attempt success.

8 Scalar classes and central inspections

Cheap central inspections maintain representatives without fine leverage reads. Signed deadlines locate inspections; a two-window scheduler charges them to motion, paid visits, or handle moves. Joint classes supply feedback and potential certificates. Clocks and deadlines use a fixed common integer grid with guarded polynomial ranges. Fix their key universe before the arithmetic word, whose enlargement must leave the keys unchanged.

8.1 Signed deadlines and lazy initialization

Lemma 8.1. *For keys in $\{0, \ldots, 2^b - 1\}$ with $b = O(\log n)$, a word RAM with polynomial address space supports insertion, deletion, minimum, and maximum in $O(\log b)$ worst-case time per changed*

key, except that minimum and maximum are constant time. A new empty instance costs $O(1)$ after one shared $O(\log n)$ size-table preparation. Its actual touched memory is charged to its actual operations; no universe-sized initialization is performed.

Proof. For $b \le 1$, use a constant-size record. Otherwise use a van Emde Boas recursion, separating a key into a high $b/2$-bit cluster number and a low $\lfloor b/2 \rfloor$-bit key. A node stores its minimum, its cached maximum, a summary of nonempty clusters, and recursively represented children. Only the minimum is excluded from the child clusters; the maximum remains there unless the set is a singleton. Insert into an empty child in constant time and recurse only on the summary; otherwise recurse only on that child. The analogous deletion rule makes its sole nonconstant recursive call on the child if it stays nonempty, and on the summary if that child becomes empty. Deleting the minimum obtains the next minimum from the summary minimum and that child's minimum, then follows this same rule. Thus each operation satisfies $T(b) \le T(b/2) + O(1)$, including cases with an empty or singleton set. Store duplicate-key labels in a doubly linked bucket; only the first insertion or last deletion changes the integer key set.

For explicit addressing, precompute the reserved recursive block sizes $S(b)$ by

$$S(b) = O(1) + S(b/2) + 2^{b/2} S(\lfloor b/2 \rfloor).$$

The child address is its parent's base plus the summary block and its high-key index times the fixed child-block size. The recurrence gives polynomial address space $S(b) = 2^{O(b)}$; one shared table for $b = 0, \dots, b_{\max}$ costs $O(\log n)$.

To avoid clearing reserved blocks, use an initializable-array dictionary for node cells. An array $\mathrm{pos}[i]$ may initially contain arbitrary machine words. A stack $\mathrm{key}[0], \dots, \mathrm{key}[k-1]$ records initialized cells. Cell i is initialized exactly when $0 \le \mathrm{pos}[i] < k$ and $\mathrm{key}[\mathrm{pos}[i]] = i$; only after the range test is that stack entry read. Otherwise its value is the empty default. Initializing a cell appends i and fills its constant-size record. Each tree gets a fresh disjoint reserved block, abandoned on discard; old descendants are never reused by resetting only the root. The polynomial instance count bounds all addresses. Fix that count and the key width before choosing the arithmetic word. □

A signed clock C accumulates field increments. For each attached label, B_e^{clk} is its baseline at the last inspection and R_e its allowed radius; the signed debt is $C - B_e^{\mathrm{clk}}$. Keep two deadline indices, storing upper deadlines $B_e^{\mathrm{clk}} + R_e$ and lower deadlines $B_e^{\mathrm{clk}} - R_e$. When the clock is changed, repeatedly report upper deadlines at most C and lower deadlines at least C. The first report queues the original label and removes all its fixed-number deadline handles, suppressing duplicate service in that frame. There is $\ddot{O}(1)$ work per insertion, deletion, or report and $O(1)$ work for a clock change that reports nothing. Shift signed integers by a fixed offset into the key universe.

When a handle moves to a new clock, preserve its signed debt by setting $B_e^{\mathrm{clk,new}} = C^{\mathrm{new}} - (C^{\mathrm{old}} - B_e^{\mathrm{clk,old}})$. Preserve an already due flag. Its radius stays fixed until an actual inspection. If a change of occurrence representation does not preserve individual debts (for example two core endpoint occurrences becoming one explicit occurrence), queue a cheap exact central inspection at that paid handle move. That inspection has the same scalar price as the move and uses no fine-score read.

8.2 The two-window central scheduler

For a central representative $b \in (-1, 1)$ choose a downward dyadic radius within a factor two of

$$\rho(\overline{\varsigma}) = c_{\mathrm{cen}} P^{-8} \min\{1 - |\overline{\varsigma}|, P^{-1}\}.$$

Maintain the outer invariant $|e-b| \le 2\rho(\overline{\boldsymbol{\varsigma}})$ between inspections. At an inspection, enclose the current e in a certified interval of radius at most $\rho_{\min}/1000$. Retain b only if this interval lies in $[b-\rho(\overline{\boldsymbol{\varsigma}}), b+\rho(\overline{\boldsymbol{\varsigma}})]$; otherwise choose a new dyadic center b' for which it lies in $[b'-\rho(b')/4, b'+\rho(b')/4]$. Passing inspections reset clock baselines but retain the center and tracking-score epoch.

Such centers can be located with exponent comparisons and quotient. For $d_Z = 1-|\overline{\boldsymbol{\varsigma}}| \le P^{-1}$ use overlapping dyadic margin shells and spacing a fixed small fraction of $c_{\mathrm{cen}}P^{-8}2^{-j}$ in the shell $d \asymp 2^{-j}$. For $d > P^{-1}$ use spacing a small fraction of $c_{\mathrm{cen}}P^{-9}$. There are $\ddot{O}(\log n)$ centers on a polynomial margin range. The overlaps and spacing smaller than $\rho/100$ ensure the inner-quarter placement, including a shell-boundary point.

Lemma 8.2. *Suppose selected labels receive a final central inspection. Suppose the physical owner supplies a fixed-number occurrence decomposition for the clock field satisfying*

$$\sum_{e,o} \phi_{eo}^2 F_{eo}^2 \le C\beta_f^2, \qquad F_{e,v} = \mathbf{z}_e |h_v - g_C|/\sqrt{\mathbf{w}_e}, \qquad h = v_y/(\varepsilon\mu\theta),$$

using one common core gauge at both endpoints. An explicit route uses its actual edge difference and $\phi = 1$; its per-field header evaluation is paid in the sparse row. Assume $\sum \phi^{-4} = \ddot{O}(m)$ and $\Psi_Z \le CP^2 m$. Then the scheduler preserves its outer invariant, and through a clean attempt its autonomous central inspection work is

$$O\left(\ddot{O}(1)[m + \sum_t B_{*,t} + N_{\mathrm{move}}] + \ddot{O}(1)N_{\mathrm{paid}}\right).$$

Here N_{move} counts paid occurrence-representation moves that require an inspection, and N_{paid} counts selected or already paid fine visits. Under the owner and observer attempt budgets this is $\ddot{O}(m)$ plus the supplied sparse work. Changing the central class as part of such an inspection does not itself create a new occurrence-representation move: the same core-vertex clocks and gauge are retained, and only list membership changes. Thus N_{move} does not recursively count the inspections whose number this bound is proving.

Proof. Clocks and finite representation. An unselected label has fixed $\boldsymbol{\varrho}_e = \mathbf{z}_e/\mathbf{w}_e$ and its exact central increment is

$$\Delta\boldsymbol{\varsigma}_e = \boldsymbol{\varrho}_e B_e h + \frac{(1-\theta)C_e}{\varepsilon\theta}.$$

Maintain signed clocks for its endpoint occurrences and a global positive clock with increment $(1+\varepsilon)(1-\theta)/(\varepsilon\theta)$. At an inspection save their baselines. Use endpoint radii within fixed factors of $\rho(\overline{\boldsymbol{\varsigma}})/\boldsymbol{\varrho}_e^+$ and a common radius within fixed factors of $\rho(\overline{\boldsymbol{\varsigma}})$, where $\boldsymbol{\varrho}_e \le \boldsymbol{\varrho}_e^+ \le 2\boldsymbol{\varrho}_e$. Choose the fixed factors so that the sum of all pre-deadline possible movements is at most $\rho(\overline{\boldsymbol{\varsigma}})/4$; for two endpoints, radii in $[\rho/(24\boldsymbol{\varrho}^+), \rho/(12\boldsymbol{\varrho}^+)]$ and common radius in $[\rho/24, \rho/12]$ suffice. Their positive lower bounds permit the work charge; owner moves translate signed debts exactly.

Round each frame's $h_v = v_{y,v}/(\varepsilon\mu')$ onto the fixed dyadic clock mesh Δ_C. Use one common-grid gauge at both endpoints of a core occurrence; explicit routes use the rounded endpoint difference. Clock increments, baselines and radii are integers in these units. For public horizon $T_{\max}$ and $\boldsymbol{\varrho}_e \le a_{\max}$, the unselected-coordinate discrepancy over an inspection interval is at most

$$2T_{\max}a_{\max}\Delta_C.$$

Choose the mesh as in Section 12 so this is below $\rho_{\min}/1000$ and total normalized rounding variation over all labels and frames is at most one. Include the perturbation of the energy bound and upward rounding of the common positive increment in that reserve. Baselines and due flags then move by exact integer arithmetic.

Validity through the end of a frame. In the following overshoot estimates, h is the scalar step size of Section 2; the vertex clock field above enters through its accumulated increments. At an inspection the true point is within $\rho(\overline{\boldsymbol{\varsigma}})$ of the retained or new center. Before the next deadline, the clock triangle inequality gives further motion at most $\rho/4$. At the first crossing, buffer the label until the completed frame. An unselected label has $p_e < 1$, so the wall floor gives $|\boldsymbol{\nu}_e| \le d_{Z,e}/k_P$. Its last unselected increment is $\mathbf{F}_e - \boldsymbol{\nu}_e$, hence has magnitude at most $\gamma h + d_{Z,e}/k_P$, plus the numerical reserve. The ratio

$$\frac{d_Z}{P^{12}\rho(\overline{\boldsymbol{\varsigma}})} \le C \max\{P^{-4}, P^{-3}\}$$

is uniformly small after the multiplier in k_P is chosen. Also $d_Z \ge c_\star n^{-2/p}$ from the central cap and $P \ge p$; therefore $h/\rho \le C c_h n^{-1/p} P^{-7}$. Choose c_h small and the numerical grid fine so the total last-frame overshoot is less than $3\rho/100$. The old error, unreported movement, and last-frame overshoot sum to less than 2ρ. Selected labels are forced into the final queue, and no next physical query uses their old representative before the queue is drained.

For the finite step use $\boldsymbol{\nu}_c$ and $\mathbf{F}_c = F + O(\Delta_D)$; $\Delta_D \ll \rho_{\min}$ preserves the bound without dividing by a sampling probability.

Charging inspections to motion. Use proof-only credit equal to absolute variation divided by each account's fixed radius, truncated at one. A deadline report consumes a fixed positive amount; a passing inspection creates none. Transport preserves credit and debt, while selected/fine visits and paid representation moves pay their own resets. Thus autonomous reports are bounded pathwise by normalized variation, initialization, and paid resets, including credit accumulated in earlier frames.

For endpoint clocks, Cauchy–Schwarz in the owner energy bound gives

$$\sum_{e,o} \frac{\boldsymbol{\varrho}_e^{+} |\Delta C_{eo}|}{\rho(\overline{\boldsymbol{\varsigma}}_e)} \le C\beta_f \left(\sum_{e,o} \frac{\phi_{eo}^{-2}}{\mathbf{w}_e \rho(\overline{\boldsymbol{\varsigma}}_e)^2} \right)^{1/2} \le \ddot{O}(B_*).$$

Indeed $\rho^{-2} \le C P^{18} d_Z^{-2}$ and $\sum \phi^{-2} d_Z^{-2} \le (\sum \phi^{-4})^{1/2} (\sum d_Z^{-4})^{1/2}$, while $w \ge 3s$ and $\sum d_Z^{-4} = O(m)$ uniformly in P. For the common clock,

$$\sum_{e} \frac{(1-\theta)}{\rho(\overline{\boldsymbol{\varsigma}}_e)} \le C P^9 (\beta_s/\sqrt{r}) \sum_{e} d_{Z,e}^{-1} \le C P^9 B_*.$$

Multiplying the lifetime reports by the $\ddot{O}(1)$ deadline and certified-point price proves the result. Central evaluation uses only stored $\mathbf{f}, \boldsymbol{\lambda}, c_\star, \mu$ and two vertex potentials. □

Transaction order. Use Steps 6–8: finish old-view increments before the physical/owner commit, preserve queued flags through moves, and inspect the final queue before installing final classes. Clock moves preserve their debts or request the paid cheap inspection above. No new physical query uses central or tracking metadata whose final service is pending.

8.3 Joint class implementation and certificates

Membership and class totals. The joint metadata key is the tuple of a certified tracking cell, central center, narrow actual-$\boldsymbol{\lambda}$ bin, and guard flag. Keep the tracking cell until an already paid fine inspection supplies its replacement. A central-only visit changes only the central key. Actual-$\boldsymbol{\lambda}$ bins and guard flags change only on selected physical labels. Recomputing waterfill branches changes headers, not member lists.

The tracking centers can be chosen in overlapping dyadic margin shells with spacing $c_\star\delta/p$, giving $O(p\log n)$ centers. A supplied fine observation must guarantee $|\boldsymbol{\varpi}_e/\kappa - x_{j(e)}| \le \boldsymbol{\delta}_e/(256p)$ until its next inspection. Narrow multiplicative bins for $\boldsymbol{\lambda}$ have $O(\log n)$ boundaries and can be located by binary search in $\ddot{O}(1)$ comparisons, or by certified scalar logarithms. The central construction supplies $\ddot{O}(\log n)$ centers. Thus the total number of possible joint headers is $n^{o(1)}$, independent of m.

Store each class and homogeneous incidence list in a compact array with original-label back-pointers. Append insertion and swap deletion cost $O(1)$ for each of a label's fixed number of roles. Header scans, aliases, and powers belong to preparation. Test the exact count $N_j = |j|$ first: empty headers contribute zero and evaluate no barrier, gradient, curvature, breakpoint, or reciprocal. The potential bound controls only occupied headers.

Sampling from a class. For supplied rational member weights within a fixed factor, propose a uniform class position and retain it with its exact weight divided by the class upper bound. The conditional expected trial count is constant. For weights $H_j/\mathbf{w}_e$, the class header factor is constant and the narrow $\boldsymbol{\lambda}$ bin makes $\mathbf{w}_e$ narrow; use $H_j/w_{j,\min}$ as the envelope. Round the class upper envelopes upward onto one dyadic common grid within a fixed factor (zero stays zero). Their count-weighted masses are then integer multiples of one mesh and can be summed exactly in a fixed number of words; no product of all class denominators is formed. Choose classes from those masses using a header alias table and exact rational coins. Thinning gives the exact requested categorical law, including a null outcome when the prescribed packet law requires one. Use the envelope normalizer directly; no dense exact $\sum 1/\mathbf{w}_e$ is needed.

Computing the tracking feedback. Use $\pi = \boldsymbol{\lambda}(1+\varepsilon\boldsymbol{\varsigma})$, $\mathbf{G} = \nabla_{\boldsymbol{\varpi}}\Psi_R$, and the mixed norm $\mathcal{N}$ of Section 2. The next lemma constructs the required feedback before the physical mask. Its optimization coefficients b_j are local to the proof; they are unrelated to the central representatives denoted by b_j below.

Lemma 8.3. *Suppose there are $J = n^{o(1)}$ joint classes. Within a class the weights π_e differ by a factor at most $1+\eta_c$, and a class-constant signed observation obeys $|\widehat{G}_e - G_e| \le \eta_c(|G_e| + C_{p,\kappa})$. Suppose the class counts and positive masses $W_j = \sum_{e\in j}\pi_e$ are available with arbitrarily small fixed relative error. For sufficiently small fixed η_c, one can construct (1), including its guard, in $\ddot{O}(nJ)$ frame work. This does not require a computed gradient-normalized sampling rate.*

Proof. Reduction to class magnitudes. For the observed objective, replacing a feasible vector by its unweighted class mean preserves the objective and cannot increase its infinity norm. Jensen's inequality and the within-class weight ratio increase its weighted norm by at most $\sqrt{1+\eta_c}$. A small rescaling thus makes class-constant optimization arbitrarily close to optimization over the full mixed ball. Replacing G by $\widehat{G}$ changes any mixed-ball objective by at most $\eta_c h(\|G\|_{\mathcal{N}^*} + B_R)$, by coordinatewise sign invariance of the mixed norm and duality.

A one-parameter optimization. Let $N_j = |j|$ be the number of coordinates in class j. For a trial maximum coordinate $v \in [0,h]$, write $b_j = N_j|\widehat{G}_j|$ and solve

$$\max \sum_j b_j x_j : \quad 0 \le x_j \le v, \quad \sum_j W_j x_j^2 \le ((h-v)/M)^2.$$

The solution is $x_j = \min\{v, \alpha b_j/W_j\}$. Sort the J breakpoints in α and solve the resulting quadratic norm constraint with certified arithmetic. Let v_* be the infinity norm of a nonzero optimum on the

mixed-ball boundary. It is at least $v_{\min} = h/(1 + M\sqrt{\sum_j W_j})$. Test a geometric grid of v between $v_{\min}$ and h with fixed ratio $1 + \eta_c$. If $v \le v_* < (1 + \eta_c)v$, scaling the optimum by v/v_* keeps it feasible for this trial, because $(v/v_*)(h - v_*) \le h - v$. It retains at least a $1/(1 + \eta_c)$ objective fraction. There are $O(\log n)$ trials. Choose the best, assign the opposite observed-gradient sign, and reduce the radius slightly to cover the mass and scalar enclosures. These fixed small losses can leave direct descent greater than $0.98h\|G\|_{\mathcal{N}^*}$.

Use hysteresis guard thresholds $(1+\kappa/10)s$ and $(1+\kappa/5)s$. For $a_0 \le \kappa/20$ these give the required lower-wall buffer. Guarding only deletes downward coordinates and cannot increase the norm. A deleted beneficial downward coordinate has positive true gradient only for $0 < t \le \log(1 + \kappa/5)$, where it is bounded by the fixed baseline. Its total lost objective is at most ChB_R. Update flags only on selected labels, since the other $\boldsymbol{\lambda}_e$ are literal no-ops.

Maintaining the class masses. The positive class mass has the identity

$$W_j = \sum_{e \in j} \boldsymbol{\lambda}_e (1 - \overline{\boldsymbol{\Xi}}_e) + \mu^{-1} \left[\sum_{e \in j} \boldsymbol{\lambda}_e z_e c_e / \mathbf{w}_e - \mathbf{y}^T B^T (\mathbf{1}_j \boldsymbol{\lambda} \mathbf{z} / \mathbf{w}) \right].$$

Cache the scalar and incidence-vector coefficients by rounding each square-root expression to a common integer grid. A selected change or enumerated migration updates only its own scalar entries and two incidence endpoints; the J vertex dot products cost $O(nJ)$. The true mass is at least $(1 - \varepsilon)sN_j$. Choose each term's absolute precision after polynomially bounding the number of terms and the current $\mu^{-1}, \mathbf{y}$, so the entire header error is at most $\eta_c(1 - \varepsilon)s/100$. This gives the required relative enclosure for every nonempty class, even when large terms cancel. Remove exactly the rounded contribution previously inserted; the cache then sums only current terms and accumulates no historical rounding error. Certified classes and enumerated migrations have their separately charged maintenance costs. □

Computing the central source. Here $G_j = |\psi'_P(b_j)|$, $J_j = \psi''_P(b_j)$, and G_j^c, J_j^c are certified upper enclosures. For this waterfill $W_j = \sum_{e \in j} \mathbf{w}_e$ and $\widetilde{w}_j = W_j/N_j$; selected scalar updates and enumerated migrations maintain them. Sort the nonzero-G_j^c header breakpoints:

$$t_j = \frac{\widetilde{w}_j}{G_j^c} \min\{c_Z d_{Z,j}/P^2, \gamma h\}$$

Between breakpoints the squared source norm is a fixed capped sum plus ϑ^2 times a fixed positive uncapped sum. Choose the rational parameter by the rule below and store the literal minimum

$$F_j = -\textsc{Sgn}(b_j) \min\{c_Z d_{Z,j}/P^2, \gamma h, \vartheta G_j^c/\widetilde{w}_j\};$$

headers with $G_j^c = 0$ return $F_j = 0$. Each coordinate is a fixed-depth rational with a fixed word bound; no further coordinatewise rounding is used.

First test the fully capped branch of (3). For fixed $0 < \eta < 1/16$, enclose its squared norm with absolute error at most $\eta\beta_s^2/100$. If its upper endpoint is at most β_s^2, return the fully capped source with $\vartheta^{-1} = 0$, also when it lies in the filled range. If this test fails, the fully capped source has true norm greater than $(1 - \eta/4)\beta_s$. Monotonicity permits a rational ϑ with certified norm in $[(1 - \eta)\beta_s, (1 - \eta/2)\beta_s]$. Use sorted breakpoints and interval bisection, with the interval-sum mesh smaller than a fixed fraction of this gap. All nonzero breakpoints, coefficients and the search interval have polynomial upper and lower bounds, so the header work is $n^{o(1)}$. Fixed-grid interval sums avoid products of unrelated denominators. They choose only ϑ; the literal minimum above defines every coordinate.

Derivative enclosures and guards. For each occupied header, the central-power construction constructs fixed-grid rational enclosures $G_j \le G_j^c \le (1+\eta)G_j$ and $J_j \le J_j^c \le (1+\eta)J_j$ in $n^{o(1)}$ work; at $b_j = 0$, set $G_j^c = 0$ exactly. Skip empty headers first. The cited construction uses rounded powers and pre-operation guards on every tape, aborting before overflow; valid inputs pass. Source caps and uncapped membership are exact rational comparisons. The literal rational source therefore meets the hypotheses of Theorem 2.6 and Lemma 2.7. Their waterfill and comparison-source bounds apply to this construction.

Potential certificates. Compute the potential certificates by summing the positive header counts times their center potentials, skipping zero counts before evaluating those potentials. Tracking relative margin error gives a summand ratio less than $e^{1/128}$; central outer error 2ρ gives ratio $\exp(O(P^{-7}))$. Outward fixed-grid power evaluation therefore gives computable U_R, U_Z with

$$\Psi_R \le U_R \le 2\Psi_R, \qquad \Psi_Z \le U_Z \le 2\Psi_Z.$$

Repeated squaring uses $\ddot{O}(1)$ scalar operations for a central power, with all errors enclosed on one fixed grid. These computations belong to the header preparation row. Their validity still depends on the supplied tracking observation invariant; a central-only inspection never silently renews that separate certificate.

9 Observing leverage changes within a clean attempt

The observer discovers stale leverage references without scanning the dense graph. Random source fields capture actual coefficient changes; a filter frozen before the mask separates vertex-clock work from a small residual handled by random marks. A lifetime credit identity pays delayed reads, and packed arithmetic removes the confidence-lane factor from dense work. Theorem 9.1 gives the attempt bound; Lemma 14.1 verifies its inputs.

Put $m = \Theta(n^2)$, $r = \Theta(n)$, $s = r/m$, and $p = 512$. All graph matrices below act on the quotient by component constants; where a full-space norm is needed, it is stated explicitly. Parallel original-label multiplicity is bounded by a fixed constant. The ordinary graph $\mathbf{H}$ uses the actual coefficients $\boldsymbol{\varrho}_e^2$: fixed-depth, word-sized rationals, equal to $\mathbf{z}_e^2/\mathbf{w}_e^2$ in the flow application. Their denominators may differ.

State, direction, and coefficient-change inputs. At each valid frozen pre-mask state, write

$$\mathbf{H} = \mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2)\mathbf{B}, \qquad \mathbf{U} = \textsc{Diag}(\boldsymbol{\varrho})\mathbf{B}\mathbf{H}^{\dagger/2}, \qquad \boldsymbol{\Pi} = \mathbf{U}\mathbf{U}^T, \qquad \boldsymbol{\tau}_e = \boldsymbol{\Pi}_{ee}, \qquad S = \boldsymbol{\Pi} \circ \boldsymbol{\Pi}.$$

Assume the frozen state and direction satisfy

$$\begin{gathered} \boldsymbol{\lambda}_e \asymp s + \boldsymbol{\tau}_e, \quad \boldsymbol{\lambda}_e \ge s, \quad \mathbf{w}_e \asymp \boldsymbol{\lambda}_e, \quad \sum_e \boldsymbol{\lambda}_e = O(r), \\ \sum_e \boldsymbol{\delta}_e^{-p} = O(m), \quad 0 < \overline{\boldsymbol{\delta}}_e \le 1, \quad \overline{\boldsymbol{\delta}}_e \asymp \boldsymbol{\delta}_e, \quad \overline{\boldsymbol{\varpi}}_e \asymp \boldsymbol{\lambda}_e \asymp s + \boldsymbol{\tau}_e, \quad \|d\|_{\mathbf{w}} \le C\beta, \quad d_e \ge 0. \end{gathered} \tag{18}$$

Here $\beta = \beta_f$ is the physical-step scale. The positive reference $\overline{\boldsymbol{\varpi}}_e$ estimates $s + \boldsymbol{\tau}_e$; it is distinct from the logarithmic tracking error $\boldsymbol{\varpi}_e$ of Section 2. In particular, every fixed inverse-margin moment through order p is $O(m)$, and $\boldsymbol{\delta}_e \ge c_\star n^{-2/p}$. The tracking tube supplies both comparisons throughout each reference epoch; both are used to convert additive score error to logarithmic error. The nonnegative vector d controls the mean absolute coefficient changes in (O1) below. Assume

every final stored candidate is frozen in a pending transaction before source construction, with actual coefficient changes satisfying

$$\boldsymbol{\sigma}_e = \frac{\boldsymbol{\varrho}_e'^2}{\boldsymbol{\varrho}_e^2} - 1, \qquad \boldsymbol{\sigma}_e = 0 \text{ on every unselected label}, \qquad \mathbb{E}[|\boldsymbol{\sigma}_e| \mid \mathcal{F}] \le C d_e. \tag{O1}$$

Frozen state and commit order. Here $\mathcal{F}$ includes the filter fixed before the mask. Keep the old $\mathbf{H}$, original-product owner, and inverse until all source and observer batches finish. Then commit, update owners and $\boldsymbol{\lambda}$ bins, and prepare final-graph fine fields, without an old-owner copy. Use the unconditioned physical law: the RAM algorithm has independent label coins and aborts the attempt on a failed residual or normalizer test, without resampling. The observer expectation itself does not require independence between physical labels.

The normalizer is an actual upward common-grid enclosure satisfying

$$C_\Gamma \sum_e \boldsymbol{\lambda}_e |\boldsymbol{\sigma}_e| \le \Gamma \le 2C_\Gamma \sum_e \boldsymbol{\lambda}_e |\boldsymbol{\sigma}_e|, \qquad \Gamma = 0 \iff b = 0. \tag{19}$$

Valid frames also satisfy the separately tested bound $\Gamma \le C_{\mathrm{cap}} \beta \sqrt{r}$. Correctness of the score certificate also assumes a first-failure event for the physical update, outside which

$$\mathbf{H}^{\dagger/2} \mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2 |\boldsymbol{\sigma}|) \mathbf{B} \mathbf{H}^{\dagger/2} \preceq q_{\mathrm{op}} \mathbf{I}, \quad q_{\mathrm{op}} < 1/2,$$

and each old certified margin remains comparable to the final margin through the frame. These correctness inputs do not condition (O1) on future success.

The ordinary graph, base graph, and observer metric. The ordinary graph supplies source products, a maintained base graph supplies proxy and residual proposals, and the observer metric prices fine reads. Their coefficients are defined separately below. The maintained base has coefficients $\mu_e \asymp \boldsymbol{\varrho}_e^2 / \boldsymbol{\lambda}_e$. Scale it once so that the actual metadata coefficient

$$\omega_e \asymp \frac{\boldsymbol{\varrho}_e^2}{\overline{\boldsymbol{\varpi}}_e \overline{\boldsymbol{\delta}}_e^4}, \qquad \mathbf{M} = \mathbf{B}^T \textsc{Diag}(\omega) \mathbf{B}$$

satisfies

$$\mu_e \le \omega_e \le R_0 \mu_e, \qquad R_0 = \ddot{O}(n^{8/p}), \qquad M \preceq C(R_0/s)\mathbf{H}. \tag{20}$$

The ordinary and base owners supply compact incidence arrays and, for every frozen field $\mathbf{z}$, nonnegative endpoint envelopes with

original-label target	total envelope mass
$\boldsymbol{\varrho}_e^2 (A_e z)^2$	$A_L \mathbf{z}^T H z$
$\mu_e (A_e z)^2$	$A_L \mathbf{z}^T G_0 \mathbf{z}$

$$G_0 = \mathbf{B}^T \textsc{Diag}(\mu) \mathbf{B}, \quad A_L = n^{o(1)}. \tag{O2}$$

For the owner of Section 5, (O2) follows from (15) and the worst-case bound on ϕ^{-1} in Lemma 5.1. Envelope construction is included in the sparse-preparation work S_{sparse} below; evaluating one returned label costs $\ddot{O}(1)$. The maintained actual-$\boldsymbol{\lambda}$ bins supply a nonnegative envelope of total mass $O(r)$ for the target $\boldsymbol{\lambda}_e$. Keep the ordinary owner for source products: using the base metric alone could lose a factor $1/s$. Use the fixed symmetric maps and finite energy-error guarantees of Section 6. Retain the source inverse's old-graph proxy $3\mathbf{H}/4 \preceq P_H \preceq 5\mathbf{H}/4$, on the same spectral event, for the executable energy guard.

Numerical conventions. Apply the finite arithmetic convention and the declared inverse-polynomial nonzero lower bounds. Fix state grids and coefficient ranges before source/packet grids, then source output before the fixed-depth filter ladder. On every reference outcome, local errors have total normalized energy and debt charge at most one over the declared horizon. Accumulated vectors, Γ, ω, μ and proxy packets use common grids; ordinary coefficients and change ratios need not. Let $\delta_{\min}, \lambda_{\min} > 0$ be the deterministic lower bounds on permitted margins and $\boldsymbol{\lambda}_e$, fixed by the parameter hierarchy and $\boldsymbol{\lambda}_e \geq s$, before score reads.

Theorem 9.1. *Under the state, update, owner, inverse, and numerical hypotheses above, start a mathematical attempt at a paid clean checkpoint and stop it before beginning any frame with an invalid premise. Let J_t indicate that frame t is begun. Suppose its predictable budgets and selected-read charges obey*

$$\sum_t J_t B_t \leq \ddot{O}(m), \qquad B_t = \beta_t m/\sqrt{r}, \qquad \mathbb{E}\sum_t J_t \sum_{e \text{ selected}} \bar{\delta}_{e,t}^{-3} \leq \ddot{O}(m). \tag{O3}$$

Assume at most $T = n^{3/4+3/p+o(1)}$ frames and aggregate sparse-preparation work $S_{\text{sparse}} = o(m)\ddot{O}(1)$; a fixed polynomial gap below n^2 with arbitrary fixed powers of $\log n$ is sufficient.

There is a finite word-RAM source and observer implementation with expected work

$$\begin{aligned} O\Big(&(\log\log n)^2 \mathbb{E}\sum_t J_t \sum_{e \text{ selected}} \bar{\delta}_{e,t}^{-3} + \ddot{O}(1)\mathbb{E}\sum_t J_t B_t \\ &+ \log\log n\, \mathbb{E} M_{\text{pkt}} + nT\log n + S_{\text{sparse}} + \ddot{O}(m)\Big) \\ &= O\Big(\ddot{O}(m)\Big). \end{aligned} \tag{O4}$$

Here M_{pkt} counts original selected endpoint packets, without a factor for confidence replicas; it is at most twice the selected-label count. The last dense term pays the final clean refresh, including when the endpoint is subsequently rejected. The source/capture and numerical implementation failure probability can be made n^{-c} for any fixed $c_\star$, in addition to the explicitly supplied primitive failure probabilities.

Positive-account read payments use a lifetime ledger. No bound on a frame's payments in terms of that frame's new credit is asserted. The theorem supplies expected work inside one stopped attempt; phase success, optimization progress, the selected-rate bound (O3), and the sparse owner/inverse work remain separate obligations.

9.1 From selected changes to source fields

Construct a stored field Y with two guarantees. Its raw second moment (23) will bound work after averaging over the physical mask. Its error relative to the exact field with the same source, (24), will transfer the source's signal to Y and prove the per-label capture bound (25).

Conserved source coefficients. Set

$$\Delta = \mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2|\boldsymbol{\sigma}|)\mathbf{B}, \qquad \zeta = (n\log n)^{-1}.$$

For label e, the positive certificate $(S|\boldsymbol{\sigma}|)_e = \boldsymbol{\varrho}_e^2 A_e H^\dagger \Delta \mathbf{H}^\dagger \mathbf{B}_e^T$ controls the inverse-induced score change by the inverse-order argument below. We estimate it without scanning all edges. If $b = 0$, return zero fields and no clock increments. Otherwise traverse the selected list and obtain Γ from

the common-grid normalizer construction, using the exact old and final rational coefficients. Zero terms stay zero. For a fixed sufficiently large constant A_{src} and $b \neq 0$, prepare short dyadics

$$(1-u)\boldsymbol{\varrho}_e\sqrt{|\boldsymbol{\sigma}_e|/\Gamma} \le \widehat{\mathbf{c}}_e \le \boldsymbol{\varrho}_e\sqrt{|\boldsymbol{\sigma}_e|/\Gamma}, \qquad u = \log^{-A_{\mathrm{src}}} n \le 1/4, \qquad \mathbf{f} = \mathbf{B}^T \textsc{Diag}(\widehat{c}_\star)\xi.$$

These dyadics define the mathematical source; the covariance sandwich, not an additive ζ-reserve, controls their error. Compare squares with $\boldsymbol{\varrho}_e^2|\boldsymbol{\sigma}_e|/\Gamma$, without storing an exact $\boldsymbol{\varrho}_e$. Binary search over $O(\log n)$ exponents and an $\ddot{O}(1)$-bit mantissa costs $\ddot{O}(1)$ per selected label. Give each nonzero selected coefficient an independent uniform sign ξ_e and use opposite incidence signs at its endpoints. Then

$$(1-u)^2\Delta/\Gamma \preceq \mathbb{E}ff^T \preceq \Delta/\Gamma, \qquad \mathbb{E}\|\mathbf{f}\|_{H^\dagger}^2 \le \frac{\sum_e \boldsymbol{\tau}_e|\boldsymbol{\sigma}_e|}{\Gamma} \le 1, \tag{21}$$

after fixing C_Γ.

Sampling an original-graph product. The source estimates hold for any valid fixed symmetric map

$$\tfrac{3}{10}\mathbf{H}^\dagger \preceq T_H \preceq \tfrac{2}{3}\mathbf{H}^\dagger.$$

It may depend on the actual b, but is fixed before source signs and product tapes; all estimates are uniform in this choice. The joint implementation uses the retained old inverse, already fixed before the mask. For an actual stored iterate x, encode

$$\boldsymbol{\varrho}_e^2(A_e x)^2/\zeta + C\log n\boldsymbol{\lambda}_e \le \rho_e \le 2\big(\boldsymbol{\varrho}_e^2(A_e x)^2/\zeta + C\log n\boldsymbol{\lambda}_e\big).$$

Choose a fixed public count B_{prod} dominating all permitted target and envelope masses, and use that many independent categorical/null positions. A position returns e with probability ρ_e/B_{prod}, with packet

$$\frac{\boldsymbol{\varrho}_e^2(A_e x)}{\rho_e}\mathbf{B}_e^T.$$

Use the envelope thinning and exact skipping of the public-count kernel.

Product guards and the analytical continuation. Before reading the product tape, test $x^T P_{\mathbf{H}}\mathbf{x} \le C_{\mathrm{en}}\log n$, coordinate ranges, and prepared envelope mass at most $C_{\mathrm{env}} n\log^2 nA_L$, including actual-$\boldsymbol{\lambda}$ bins. Use an outward common-grid quadratic-form enclosure of absolute error at most one with a fixed threshold reserve; header masses give the other test without an original-edge scan. A failed guard aborts execution. Its analytical continuation uses the exact product $\mathbf{Hx}$, then rounds $\mathbf{x}_i + T_H(\mathbf{f} - \mathbf{H}\mathbf{x}_i)$ on the same iterate grid with the same componentwise zero-sum closure; put this defect in ν_i. Write $\widehat{\mathbf{Hx}}$ for the unrounded packet sum at the stored input when guards pass, and for $\mathbf{Hx}$ otherwise. For every valid private representation,

$$\mathbb{E}[\widehat{\mathbf{Hx}} \mid x] = \mathbf{Hx}, \qquad \textsc{Cov}(\widehat{\mathbf{Hx}} \mid x) \preceq \zeta\mathbf{H}. \tag{22}$$

Indeed $\sum_e \boldsymbol{\varrho}_e^4(A_e x)^2\mathbf{B}_e^T A_e/\rho_e \preceq \zeta\mathbf{H}$, and subtracting position means decreases covariance. Occupancy counts need not be independent. Store rounded packets; their possibly biased deviation from this ideal sum belongs to ν_i in (O5).

The stored recurrence and its two source estimates. Start $\mathbf{x}_0 = 0$ and use $j_R = C_R \log n$ iterations

$$x_{i+1} = \mathbf{x}_i + T_H(\mathbf{f} - \widehat{\mathbf{H}\mathbf{x}_i}),$$

including the prescribed finite evaluation and common-grid errors. Put $C_H = \mathbf{H}^{1/2} T_H \mathbf{H}^{1/2}$, $R = \mathbf{I} - C_H$, and $e_i = \mathbf{H}^{1/2}(\mathbf{x}_i - \mathbf{H}^\dagger \mathbf{f})$. With $\eta_i = \mathbf{H}^{\dagger/2}(\widehat{\mathbf{H}\mathbf{x}_i} - \mathbf{H}\mathbf{x}_i)$, the normalized error has the exact form

$$e_{i+1} = Re_i - C_H\eta_i + \nu_i, \qquad \mathbb{E}[\eta_i \mid \mathcal{F}_i, \mathbf{f}] = 0, \qquad \mathbb{E}[\eta_i \eta_i^T \mid \mathcal{F}_i, \mathbf{f}] \preceq \zeta \mathbf{I}. \tag{O5}$$

Evaluate kernels at the stored iterate. Round each packet once to a common grid and add opposite endpoint integers; the public position count bounds total error on every outcome. Choose $\max_i \|\nu_i\|$ so its geometric convolution has squared norm $o(\zeta)$. All iterations share one absolute iterate grid, closing the zero sum in each component after coordinatewise rounding.

The martingale cross terms in (O5) vanish conditionally on $\mathbf{f}$. Since R and C_H commute,

$$\sum_{j \geq 0} R^j C_H^2 R^j = C_H(2\mathbf{I} - C_H)^{-1} \preceq \mathbf{I}.$$

The normalized source covariance has trace at most one by (21); no commutation with R, C_H is needed. Choose C_R so the initial-bias second moment is at most $\zeta\mathbf{I}$. The final stored field $Y = x_{j_R}$ then satisfies, uniformly in b, T_H,

$$\mathbb{E}[\Gamma Y Y^T \mid \boldsymbol{\sigma}, \mathcal{F}] \preceq C\mathbf{H}^\dagger \Delta \mathbf{H}^\dagger + C\Gamma\zeta\mathbf{H}^\dagger, \tag{23}$$

$$\mathbb{E}[(Y - \mathbf{H}^\dagger\mathbf{f})(Y - \mathbf{H}^\dagger\mathbf{f})^T \mid \boldsymbol{\sigma}, \mathcal{F}] \preceq C\zeta\mathbf{H}^\dagger. \tag{24}$$

The numerical errors may be biased and correlated with $\mathbf{f}$; their deterministic norm reserve and $(x+\mathbf{y})(x+\mathbf{y})^T \preceq 2xx^T + 2yy^T$ suffice.

Capturing each label's change. For $Z = a_e A_e H^\dagger \mathbf{f}$, the Rademacher fourth moment is at most $3(\mathbb{E}Z^2)^2$, and $\mathbb{E}Z^2 \asymp (S|\boldsymbol{\sigma}|)_e/\Gamma$. Paley–Zygmund gives a fixed lower probability for $Z^2 \geq \mathbb{E}Z^2/2$. Equation (24) and Markov's inequality preserve a fixed fraction of that probability whenever the signal is larger than a sufficiently large multiple of $\zeta\boldsymbol{\lambda}_e$. Below that scale, the background term in (25) already bounds $(S|\boldsymbol{\sigma}|)_e$. For $J = C_J \log n$ fresh replicas, Chernoff's bound for the capture indicators and an all-label union bound therefore give

$$(S|\boldsymbol{\sigma}|)_e \leq C\Gamma \left[\frac{\boldsymbol{\varrho}_e^2}{J} \sum_{j \leq J} (\mathbf{B}_e Y_j)^2 + \zeta\boldsymbol{\lambda}_e \right] \quad \text{simultaneously for every original label.} \tag{25}$$

Replicas are independent conditional on the frozen source prestate, including b, Γ and the shared T_H. Each uses fresh source signs and product tapes.

Why the executable guards pass. Write $\mathbf{H}^{\dagger/2}\mathbf{f} = A_f \mathbf{x}^{\mathrm{rand}}$, so $\mathrm{Tr}(A_f A_f^T) \leq 1$. For a proof-only standard Gaussian g, its exponential integral and the Rademacher exponential-moment inequality give

$$\mathbb{E}_{\mathbf{x}^{\mathrm{rand}}} e^{\|A_f \mathbf{x}^{\mathrm{rand}}\|^2/4} = \mathbb{E}_g \mathbb{E}_{\mathbf{x}^{\mathrm{rand}}} e^{g^T A_f \mathbf{x}^{\mathrm{rand}}/\sqrt{2}} \leq \mathbb{E}_g e^{\|A_f^T g\|^2/4} \leq e^{1/2}.$$

Thus the source energy is $O(\log n)$ at polynomial confidence. For one product packet, putting $\mathbf{z} = \boldsymbol{\varrho}_e |A_e x|$, its normalized norm is $\mathbf{z}\sqrt{\boldsymbol{\tau}_e}/\rho_e \leq C\sqrt{\zeta/\log n}$: use the two terms in ρ_e, AM–GM, and $\boldsymbol{\tau}_e \leq C\boldsymbol{\lambda}_e$. In each vector of a proof-only orthonormal basis, the complete centered product

error has variance at most ζ, while each centered position is bounded by $C\sqrt{\zeta/\log n}$. Scalar Bernstein at threshold

$$C_{\mathrm{conf}}\big(\sqrt{\zeta\log n}+\log n\sqrt{\zeta/\log n}\big)=C'_{\mathrm{conf}}/\sqrt{n}$$

has any prescribed inverse-polynomial failure probability after fixing the constant. A union over at most n coordinates then bounds the complete error in $\mathbf{H}^{\dagger}$-norm by $O(1)$. The basis is analytical and is never computed. The contraction in (O5) therefore gives every iterate $\mathbf{H}$-energy $O(\log n)$. Its target mass is

$$\sum_e \rho_e \le C\big(\|x\|_H^2/\zeta+\log n\sum_e \boldsymbol{\lambda}_e\big)=O(n\log^2 n).$$

The owner and actual-$\boldsymbol{\lambda}$ bins add $n^{o(1)}$. On the proxy's spectral event this also bounds $x^T P_{\mathbf{H}}\mathbf{x}=O(\log n)$ and the prepared envelope, so fixed reserves make the guards pass. Scalar count bounds and the declared replica/stage counts pay the sparse-preparation work. Every reference outcome has polynomial coordinates: passing products have at most B_{prod} polynomial-size packets; failing products use deterministic contraction. Choose grids from this bound.

Expected product work. The preceding $O(\log n)$ energy bound controls guard failures. The mean energy on the full reference law gives a sharper work estimate. At every reference stage,

$$\mathbf{H}^{1/2}\mathbf{x}_i=(\mathbf{I}-R^i)\mathbf{H}^{\dagger/2}\mathbf{f}-\sum_{j<i}R^{i-1-j}C_H\eta_j+\sum_{j<i}R^{i-1-j}\nu_j.$$

Since $0\preceq R\preceq\mathbf{I}$, the source term has expected squared norm at most one. The martingale term has covariance at most $\zeta\mathbf{I}$, hence trace at most $r\zeta=O(1/\log n)$, and the numerical term is uniformly small. Therefore $\mathbb{E}\|\mathbf{x}_i\|_H^2=O(1)$ on the full reference law and $\mathbb{E}\sum_e\rho_e=O(n\log n)$. Before abort, executed steps agree with this reference; multiplying their nonnegative charges by reach indicators only decreases expectation. Guards cost $O(\textsc{Nnz}P_H+n)$ beyond shared header preparation, in the sparse row. The dense continuation has no executed cost.

9.2 Averaging the source moment before filtering

Now average the raw moment over the mask and source fields, keeping the metric and filter fixed. The scalar B below is the owner's per-frame budget B_*; the fields need not have mean zero.

Lemma 9.2. *Under (18), (O1), and (23), let the filter and M be fixed before the physical mask. For analytical notation only, put $h=\sqrt{\Gamma}Y$. Then*

$$\mathbb{E}\Gamma\le C\beta\sqrt{r},\qquad \mathbb{E}h^T\mathbf{M}h\le CB,\qquad \Big\|\mathbf{M}^{1/2}\mathbb{E}[hh^T\mid\mathcal{F}]\mathbf{M}^{1/2}\Big\|_F\le C\beta R_0/s,\qquad B=\beta m/\sqrt{r}. \tag{26}$$

No covariance bound or independence hypothesis on b is needed.

Proof. Let $\mathbf{c}_e=\mathbb{E}[|\boldsymbol{\sigma}_e|\mid\mathcal{F}]$, so $0\le\mathbf{c}_e\le Cd_e$. The normalizer bound and weighted Cauchy give

$$\mathbb{E}\Gamma\le C\sum_e\boldsymbol{\lambda}_e d_e\le C\|d\|_{\mathbf{w}}\left(\sum_e\boldsymbol{\lambda}_e^2/\mathbf{w}_e\right)^{1/2}\le C\beta\sqrt{r}.$$

The projection identity implies $S\mathbf{1}=\boldsymbol{\tau}$, and S is symmetric with nonnegative entries. Thus

$$\Big\|\mathbf{H}^{\dagger/2}\mathbb{E}\Delta\mathbf{H}^{\dagger/2}\Big\|_F^2=\mathbf{c}^TSc\le\sum_e\boldsymbol{\tau}_e c_e^2\le C\beta^2. \tag{O6}$$

The middle inequality is also the positivity of the weighted Laplacian $\textsc{Diag}(\boldsymbol{\tau}) - S$. Apply (23) and average over the mask before taking the Frobenius norm. No bound on $\mathbb{E}\|\mathbf{H}^{\dagger/2}\Delta\mathbf{H}^{\dagger/2}\|_F$ is needed. Writing $D = \mathbf{M}^{1/2}\mathbf{H}^{\dagger/2}$, we have $\|D\|_{\mathrm{op}}^2 \le CR_0/s$. Equation (O6) bounds the Frobenius norm of the signal by $C\beta R_0/s$. The background contributes at most

$$C\,\mathbb{E}\Gamma\,\zeta\sqrt{r}\,R_0/s \le C\beta R_0/s,$$

because $r\zeta = O(1/\log n)$.

For the trace bound set $v_e = \overline{\boldsymbol{\delta}}_e^{-4}/\overline{\boldsymbol{\varpi}}_e$. The signal is at most Cd^TSv. Weighted Cauchy and $(Sv)_f^2 \le \boldsymbol{\tau}_f \sum_e S_{fe}v_e^2$ give

$$d^TSv \le \|d\|_{\mathbf{w}} \left(\sum_f (Sv)_f^2/w_f\right)^{1/2}, \qquad \sum_f (Sv)_f^2/w_f \le C\sum_e \boldsymbol{\tau}_e v_e^2 \le C\sum_e \boldsymbol{\delta}_e^{-8}/\boldsymbol{\lambda}_e \le Cm/s.$$

Hence the signal is CB. The background trace is at most

$$C\,\mathbb{E}\Gamma\,\zeta\,\textsc{Tr}(\mathbf{M}\mathbf{H}^\dagger) \le C\beta\sqrt{r}\,\zeta\sum_e \boldsymbol{\delta}_e^{-4} \le CB/\log n.$$

The trace calculation avoids the extra R_0 factor in the coarse comparison $M \preceq C(R_0/s)\mathbf{H}$. □

9.3 A filter built from the maintained base

Freeze the following proxies and filter before the mask. It splits $Y = v_Y + r_Y$: residual-degree energy pays vertex clocks for v_Y, and small expected M-energy pays marks for r_Y.

Choosing and accessing the top-list set. Choose $d_0 = n^{64/p}\log^{D_0} n$. Enumerate the union C of the endpoint top-d_0 base-weight lists, using dyadic bins; let $R_E = E \setminus C$. This edge set is distinct from the owner's expander pieces. Write L_C for its weighted Laplacian and D_{R_E}, A_{R_E} for residual degree and adjacency (not incidence). Put

$$\mathbf{Q}_\omega = L_C(\omega) + D_{R_E}(\omega), \qquad \mathbf{Q}_\mu = L_C(\mu) + D_{R_E}(\mu).$$

Thus $M = \mathbf{Q}_\omega - A_{R_E}(\omega)$; replacing residual edges by their endpoint diagonals permits the trace bound below. Maintain full base degrees at genuine coefficient changes. Subtract the enumerated top-list incidences to obtain $D_{R_E}(\mu)$ in $O(nd_0)$ work, then add the explicit $L_C(\mu)$ to form $\mathbf{Q}_\mu$. Handle sparse components without residual edges directly; on the others $\mathbf{Q}_\omega, \mathbf{Q}_\mu$ are positive definite.

At each original vertex, maintain the integer index of Lemma 8.1, keyed by the negative dyadic exponent of μ_e and the original-label identifier. Shift the exponent by its fixed range bound to obtain a nonnegative polynomial-universe key. Initial insertion costs $\ddot{O}(m)$. Each genuine base-bucket change updates two keys at $\ddot{O}(1)$ cost each; private owner migrations do not change these keys.

Repeatedly read and temporarily delete the minimum key until d_0 entries or exhaustion, then reinsert those entries. Only this auxiliary index changes; owners, coefficients, physical state and output laws stay fixed. These operations cost $\ddot{O}(nd_0+n)$, without successor queries or dense scans. Generation marks and direct incidence handles deduplicate the union and reconcile roles within the same sparse query row. Identifier ties make the set canonical at the frozen base.

The top-list trace bound. If $0 < c_\star \le 1$, selected weights at a vertex are at least $c_\star$ times every residual weight there, and parallel-label multiplicity is at most b_0, then

$$\textsc{Tr}(\mathbf{Q}_\omega^{-1} A_{R_E} \mathbf{Q}_\omega^{-1} A_{R_E}) \le 32 b_0 n/(c_\star d_0). \tag{27}$$

To prove it, let t_v be the largest residual weight at v, or zero if no residual edge meets v. Give a chosen edge uv auxiliary weight

$$\tfrac{c_\star}{2}\big(t_u \mathbf{1}_{\{u \text{ chooses } uv\}} + t_v \mathbf{1}_{\{v \text{ chooses } uv\}}\big).$$

Let $\mathbf{H}'$ be the graph of these auxiliary selected edges and D the degree matrix of $\mathbf{H}' + R_E$. Orient a selected edge from a choosing endpoint of largest t, and orient each residual edge arbitrarily. In $\mathbf{H}' + R_E$, each oriented edge weight divided by its tail degree is at most $2/(cd_0)$. Its normalized adjacency therefore has squared Frobenius norm at most $4b_0 n/(cd_0)$. Writing $B_0 = D^{-1/2} A_{H'} D^{-1/2}$ and $C_0 = D^{-1/2} A_{R_E} D^{-1/2}$, nonnegative entries give the same bound on $\|B_0\|_F^2 + \|C_0\|_F^2$. The matrix

$$Z = (\mathbf{I} - B_0)^{-1/2} C_0 (\mathbf{I} - B_0)^{-1/2}$$

is a contraction since $L_{H'} + D_{R_E} \pm A_{R_E} \succeq 0$. The subspace spanned by eigenvectors of B_0 with eigenvalues above $1/2$ has dimension at most $4\|B_0\|_F^2$. Since Z is a contraction, the blocks meeting this subspace have total squared Frobenius norm at most twice its dimension. The block on its orthogonal complement has squared Frobenius norm at most $4\|C_0\|_F^2$, proving the auxiliary bound. Finally $\mathbf{Q}_\omega \succeq L_{H'} + D_{R_E}$. For symmetric $\mathbf{B}$, $\textsc{Tr}(YAYA)$ is increasing in PSD Y, since its directional derivative is $2\textsc{Tr}(\Delta AYA) \ge 0$. Inverse order transfers the bound to $\mathbf{Q}_\omega$. Regularization handles isolated auxiliary rows annihilated by A_{R_E}. Clamp any supplied comparison constant to $\min\{c, 1\}$, since a selected auxiliary edge has tail ratio at most $2/d_0$. Here $c_\star = \min\{1, c'_\star/R_0\}$ gives $O(R_0 n/d_0)$.

Sparse proxy construction. Prepare proxies $P_M = (1 \pm \epsilon)M$ and $\widehat{\mathbf{Q}}_\omega = (1 \pm \epsilon)\mathbf{Q}_\omega$, with $\epsilon = c_\star d_0^{-1/32}$, before the physical mask. The filter multiplies by P_M and approximately inverts $\widehat{\mathbf{Q}}_\omega$. Both target actual ω; the μ-weighted base supplies sampling access. Base leverage upper scores $\mathbf{B}_e$, distinct from incidence rows $\mathbf{B}_e$, select residual edges for P_M; coordinate scores $z_v \asymp (\mathbf{Q}_\mu^{-1})_{vv}$ select residual endpoint atoms for $\widehat{\mathbf{Q}}_\omega$. Encode each set before its sampling tape.

The base scores come from the immediately preceding frame's cached proxy. The same construction refreshes this cache to approximate the current G_0 for the next frame. Lemma 6.2 supplies the first cache and its finite categorical construction. Discard or restore aborted caches consistently with the clean checkpoint.

Base-score fields from the previous cache. Write the immediately preceding base proxy as the stored sparse sum

$$P_{\text{old}} = \sum_\alpha c_\alpha b_\alpha b_\alpha^T.$$

Here b_α is the signed incidence vector of copy α. The supplied one-frame coefficient ratio makes this proxy constant-comparable to the current G_0. Freeze this proxy and a fixed symmetric inverse map T_{old} before the field signs. For a fixed small constant $\eta_s > 0$, choose short dyadics

$$(1 - \eta_s)\sqrt{c_\alpha} \le d_\alpha \le \sqrt{c_\alpha}, \qquad \widetilde{P}_{\text{old}} = \sum_\alpha d_\alpha^2 b_\alpha b_\alpha^T.$$

The coefficients d_α define the source exactly. With fresh independent signs on the sparse atoms, define the analytical lanes

$$\mathbf{x}_j^{\mathrm{rand}} = \sum_\alpha d_\alpha \sigma_{\alpha j} b_\alpha, \qquad h_j = T_{\mathrm{old}} \mathbf{x}_j^{\mathrm{rand}}, \qquad J_0 = 2^{\log_2(C_J \log n)}.$$

Conditional on the frozen proxy and inverse,

$$\mathbb{E} h_j h_j^T = T_{\mathrm{old}} \widetilde{P}_{\mathrm{old}} T_{\mathrm{old}} \asymp G_0^\dagger.$$

Indeed $\widetilde{P}_{\mathrm{old}} \asymp P_{\mathrm{old}}$, and if $C_{\mathrm{old}} = P_{\mathrm{old}}^{1/2} T_{\mathrm{old}} P_{\mathrm{old}}^{1/2}$ satisfies $\|C_{\mathrm{old}} - \mathbf{I}\|_{\mathrm{op}} \le \eta_T < 1$, then

$$(1-\eta_T)^2 P_{\mathrm{old}}^\dagger \preceq T_{\mathrm{old}} P_{\mathrm{old}} T_{\mathrm{old}} \preceq (1+\eta_T)^2 P_{\mathrm{old}}^\dagger.$$

Stored base fields and canonical upper scores. Before the proxy-position tapes, store the finite lanes $\widehat{h}_j$ on a predetermined dyadic grid. For a predetermined dyadic score floor $\gamma > 0$, choose inverse accuracy and the grid so that, on the valid-proxy/inverse reference,

$$\left| \mu_e(\mathbf{B}_e \widehat{h}_j)^2 - \mu_e(A_e h_j)^2 \right| \le \gamma \quad \text{for every original label, lane, and reference outcome.} \tag{CF1}$$

Here γ is unrelated to the physical step ratio β_s/β_f. Public polynomial bounds give $|A_e h_j| \le U$ and $\mu_e \le \mu_{\max}$. A uniform coordinate error $\varepsilon_{\mathrm{abs}}$ changes the squared score by at most $4\mu_{\max}\varepsilon_{\mathrm{abs}} U + 4\mu_{\max}\varepsilon_{\mathrm{abs}}^2$. Choose the inverse accuracy and the cache and scratch grids of Section 12 to make this at most γ. Other factorization tapes obey the finite arithmetic guards or abort.

Let $\mathrm{up}_\gamma(x) = \gamma \lceil x/\gamma \rceil$. Define the actual stored-label upper-score function to be

$$\mathbf{B}_e = C_A \mathrm{up}_\gamma \left(\frac{\mu_e}{J_0} \sum_{j \le J_0} (\mathbf{B}_e \widehat{h}_j)^2 + \gamma \right), \tag{CF2}$$

where C_A is one sufficiently large fixed integer. Frozen finite fields and stored μ_e fix this value independently of proposals and reads. The γ term covers (CF1), even when the stored squared distance is zero, and

$$\mathbf{B}_e \le C_A \left\lceil \frac{\mu_e}{J_0} \sum_j (\mathbf{B}_e \widehat{h}_j)^2 + 2\gamma \right\rceil. \tag{CF3}$$

Choose $m\gamma \le 1$. Scalar Bernstein for averages of squared Rademacher linear forms [Tro12, Theorem 6.2, scalar case], followed by the all-label union bound, gives simultaneously

$$\boldsymbol{\tau}_e(G_0) \le \mathbf{B}_e, \qquad \sum_e \mathbf{B}_e = O(n).$$

Score envelopes and cache refresh. Mix the J_0 base-owner envelopes with weights C_A/J_0, and add a uniform-label envelope of rate $2C_A\gamma$. By (CF3), this is an upper envelope for the actual $\mathbf{B}_e$. Its mass is bounded by

$$\frac{C_A A_L}{J_0} \sum_j \|\widehat{h}_j\|_{G_0}^2 + 2C_A m\gamma. \tag{CF4}$$

Prepared header masses give the normalizer without a dense score sum or original-graph quadratic form. Covariance and (CF1) give unconditional expectation $O(nA_L)$ for (CF4), before any field-work cap. Thin each proposal to (CF2) and use that actual rate also in the packet denominator. The refresh and residual rates are

$$\rho_e^{\mathrm{cache}} = C \log n \mathbf{B}_e, \qquad \rho_e^{M} = C R_0 \epsilon^{-2} \log n \mathbf{B}_e \mathbf{1}_{\{e \in R_E\}}.$$

Retain the top-list edges C exactly in P_M and sample residual atoms at rate ρ_e^M. At rate ρ_e^{cache}, use fresh positions to sample a constant-accuracy proxy of the entire current G_0, with the prescribed positive common-grid packet rounding. Store it as the next frame's base cache. This refresh occurs before the physical mask and adds $n^{1+o(1)}$ sparse atoms and work after the shared $n(R+\log n)n^{o(1)}$ owner-header preparation, where $R \asymp n^{1/16}$ is the counter radix. Use the same public-count kernel with fresh position tapes after fixing targets, without redrawing fields. Each cache samples current G_0, so approximation factors do not multiply across frames.

Diagonal scores for the second proxy. For the coordinate scores of $\mathbf{Q}_\mu$, use its explicit top-list edge atoms and the diagonal atoms of $D_{R_E}(\mu)$, with the latter regarded as edges to ground. The same atom-sign construction has analytical covariance comparable to $\mathbf{Q}_\mu^{-1}$. Compute all nJ_0 finite coordinates in the sparse row, and define z_v by the counterpart of (CF2), with no edge coefficient. Choose its numerical mesh γ_Q so that

$$\gamma_Q \sum_v (D_{R_E}(\mu))_{vv} \le 1, \qquad \gamma_Q \le c_\star / \lambda_{\max}(\mathbf{Q}_\mu).$$

Fix these meshes from public polynomial bounds. Since $(\mathbf{Q}_\mu^{-1})_{vv} \ge 1/\lambda_{\max}(\mathbf{Q}_\mu)$, the resulting scores satisfy the stated two-sided coordinate comparison and

$$\sum_v z_v (D_{R_E}(\mu))_{vv} \le C \mathrm{Tr}(D_{R_E}(\mu)\mathbf{Q}_\mu^{-1}) + C \le Cn.$$

Sample residual endpoint atoms at rates $CR_0\epsilon^{-2}\log n \mu_e z_v$. Endpoint upper weights within a dyadic bucket differ by at most a factor two. Excluded top-list positions split each compact array into gaps; sampling the gaps avoids rejection from an array almost entirely in C. Each proposal thins to its exact original-label target. Endpoint packets use the literal stored z_v in both their position probabilities and their inverse-rate weights.

The $O(J_0) = O(\log n)$ point cost of (CF2) and the J_0 field/header preparations occur only in sparse proxy rows, where fixed powers of $\log n$ are allowed. Dense selected events and fine reads priced by (O8) never invoke this lane scan.

For either observer proxy, normalize by its target matrix. A retained categorical-position PSD atom has norm at most $R_{\mathrm{pkt}} = C\epsilon^2/\log n$; its centered version has norm at most $2R_{\mathrm{pkt}}$, and the sum of its second moments is at most $R_{\mathrm{pkt}}\mathbf{I}$. Independent matrix Bernstein [Tro12, Theorem 1.4], applied to both signs, gives the desired two-sided error at polynomial confidence. This use allows the exact retained top-list edges to have arbitrary atom size: only random residual atoms enter the variance bound.

The filter's operator and Frobenius bounds. Let T_Q be a fixed symmetric inverse of $\widehat{\mathbf{Q}}_\omega$ refined to accuracy $O(\epsilon)$, and set $F = \mathbf{I} - T_Q P_M$, $k_{\mathrm{it}} = 16$. For an operator on the Laplacian quotient, use $X_M = \mathbf{M}^{1/2} X M^{\dagger/2}$ for its M-normalized form. Since $P_M\mathbf{1} = 0$, F preserves constants and induces this quotient operator. With $K = \mathbf{I} - \mathbf{M}^{1/2}\mathbf{Q}_\omega^{-1}\mathbf{M}^{1/2}$ on the quotient,

$$\|K\|_{\mathrm{op}} \le 1, \qquad \|K\|_F^2 \le CR_0 n/d_0, \qquad F_M = K + E, \quad \|E\|_{\mathrm{op}} \le C\epsilon.$$

Expand $(K+E)^{k_{\mathrm{it}}}$ in the original factor order. Every term except E^k contains a K; bound one such factor in Frobenius norm and all remaining factors in operator norm. The remaining term has squared Frobenius norm at most $n(C\epsilon)^{2k} = O(n/d_0)$. Hence the conservative bounds

$$\|(F^{k_{\mathrm{it}}})_M\|_{\mathrm{op}} \le C, \qquad \|(F^{k_{\mathrm{it}}})_M\|_F^2 \le CR_0^2 n/d_0 \tag{28}$$

hold without any commutation assumption.

The complement, including finite evaluation. Since $D_{R_E} \preceq \mathbf{Q}_\omega$, $\|\mathbf{Q}_\omega^{1/2} F Q^{-1/2}\|_{\mathrm{op}} = O(1)$, and $P_M Q^{-1} P_M \preceq CM$, the identity

$$\mathbf{I} - F^{k_{\mathrm{it}}} = \sum_{j<k} F^j T_Q P_M$$

gives $\|(\mathbf{I} - F^{k_{\mathrm{it}}})h\|_Q^2 \le Ch^T \mathbf{M} h$. For finite evaluation use a fixed k-layer grid ladder. Fix a positive dyadic $\nu_{\mathrm{num}} \le c_{\mathrm{num}} d_0^{-1/4}$, also below the ordinary numerical reserve, before choosing that ladder. Here $c_{\mathrm{num}} > 0$ is a sufficiently small fixed constant. Since d_0 has a fixed polynomial range, this requires only $O(\log n)$ bits and does not change the fixed number of layers. At layer j, form $P_M z$ exactly on the product grid; if it is zero, use the exact no-op. Otherwise its grid and the polynomial upper eigenvalue bound on $\mathbf{Q}_\omega$ provide a positive lower bound on $\|P_M z\|_{Q^{-1}}$. A sufficiently finer predetermined next grid gives action error at most $\nu \|P_M z\|_{Q^{-1}}$ in $\mathbf{Q}_\omega$-norm. Because k is fixed, all grid exponents remain $O(\log n)$. For the stored output r_Y and literal subtraction $v_Y = Y - r_Y$,

$$v_Y^T D_{R_E}(\omega) v_Y \le CY^T M Y, \qquad \|r_Y - F^k Y\|_Q \le C\nu \|Y\|_M. \tag{29}$$

The numerical error is controlled in the full-space norm.

The expected residual energy. Apply Lemma 9.2 with the fixed filter, and use $\|X^T X\|_F \le \|X\|_{\mathrm{op}} \|X\|_F$. For analytical notation $r_h = \sqrt{\Gamma} r_Y$, Schatten Cauchy and (28) give

$$\mathbb{E}\|r_h\|_M^2 \le C(\beta R_0/s) R_0 \sqrt{n/d_0} + C\nu^2 \mathbb{E} h^T \mathbf{M} h \le CBR_0^2/\sqrt{d_0}. \tag{30}$$

The numerical term uses $\mathbb{E} h^T \mathbf{M} h = O(B)$, $\nu^2 \le c_{\mathrm{num}}^2/\sqrt{d_0}$, and $R_0 \ge 1$. Implement the filter at Y and multiply its squares by Γ; this avoids rounding $\sqrt{\Gamma} Y$ or assuming homogeneity of finite evaluation.

9.4 From fields to read requests and lifetime payments

A public score epoch runs between consecutive public fine reads of one label. Its price and threshold stay fixed:

$$\pi_e = \overline{\boldsymbol{\delta}}_e^{-3}, \qquad \Theta_e = c_\star \, \overline{\boldsymbol{\varpi}}_e \overline{\boldsymbol{\delta}}_e / \boldsymbol{\varrho}_e^2, \qquad \pi_e/\Theta_e \asymp \omega_e.$$

The *physical debt* below is in squared-field units. Multiplying by epoch-fixed $\boldsymbol{\varrho}_e^2$ and the capture constants bounds score variation; multiplying by π_e/Θ_e gives read cost.

Positive accounts. For replica Y_j, write $r_j = r_{Y_j}$ and $v_j = Y_j - r_j$. Labels in the top-list set C use a direct account; residual labels use the split

$$(\mathbf{B}_e Y_j)^2 \le 4\big(v_j(u)^2 + v_j(v)^2\big) + 2(A_e r_j)^2.$$

Vertex complement clocks pay the first term and marks the second. The direct account receives $\Gamma J^{-1}\sum_j(\mathbf{B}_eY_j)^2$, each endpoint complement account receives $4\Gamma J^{-1}\sum_j v_j(v)^2$, and the global background clock increments by $\Gamma\zeta$, with the epoch-fixed conversion $\boldsymbol{\lambda}_e/\boldsymbol{\varrho}_e^2$. Thus its account receives $\Gamma\zeta\boldsymbol{\lambda}_e/\boldsymbol{\varrho}_e^2$, all in physical-debt units. A label in C has direct and background accounts; a residual label has two endpoint accounts and background. Give each account ceiling $\vartheta_e=\Theta_e/8$. The timers below release a read only after $\Theta_e/16$ positive debt, so fix $c_a=1/16$.

Residual marks. For each replica, use the target

$$t_{e,j}=\frac{16\Gamma}{J\Theta_e}(A_er_j)^2\mathbf{1}_{\{e\in R_E\}}.$$

Accumulated target one corresponds to residual physical debt $\Theta_e/8$, the tolerance used below. Take fresh public-count categorical/null batches with multiplier $\lambda_{\mathrm{mon}}=C_{\mathrm{conf}}\log n$. A raw proposal reads only that replica's two stored coordinates, so its point price has no factor J. The product no-hit probability is at most $\exp(-\lambda_{\mathrm{mon}}\sum_j t_{e,j})$. The public count is fixed using the all-outcome polynomial range, not a successful future source or work event.

Buffering the first request. At a label's first fine-score request, stamp and append it once, remove its positive-account timers, and deactivate those accounts until its single final read starts a new epoch. For selected labels, do this before any observer clock increment. Later marks only test the stamp; role and owner moves preserve inactivity. Central-only requests do not suspend accounts. Only deadline metadata changes: frozen coefficients, fields, sampling lists and old read prices remain until the old batches finish. Each first request costs a fixed number of handle operations, independent of the replica count.

The no-hit guarantee. Stop a label's no-hit process at its first buffered public fine-score request; central-only requests and inspections leave it unchanged. Conditional on the generated fields and the history before the fresh mark coins, the no-request indicator times the exponential of accumulated predictable residual intensity is a nonnegative supermartingale. Stopping when the target reaches one bounds failure by $\exp(-c_\star\lambda_{\mathrm{mon}})$, conditional on a valid epoch birth at zero, uniformly over adaptive later fields. This bound is conditional on the birth, not on each intermediate no-hit history. Union over the predeclared polynomial number of adaptive births, stopping at earlier primitive failures. A clean refresh restarts epochs at zero; other fine-score requests kill the process and improve the bound. The proof-only target survives all role and gauge changes until a public read.

Protecting the reference tube. By (25) and the inverse-order consequence

$$-(1-q)^{-1}\mathbf{H}^\dagger\Delta\mathbf{H}^\dagger\preceq(\mathbf{H}')^\dagger-\mathbf{H}^\dagger\preceq(1-q)^{-1}\mathbf{H}^\dagger\Delta\mathbf{H}^\dagger,$$

the four debt terms protect the tube until the first request. The displayed inverse bounds follow from $\mathbf{H}-\Delta\preceq\mathbf{H}'\preceq\mathbf{H}+\Delta$. Choose debt tolerances as follows. Let K_{obs} bound the fixed capture and inverse-order constants, and let $C_t,C_\delta\ge1$ be fixed comparison constants with $\overline{\boldsymbol{\varpi}}_e\le C_t(s+\boldsymbol{\tau}_e)$ and $\overline{\boldsymbol{\delta}}_e\le C_\delta\boldsymbol{\delta}_e$ up to the first potential reference exit. A label that has not been requested is unselected, so its $\boldsymbol{\varrho}_e,\boldsymbol{\lambda}_e,\overline{\boldsymbol{\varpi}}_e,\overline{\boldsymbol{\delta}}_e$ stay fixed throughout this epoch. If no positive account reports, its total positive debt is below $3\Theta_e/8$. If the residual no-hit event has not failed, its accumulated residual physical

debt $2\sum_{t,j}\Gamma_t(A_e r_{t,j})^2/J$, restricted to its residual roles, is below $\Theta_e/8$. Consequently its total captured score variation is at most $K_{\mathrm{obs}}\varrho_e^2\Theta_e/2$, and

$$\frac{|\log(s+\boldsymbol{\tau}_e)-\log(s+\boldsymbol{\tau}_{e,\mathrm{birth}})|}{\kappa}\le\frac{K_{\mathrm{obs}}C_t c_{\overline{\boldsymbol{\delta}}}\overline{\boldsymbol{\delta}}_e}{2\kappa}\le\frac{\overline{\boldsymbol{\delta}}_e}{1024pC_\delta}$$

when $c_\star\le\kappa/(512pC_\delta K_{\mathrm{obs}}C_t)$. Choose the fresh fine-read and tracking-cell rounding constants to give the same $\overline{\boldsymbol{\delta}}_e/(1024pC_\delta)$ birth error in normalized tracking coordinates. Their sum is at most $\boldsymbol{\delta}_e/(512p)$, leaving a strict reserve in the required $\boldsymbol{\delta}_e/(256p)$ observation window. All clock quantities are upper enclosures of their physical contributions; the prescribed source, capture, and fine-read numerical reserves are included in these constants. The argument uses the comparisons only before their first possible violation. The physical one-frame margin reserve protects the wait until the final read.

Transferring debt when a role changes. At a top-list role change, materialize each touched label's direct/complement debt and due flags, then transfer them to one canonical account without duplication. Preserve the background baseline and conversion. Queue a read paid by old debt exactly when the new $\Theta_e/8$ ceiling or early-report rule applies. In either case, preserve debt and reference metadata until the final public read starts the next epoch. Touch a fixed number of accounts for each of $O(nd_0)$ labels. Constant-unit transfers are exact on the common debt grid; upward conversions may instead add total weighted numerical credit at most one under the fixed horizon and precision. Never revalue debt at a new π_e.

The lifetime payment identity. Here D_e denotes accumulated physical debt, distinct from the owner's density deficit. Cap its proof-only active credit at the price of one read:

$$C_e=\pi_e\min\{1,D_e/(c_a\Theta_e)\}.$$

A first positive-account request releases π_e. A first selected or residual-mark request discards any old credit, because its read has another payment. The label then stops accruing credit until its final read starts a new epoch with zero debt and zero credit. Role changes preserve C_e, including immediately due ones. Writing A_t for new credit and R_+ for all released positive-account prices, exact bookkeeping gives, at every stopped prefix,

$$R_+\le C_{\mathrm{initial}}+\sum_t A_t,\qquad\mathbb{E}[A_t\mid\mathcal{F}_t]\le CB_t. \tag{31}$$

For the mean, an added debt increases credit by at most $C\omega_e\,\Delta D_e$. Summing the direct and complement increments gives

$$\frac{C\Gamma}{J}\sum_j\left(Y_j^TL_C(\omega)Y_j+v_j^TD_{R_E}(\omega)v_j\right).$$

Equations (29) and (26) bound its conditional expectation by CB_t. The background contributes at most

$$C\,\mathbb{E}\Gamma\,\zeta\sum_e\boldsymbol{\lambda}_e/(\overline{\boldsymbol{\varpi}}_e\overline{\boldsymbol{\delta}}_e^4)\le CB_t/\log n.$$

Suspension, clipping, and discarding credit can only decrease it. The credits are analytical; no dense clock update is performed.

Paying for marks and their raw proposals. The expected residual-mark price, counting even duplicate marks, is

$$\mathbb{E}\sum_e \pi_e \,\#\{\text{marks of } e\} \le C\log n\,\mathbb{E}\|r_h\|_M^2.$$

Generate marks from the full base envelope (O2), testing residual membership during thinning. Since π/Θ is comparable to the actual encoded ω, fix $K_\omega \ge \sup_e \pi_e/(\Theta_e\omega_e)$. For replica r_j, this envelope has per-label mass $E_e \ge \mu_e(A_e r_j)^2$ and total mass at most $A_L r_j^T G_0 r_j$. Give it raw rate $(16\lambda_{\mathrm{mon}}\Gamma/J)K_\omega R_0 E_e$ in the public-count categorical construction, and thin each proposed label by

$$\frac{\mu_e(A_e r_j)^2}{E_e}\,\frac{\omega_e}{R_0\mu_e}\,\frac{1}{\pi_e}\,\frac{\pi_e}{K_\omega\Theta_e\omega_e}\,\mathbf{1}_{\{e\in R_E\}}.$$

Each factor lies in $[0,1]$, and a zero envelope returns zero, giving exactly rate $\lambda_{\mathrm{mon}}t_{e,j}$. Comparability of encoded weights therefore affects domination and work only. Enlarge the public count by K_ω before the run; its point correction is fixed-depth rational arithmetic without a fine read. Thus no spectral certificate is restricted to the adaptive residual. The raw-proposal expectation is at most $C\log nR_0A_L\,\mathbb{E}\|r_h\|_M^2$. Consequently the sum of new credit, marked-read price, and raw proposals has expectation

$$CB_t\left(1+\frac{\log nA_LR_0^3}{\sqrt{d_0}}\right)=O(B_t). \tag{32}$$

The ratio has polynomial factor $n^{-8/p}$; choose D_0 after the fixed logarithmic degree of the envelope. The normalization $1/J$ pays summing all J replicas. Selected reads are charged separately through (O3).

Implementing the positive clocks. Use dyadic FIFO timers. For an account, let ϑ be its physical-debt ceiling, D its current physical debt, and $k>0$ its epoch-fixed conversion from clock increments to physical debt. In particular, the background account has $k=\boldsymbol{\lambda}_e/\boldsymbol{\varrho}_e^2$. Use a common clock mesh Δ_C so fine that $k\Delta_C \le \vartheta/2$ for every permitted account. If $D\ge\vartheta$, or if $0<\vartheta-D<k\Delta_C$, queue a request immediately. Otherwise install a timer at the current clock plus

$$R=\Delta_C\,2^{\left\lfloor \log_2 \frac{\vartheta-D}{k\Delta_C}\right\rfloor}.$$

Each radius R has a FIFO: the clock and new equal-radius deadlines are nondecreasing. Direct handles remove a requested label's timers; an update scans $O(\log n)$ radius headers and reports due prefixes. A timed report has accumulated physical debt at least $D+kR\ge\vartheta/2$; the immediate submesh branch also has $D\ge\vartheta-k\Delta_C\ge\vartheta/2$. Absence of a report implies physical debt below ϑ. These half-ceilings are at least $c_a\Theta_e$ for the choice $c_a=1/16$ above. Role changes apply the debt-transfer rule and reinstall timers, preserving due flags, background baseline and conversion, and single payment. Insertion costs $\ddot{O}(1)$; scanning $O(\log n)$ headers per updated clock belongs to the $O(nd_0+n)$ sparse clock row.

9.5 Packed fine reads

After committing selected coefficients and before new metadata transactions, prepare final-graph fine fields from the actual-$\boldsymbol{\lambda}$ index. Rates $C\log n\epsilon_{\min}^{-2}\boldsymbol{\lambda}_e$, with $\epsilon_{\min}\asymp n^{-2/p}$, give a fine proxy without fine-score queries. Freeze graph, proxy and inverse before drawing the $J_{\max}=\Theta(\log n\epsilon_{\min}^{-2})$ sign lanes; retain this bank through all adaptive metadata requests and prefix choices.

Covariance and canonical values in the fine bank. For the frozen final graph write its fine sparse proxy as

$$P_f = \sum_\alpha c_\alpha b_\alpha b_\alpha^T = (1 \pm \epsilon_f) H_{\text{final}}, \qquad \epsilon_f \le c_f \delta_{\min}.$$

Take a fixed symmetric inverse T_f of P_f with relative error at most $c_f\delta_{\min}$. Choose source dyadics d_α with relative error at most $c_f\delta_{\min}$ from $\sqrt{c_\alpha}$, and define, before drawing the sign lanes,

$$\widetilde{P}_f = \sum_\alpha d_\alpha^2 b_\alpha b_\alpha^T, \quad \mathbf{x}_j^{\text{rand}} = \sum_\alpha d_\alpha \sigma_{\alpha j} b_\alpha, \quad h_j = T_f f_j.$$

The signs are independent over sparse atoms and lanes. The exact analytical covariance is

$$\mathbb{E}[h_j h_j^T \mid P_f, T_f, (d_\alpha)] = T_f \widetilde{P}_f T_f = (1 \pm C c_f \delta_{\min}) H_{\text{final}}^\dagger. \tag{CF5}$$

Store $\widehat{h}_j$ on a predetermined common grid. As in (CF1), polynomial ranges give deterministic squared-score error small compared with $\delta_{\min}\lambda_{\min}$ on every outcome of the frozen valid-proxy/inverse reference.

The certified old margin δ determines how many lanes of the fixed bank to read. Its permitted prefix is the canonical integer

$$J(\delta) = 2^{\log_2(C_{\text{fine}} \log n \delta^{-2})}, \qquad J_{\max} \ge \max_{\delta \ge \delta_{\min}} J(\delta).$$

Conditional on the frozen proxy, inverse, and source coefficients, scalar Bernstein for the squared Rademacher linear forms in (CF5) and a union bound over the polynomial family of labels and prefixes give additive score error $c_\star \delta \boldsymbol{\lambda}_e$ and empirical energy $O(\boldsymbol{\lambda}_e)$, at polynomial confidence after increasing C_{fine}. On the same event the Rademacher moment bound gives

$$|a_e A_e h_j| \le C\sqrt{\log n \boldsymbol{\lambda}_e}. \tag{O7}$$

This uses $\boldsymbol{\tau}_e \le C\boldsymbol{\lambda}_e$; fields and prefixes require no current fine-score query.

The old stored margin and current stored $\boldsymbol{\varrho}_e^2, \boldsymbol{\lambda}_e$ determine scale, prefix and rounding. Compute the quantized squared distance by (O8), add predetermined positive reserves for proxy, sampling, quantization and evaluation errors, and round upward on a score grid finer than a fixed fraction of $\delta_{\min}\lambda_{\min}$. This canonical upper score may be capped by the certified upper bound $C\boldsymbol{\lambda}_e$. For the prefix length $J = J(\delta)$, the squared-distance perturbation from quantization is bounded by

$$2q \left(J^{-1} \sum_j (\mathbf{B}_e \widehat{h}_j)^2 \right)^{1/2} + q^2,$$

so $q = c_{\text{fine}}\delta\sqrt{\boldsymbol{\lambda}_e}/\boldsymbol{\varrho}_e$ incurs additive $O(c_{\text{fine}}\delta\boldsymbol{\lambda}_e)$ score error on the simultaneous empirical-energy event, without an extra $\sqrt{\log n}$ loss.

Simultaneous prefix accuracy permits adaptive requests without selective redraw. Sparse sources, inverse lanes and bit planes cost

$$n^{1+8/p+o(1)}$$

per frame, as in the fine-inverse preparation row.

The packed query. At a requested label use its old certified $\delta = \overline{\boldsymbol{\delta}}_e$, using the canonical prefix $J_e = J(\delta)$ fixed above, before inspecting the score. Pack $\Theta(\log n)$ lanes into each word and store coordinate bits in planes. Choose a dyadic scale satisfying

$$\frac{c_{\mathrm{fine}}\delta\sqrt{\boldsymbol{\lambda}_e}}{2\boldsymbol{\varrho}_e} \le q \le \frac{c_{\mathrm{fine}}\delta\sqrt{\boldsymbol{\lambda}_e}}{\boldsymbol{\varrho}_e},$$

with a fixed sufficiently small $c_{\mathrm{fine}} > 0$. At this scale, bound (O7) requires only $b = \ddot{O}(\log(1/\delta))$ planes. Modular bit-sliced subtraction recovers the unique signed difference; high common offsets cancel, with sign extension for negative integers. Bit-sliced squaring and masked pair additions sum its squares in

$$O\big(\delta^{-2}(\log\log n + \log(1/\delta))^2\big) \le \ddot{O}(\delta^{-3}) \tag{O8}$$

word operations. No unit-cost population-count instruction is used. The reverse comparison $\boldsymbol{\lambda}_e \le C(s + \boldsymbol{\tau}_e)$ converts the additive $O(\delta\boldsymbol{\lambda}_e)$ score error into $O(\delta)$ relative error in $s + \boldsymbol{\tau}_e$. After the comparison constants and before the bank, fix c_{fine}, proxy accuracy and lane count so the combined log-score error divided by κ is at most $\delta/(256p)$. Fixed $p = 512$ and κ leave logarithmic exponents unchanged.

The one-frame margin comparison validates this old precision on the final graph. Choose the accuracy and tracking-cell placement constants with that comparison reserve so the new epoch's birth error is at most $\bar{\delta}_{e,\mathrm{new}}/(1024pC_\delta)$ in normalized tracking coordinates, as required by the debt tolerances above. Direct dyadic lower-margin selection supplies the interior reserve without repeated precision-discovery reads.

9.6 Packed conserved-source assembly

Prepare each $\ddot{O}(1)$-bit $\widehat{\mathbf{c}}_e$ once per selected label. One independent sign word per label and lane block suffices for $J = C_J \log n$, a fixed number of blocks. Emit identical opposite endpoint packets without materializing J scalar copies.

Exact modular accumulation. Encode every short dyadic in one fixed dyadic source unit. Its integer encoding is an $\ddot{O}(1)$-bit mantissa shifted by a known nonnegative number of planes. Polynomial coefficient ranges give an $O(\log n)$-bit bound on the absolute packet sum. Choose a $B_0 = \Theta(\log n)$-bit modular range greater than twice this deterministic bound. At frame start, pad each destination accumulator independently by a uniform B_0-bit integer in each lane.

Fix packet order from selected labels and mathematical data; neither order nor mathematical stop tests inspect pads or carry lengths. Skip planes below the first nonzero bit. Apply Boolean full-adders and full-subtractors on the $\ddot{O}(1)$ mantissa planes to positive and negative masks separately, then continue until the carry/borrow mask vanishes or reaches the modular boundary. Negative packets subtract magnitudes, avoiding sign-extension work. Only Boolean operations and mask zero tests are used; modular wrap triggers neither operand failure nor a pad-dependent stop.

After the final packet has reached both endpoints, remove destination pads exactly on a fixed full-width schedule. Decode and transpose the unique signed sums before any inverse, Richardson product or source-dependent header query. The range bound ensures exact endpoint conservation without intermediate flushes. A work-cap interruption aborts the whole attempt before any incomplete-source query.

Expected carry cost. Condition on the checkpoint, preceding attempts, and the current attempt's entire mathematical history with actual-work and padding-cost caps ignored. Exact unpadding leaves fields unchanged; neither fields nor mathematical stops inspect pads or carry costs. Each lane before addition is its uniform pad plus a fixed prefix, hence uniform modulo 2^{B_0}. Its upper bits remain independent unbiased bits conditional on lower bits, including whether they initiate a carry or borrow. Survival through t extra planes has probability at most 2^{-t} per lane, and at most $\min\{1, C \log n 2^{-t}\}$ after the lane union bound. Thus each packet has expected $\ddot{O}(1)$ work, even though different packets' carry costs need not be independent. Initialization and removal cost $O(\log n)$ per destination and frame. It follows that

$$\mathbb{E}[W_{\text{source}} \mid \text{one attempt's mathematical history}] \le C(M_{\text{pkt}} \log\log n + nT \log n). \tag{33}$$

The final transpose to scalar field vectors costs $\ddot{O}(nT \log n)$, included in S_{sparse}. Fresh pads are used for each restarted attempt. Conditioning on the actual cost-aborted history, or on a whole restarted trajectory influenced by earlier pad costs, would not justify (33).

9.7 Clean checkpoints and the stopped expectation

Coalesce selections, account requests and residual marks in one public queue, without resampling physical masks. Finish old source and replica batches, commit selected coefficients, update owners and actual-$\boldsymbol{\lambda}$ bins, prepare final fields, then read each queued label once at its old price. Close its old epoch before replacing reference $\boldsymbol{\varrho}, \boldsymbol{\lambda}$, $\bar{t}, \overline{\boldsymbol{\delta}}, \pi, \Theta$ or conversions; the reference $\boldsymbol{\varrho}, \boldsymbol{\lambda}$ differ from their committed physical values. Central-only visits and private moves never close epochs.

A clean checkpoint is an actual full public refresh. One read per original label costs

$$C(\log\log n)^2 \sum_e \overline{\boldsymbol{\delta}}_e^{-3} = \ddot{O}(m).$$

Shared field preparation is subquadratic. Finish pending old-view source and clock service, then scan all labels at frozen old prices before resetting accounts. Each read also performs the central inspection of Lemma 8.2: retain its center only on the inner-window test, otherwise install an inner-quarter center. After every label receives both services, zero positive debts, residual targets and due flags, and save baselines for signed endpoint and common positive central clocks. After abort, restore saved physical arrays, certified references, point enclosures and certificates literally, without a new randomized read or central decision; discard or rebuild abandoned caches, transactions and tapes. Thus $C_{\text{initial}} = 0$ in (31).

Since J_t is measurable before frame t's fresh coins, the tower rule, (31), and (32) bound total positive-account reads, marked reads, and proposals by $C\mathbb{E}\sum_t J_t B_t$, times their $\ddot{O}(1)$ point and packed-read factors. Equation (O3) pays the selected reads and, since $\pi_e \ge 1$, the source packet count. Equation (33) pays source assembly. The end refresh pays its own dense scan. For the sparse implementation used here,

$$T\, n d_0\, n^{1/16}\, n^{16/p + o(1)} = n^{1011/512 + o(1)} = o(n^2),$$

and the fine-field, transpose, and source-inverse rows have smaller fixed exponents. This proves the work bound (O4) under the stated sparse-preparation assumptions.

Charge cap, numerical, proxy, inverse and capture failures at their first failed premise. Uniform source references and all-outcome grids justify unconditioned expectations before coupling to the actual prefix; future cap success never conditions signs, masks or no-hit processes. Couple the deterministic attempt-work cap only until its first crossing. Markov's inequality and (O4) bound that probability after choosing its constant, without a separate $m \log n$ confidence allowance per attempt.

10 Initialization and exact output

A hub supplies a feasible interior midpoint and an infeasibility alternative. We initialize tracking, prove polynomial numerical ranges, and convert a terminal state into an integral flow and integer dual. Cold initialization and exact output are each paid once.

Use the input and word model of Section 1. Here n, m are input dimensions, and N, M will denote the augmented working dimensions. An incidence row has -1 at its tail and $+1$ at its head. On the augmented graph, use the core-score proxy construction and fixed sparse inverse of Section 6: their sparse work is $N^{19/16+o(1)}$, and initial dense construction costs $\ddot{O}(M)$.

10.1 Preprocessing and a feasible midpoint

Removing lower bounds and trivial arcs. If some $l_e > \boldsymbol{\upsilon}_e$, that pair certifies infeasibility. Otherwise set

$$x = \mathbf{f} - l, \qquad u^{\text{new}} = u - l, \qquad b^{\text{new}} = b - \mathbf{B}^T l$$

and record the cost offset $\mathbf{c}^T l$. Delete zero-width arcs. Self-loops have zero incidence and are optimized independently at the cost-minimizing endpoint. Check $\sum_v b_v = 0$; a nonzero sum is an exact infeasibility certificate. All transformations take $O(m+n)$ word work, and the transformed data retain a fixed polynomial bound. For the remainder of this subsection, $0 \le x \le u$ denotes the transformed problem and every capacity is a positive integer.

The hub augmentation. In this augmentation calculation, r denotes the midpoint imbalance. Put

$$U_{\text{tot}} = \sum_e u_e, \qquad C_c = \max\{1, \|c\|_\infty\}, \qquad M_{\text{art}} = 2C_c(U_{\text{tot}} + 1) + 1, \qquad r = b - \mathbf{B}^T(u/2).$$

Add a hub h with demand zero. For each original vertex v add the two arcs

$$h \longrightarrow v: \quad u_{hv} = 2(1 + (r_v)_+), \quad c_{hv} = M_{\text{art}}, \qquad v \longrightarrow h: \quad u_{vh} = 2(1 + (-r_v)_+), \quad c_{vh} = M_{\text{art}}.$$

Since r is half-integral and sums to zero, these capacities are positive integers and the augmented midpoint is exactly feasible. It has half-integer coordinates and both box margins at least $1/2$. One incidence pass constructs the $n+1$-vertex, $m+2n$-arc instance with polynomial data.

A dense working graph. Add n zero-demand dummy vertices, all directed arcs between distinct dummy vertices, and the two opposite arcs joining one dummy to the hub. Each has capacity two and cost M_{art}. Their midpoint circulation is one on each arc. The resulting graph has $N \le 2n+1$ vertices, $M = \Theta(n^2)$ arcs, is connected, and has at most one arc for each ordered pair. It adds no route between original vertices. The dummy subgraph, including its hub attachment, carries a circulation in every feasible flow; replacing it by zero preserves all demands and strictly lowers cost unless it was zero. Thus optimum dummy flow is zero.

Lemma 10.1. *An integral optimum of the augmented problem either has zero hub flow and restricts to an original optimum, or yields in $O(M+N)$ work an original set S with*

$$b(S) < -u(\delta^+(S)).$$

Before lower-bound removal this is

$$b^{\text{orig}}(S) < l(\delta^-(S)) - u^{\text{orig}}(\delta^+(S)).$$

Proof. The cycle argument in Lemma 10.2 proves that every nonempty integer-demand box-flow polytope has an integral optimum. An original feasible integral flow has cost at most $C_c U_{\mathrm{tot}}$. Every augmented integral flow with positive hub flow costs at least $M_{\mathrm{art}} - C_c U_{\mathrm{tot}} > C_c U_{\mathrm{tot}}$. Hence original feasibility forces every integral augmented optimum to have zero hub flow.

For the other case, let A_- be the vertices with $\mathbf{f}_{vh} > 0$ and let A_+ be those with $\mathbf{f}_{hv} > 0$. Conservation at the hub and absence of dummy flow make both nonempty. Let S be the set reachable from A_- using residual arcs of the original graph only. It cannot meet A_+: a simple such path, together with the reverses of its two positive hub arcs, would be an augmented residual cycle of cost at most

$$C_c m - 2M_{\mathrm{art}} \le C_c U_{\mathrm{tot}} - 2M_{\mathrm{art}} < 0.$$

The zero-length path handles a vertex in both sets. This contradicts optimality. No original residual arc leaves S, so all original arcs leaving S are saturated and all entering arcs carry zero. Summing balance gives

$$b(S) = -u(\delta^+(S)) - \sum_{v \in S} \mathbf{f}_{vh} < -u(\delta^+(S)).$$

Reachability and checking this strict integer inequality are linear. Substituting $b^{\mathrm{new}} = b^{\mathrm{orig}} - \mathbf{B}^T l$ gives the lower-bound form. The cut is an independently checkable certificate. □

10.2 Cold tracking initialization

A contracting fixed-point map. Use the augmented dimensions and write $r = N-1$, $s = r/M$. At the midpoint, $\mathbf{z}_e = \boldsymbol{\upsilon}_e/(2\sqrt{2})$, so $\mathbf{z}_e^2 = \boldsymbol{\upsilon}_e^2/8$ is rational. For positive weights w, let

$$\mathbf{H}^\star(w) = \mathbf{B}^T \textsc{Diag}(\mathbf{z}^2/\mathbf{w}^2)\mathbf{B}, \qquad \boldsymbol{\tau}(w) = \textsc{Diag}\Big(\textsc{Diag}(\mathbf{z}/\mathbf{w})\mathbf{B}\mathbf{H}^\star(w)^\dagger \mathbf{B}^T \textsc{Diag}(\mathbf{z}/\mathbf{w})\Big).$$

The desired regularized equation is

$$\mathbf{w} = 3s + \boldsymbol{\tau}(w), \qquad \boldsymbol{\lambda} = w - 2s.$$

Define the logarithmic map

$$\Phi(v) = \tfrac{2}{3} v + \tfrac{1}{3} \log(3s + \boldsymbol{\tau}(e^v)).$$

If $\boldsymbol{\Pi}$ is the corresponding edge projection, differentiation gives

$$\frac{\partial \boldsymbol{\tau}_i}{\partial v_j} = -2\mathbf{1}_{\{i=j\}} \boldsymbol{\tau}_i + 2\boldsymbol{\Pi}_{ij}^2.$$

Consequently $D\Phi$ has nonnegative entries and every row sums to $2/3$; its diagonal lies in $[0, 2/3]$. Thus Φ is an ℓ_∞ contraction with factor $2/3$. The box $3s \le w \le 3s+1$ is invariant. Its unique fixed point is at logarithmic distance $O(\log N)$ from $\mathbf{w} = 3s$. Reaching the required fixed multiple of κ/p accuracy takes $O(\log \log N)$ iterations.

Estimating one iteration. At each iteration build a fresh static proxy $P = \sum_\alpha c_\alpha b_\alpha b_\alpha^T$ of the actual rational graph and use its fixed symmetric inverse T_P. Choose dyadics d_α approximating $\sqrt{c_\alpha}$ to the reserved relative accuracy, and prepare $J = \Theta(\log N)$ fields $h_j = T_P \sum_\alpha d_\alpha \sigma_{\alpha j} b_\alpha$, with signs independent across sparse atoms and lanes. Their analytical covariance is $T_P(\sum_\alpha d_\alpha^2 b_\alpha b_\alpha^T)T_P$, which approximates $\mathbf{H}^\star(w)^\dagger$ to the allocated relative accuracy. Finite evaluations and canonical score rounding use the same numerical reserves as the fine-bank construction in Section 9. An all-label squared-distance scan gives relative score accuracy $\eta = c_\star \kappa/p$, with any prescribed polynomial confidence, after allocating a fixed fraction of η to each of the proxy, inverse, and sampling biases. The static core construction gives these first estimates without a prior tracking-weight bound.

The finite update and installed reference. Implement the scalar map with error at most η in log coordinates and clamp the result to $3s \le w \le 3s+2$. The unrounded update is the cube root of $\mathbf{w}_e^2(3s+\widehat{\boldsymbol{\tau}}_e)$. Locate its exponent by comparisons to powers of two, then use fixed-relative-accuracy mantissa comparisons of cubes. Since η is fixed, this costs $O(\log\log N)$ word operations, with fixed-depth rational comparisons. The clamp cannot increase distance to the exact fixed point. For example, store $\mathbf{w} = 3s + q2^{-B}$, so $\boldsymbol{\lambda} = s + q2^{-B}$ and the floor is exact even when $s = r/M$ is not dyadic. These all share the fixed denominator $M2^B$. The ordinary squared coefficients are exactly $\mathbf{z}^2/\mathbf{w}^2$, evaluated as fixed-depth word-sized rationals from the stored state. Sparse packets and accumulated vectors use separate common grids, with errors in the proxy/inverse reserves. Cold ranges follow from midpoint capacities and $3s \le w \le 3s+2$, without later tube bounds.

After $O(\log\log N)$ iterations the distance to the fixed point is as small as any prescribed constant multiple of κ/p. The map $g(v) = \log(3s + \boldsymbol{\tau}(e^v))$ has ℓ_∞ Lipschitz constant at most four, so

$$\|v - g(v)\|_\infty \le 5\|v - v^\star\|_\infty.$$

Moreover $(3s+\boldsymbol{\tau})/(s+\boldsymbol{\tau}) \le 3$, so converting w to $\boldsymbol{\lambda} = w - 2s$ multiplies the relative error by only a fixed constant. Fixing all accuracy constants therefore gives

$$\left|\kappa^{-1}\log\frac{\boldsymbol{\lambda}_e}{s+\boldsymbol{\tau}_e}\right| \le \frac{1}{600p}.$$

In particular the initial tracking potential is at most $2.01M$. Prepare fresh fields and scan all initial references on the final installed graph; the previous iteration's fields do not apply. The dense cost is

$$\ddot{O}(M\log N),$$

plus a strictly subquadratic total of sparse preparations. Scalar fixed-accuracy roots and logarithms fit this budget.

The initial central state. For centrality set $\mathbf{y}_0 = 0$, keep $\mathbf{f}_0 = u/2$, and put

$$M_c = \max_e u_e|\mathbf{c}_e|.$$

Use the same actual central exponent $P \ge 512$, $P = O(\log\log(2N))$, as the phase algorithm. Let μ_0 be the least power of two at least

$$\max\{1, 100P(1+M_c)/(\varepsilon s)\}.$$

At the midpoint the barrier first derivative vanishes, and

$$|\boldsymbol{\varsigma}_{0,e}| = \frac{\mathbf{z}_e|\mathbf{c}_e|}{\varepsilon\mu_0\mathbf{w}_e} \le \frac{M_c}{6\varepsilon\mu_0 s} \le \frac{1}{600P}.$$

Thus the initial central potential is also at most $2.01M$. The residual is exactly zero. Place the central representatives and score references in their interior windows in the paid scan, with all observer debts zero. The factor P in the polynomially bounded μ_0 keeps this potential bound uniform as P grows.

10.3 Polynomial numerical ranges

Bound reduced costs from the feasible midpoint, then capacity margins and condition numbers, all before choosing arithmetic meshes.

Suppose a later state has $0 < \mathbf{f} < \boldsymbol{\upsilon}$, $\boldsymbol{\lambda} \geq s$, and tracking for the actual rational graph:

$$\left|\log \frac{\boldsymbol{\lambda}_e}{s + \boldsymbol{\tau}_e}\right| \leq \kappa, \qquad \mu_{\mathrm{low}} \leq \mu \leq \mu_0,$$

where $\mu_{\mathrm{low}} = N^{-O(1)}$ includes terminal-phase overshoot. Use the barrier quantities of Section 2, with A_{st}, u in place of $B, \boldsymbol{\upsilon}^{\max}$, and write $\mathbf{D}_0 = \mathbf{D}$ and $\boldsymbol{\Xi}_0 = \mathbf{z}\mathbf{c}^y/(\mu\mathbf{w}) = \overline{\boldsymbol{\Xi}} + \varepsilon\boldsymbol{\varsigma}$, where $\mathbf{c}^y = \mathbf{c} - \mathbf{B}\mathbf{y}$, $|\boldsymbol{\varsigma}_e| \leq 1$, and $\mathbf{w} = 2s + \boldsymbol{\lambda}$. Define

$$\mathbf{k} = \mathbf{D}_0\boldsymbol{\Xi}_0 + (3\boldsymbol{\lambda}/\mathbf{w})(1 - \mathbf{D}_0\overline{\boldsymbol{\Xi}}), \qquad W_e = \frac{3\boldsymbol{\lambda}_e\mathbf{z}_e^2}{\mu\mathbf{w}_e^2\mathbf{k}_e},$$
$$\mathbf{W} = \textsc{Diag}(W_e), \qquad \mathbf{L} = \mathbf{B}^T W\mathbf{B}, \qquad \rho = \mathbf{B}^T\mathbf{f} - b.$$

Assume $\|\rho\|_{\mathbf{L}^\dagger} \leq \beta\sqrt{\mu}$ with $\beta \leq 1$. Since $|\mathbf{D}_0|, |\overline{\boldsymbol{\Xi}}| \leq 1$, $0 \leq \mathbf{D}_0\overline{\boldsymbol{\Xi}} \leq 1$, and $1 \leq 3\boldsymbol{\lambda}/\mathbf{w} \leq 3$, we have $1 - \varepsilon \leq \mathbf{k} \leq 3 + \varepsilon$. Tracking gives $\boldsymbol{\lambda}_e \leq C$, $\sum_e \boldsymbol{\lambda}_e = O(r)$, $3s \leq \mathbf{w}_e \leq C$, and $\sum_e w_e = O(r)$.

Reduced costs. Define the nonnegative box expression

$$G(\mathbf{f}, \mathbf{y}) = \sum_e \left[\mathbf{f}_e(\mathbf{c}_e^y)_+ + (\boldsymbol{\upsilon}_e - \mathbf{f}_e)(-\mathbf{c}_e^y)_+\right].$$

Centrality alone gives $G(\mathbf{f}, \mathbf{y}) \leq C\mu r$, without feasibility. For instance, when $\mathbf{c}_e^y \geq 0$, writing $t = \mathbf{f}_e/(\boldsymbol{\upsilon}_e - \mathbf{f}_e)$,

$$\frac{\mathbf{f}_e c_e^y}{\mu\mathbf{w}_e} = 1 - \boldsymbol{\varpi} + \varepsilon\boldsymbol{\varsigma}_e\sqrt{1 + \boldsymbol{\varpi}^2} \leq 1 + \varepsilon.$$

Set $S_c = \sum_e \boldsymbol{\upsilon}_e|\mathbf{c}_e|$. The estimates $\mathbf{z}_e \leq \boldsymbol{\upsilon}_e/2$ and $\mathbf{w}_e \geq 3s$ give

$$W_e \leq C\boldsymbol{\upsilon}_e^2/(\mu s), \qquad W_e(\mathbf{c}_e^y)^2 = 3\mu\boldsymbol{\lambda}_e\boldsymbol{\Xi}_{0,e}^2/\mathbf{k}_e.$$

Therefore, using quotient Cauchy–Schwarz before any lower eigenvalue bound,

$$\|\mathbf{y}\|_{\mathbf{L}} \leq C\left(S_c/\sqrt{\mu s} + \sqrt{\mu r}\right), \qquad |\mathbf{y}^T\rho| \leq C\beta(S_c/\sqrt{s} + \mu\sqrt{r}).$$

The exactly feasible midpoint gives the identity

$$\frac{1}{2}\sum_e \boldsymbol{\upsilon}_e|\mathbf{c}_e^y| = G(\boldsymbol{\upsilon}/2, \mathbf{y}) = G(\mathbf{f}, \mathbf{y}) + \mathbf{c}^T(\boldsymbol{\upsilon}/2 - \mathbf{f}) + \mathbf{y}^T\rho.$$

For a fixed sufficiently large C_B, put

$$B_c = 1 + S_c + C_B\mu r + C_B\beta(s^{-1}S_c + \mu r).$$

Then $\|\mathbf{c}^y\|_\infty \leq 2B_c$, with a fixed polynomial upper bound independent of subsequent arithmetic meshes.

Capacity margins and condition numbers. If $\mathbf{f}_e \le 1/4$, then $\boldsymbol{\upsilon}_e \ge 1$ implies $\overline{\boldsymbol{\Xi}}_e \ge 2/\sqrt{10}$, hence $|\boldsymbol{\Xi}_{0,e}| \ge 1/2$. Using $\mathbf{z}_e \le \mathbf{f}_e$ gives

$$\mathbf{f}_e \ge \mu \mathbf{w}_e/(2|\mathbf{c}_e^y|) \ge 3\mu s/(4B_c).$$

The upper endpoint is symmetric, so

$$\min\{\mathbf{f}_e, \boldsymbol{\upsilon}_e - \mathbf{f}_e\} \ge \min\{1/4, \mu s/(4B_c)\} = N^{-O(1)}. \tag{34}$$

Ground the hub. Paths and $|A_e y| \le |\mathbf{c}_e| + |\mathbf{c}_e^y|$ bound every stored potential coordinate by a polynomial. All positive barrier and structural coefficients now have polynomial lower and upper bounds. For a connected positive coefficient graph with minimum weight a_-, path Cauchy–Schwarz gives the grounded eigenvalue lower bound a_-/N^2; twice the total coefficient mass is an upper bound. Thus all required grounded condition numbers are polynomial. These bounds determine grids for rational coefficients, sparse packets, and accumulated vectors without assuming their accuracy in advance.

Finally the potential caps imply $\delta^{-p}, (1-|\boldsymbol{\varsigma}|)^{-P} \le N^{O(1)}$. Any fixed number of differentiations adds only fixed powers of inverse margins, p, and $P = O(\log\log N)$. The resulting polynomial derivative bounds fix word and grid exponents before implementation.

10.4 An exactly feasible fractional endpoint

The interior algorithm permits a small incidence residual. We remove it while preserving the box constraints and a small duality gap: first bound the ideal feasible projection, then approximate it and correct the remaining residual exactly on a tree.

For any $\mathbf{y}$, the quantity

$$\mathcal{D}(\mathbf{y}) = \mathbf{b}^T\mathbf{y} + \sum_e \boldsymbol{\upsilon}_e \min\{0, \mathbf{c}_e - (\mathbf{B}\mathbf{y})_e\}$$

is a dual lower bound. For exactly feasible x,

$$\mathbf{c}^T x - \mathcal{D}(\mathbf{y}) = \sum_e \left[x_e(\mathbf{c}_e^y)_+ + (\boldsymbol{\upsilon}_e - x_e)(-\mathbf{c}_e^y)_+\right]. \tag{35}$$

The analytical feasible projection. Consider the ideal projection

$$\widehat{\mathbf{f}} = \mathbf{f} - \mathbf{W}\mathbf{B}\mathbf{L}^\dagger \rho.$$

Let $\ell_e^{\mathrm{phys}} = W_e A_e \mathbf{L}^\dagger \mathbf{B}_e^T$. The ordinary graph is exactly $\mathbf{H}^\star = \mathbf{B}^T \textsc{Diag}(\mathbf{z}^2/\mathbf{w}^2)\mathbf{B}$, so the tracking bound applies directly in this metric. Since $\mu\mathbf{L} \succeq csH^\star$, inverse order and the trivial bound $\ell_e^{\mathrm{phys}} \le 1$ imply

$$\ell_e^{\mathrm{phys}}/\boldsymbol{\lambda}_e \le C\min\{1/\boldsymbol{\lambda}_e, \boldsymbol{\lambda}_e/s\} \le Cs^{-1/2}.$$

Row Cauchy–Schwarz against the residual tube then gives

$$\|(\widehat{\mathbf{f}} - \mathbf{f})/\mathbf{z}\|_\infty \le C\beta s^{-1/4}, \qquad \|(\widehat{\mathbf{f}} - \mathbf{f})/\mathbf{z}\|_{\mathbf{w}} \le C\beta.$$

At the terminal trajectory state the residual invariant of Lemma 2.4 allows $\beta = \beta_s$. The physical step scale then makes the first bound at most $1/10$. Consequently $\widehat{\mathbf{f}}$ is in the box with a known inverse-polynomial margin. Since $\mathbf{z}\mathbf{c}^y = \mu\mathbf{w}C_0$,

$$0 \le \mathbf{c}^T\widehat{\mathbf{f}} - \mathcal{D}(\mathbf{y}) \le G(\mathbf{f}, \mathbf{y}) + C\mu\beta\sqrt{r} \le C_{\mathrm{gap}}\mu r.$$

Choose μ_{end} so that $C_{\mathrm{gap}}\mu_{\mathrm{end}} r \le 1/(64M)$. It is inverse polynomial and $\log(\mu_0/\mu_{\mathrm{end}}) = O(\log N)$.

Computing the projection. Freeze the exact rational weights W_e at the terminal stored state, using the square-root cancellations in the proof of Lemma 3.2. The physical owner gives a sparse constant proxy of $\mathbf{L}$. After grounding the hub, its fixed sparse inverse is scaled to obey

$$a\mathbf{L}_g^{-1} \preceq T \preceq \mathbf{L}_g^{-1}$$

for an absolute $a > 0$. Starting from zero, use

$$v_{j+1} = v_j + T(\rho_g - \mathbf{L}_g v_j).$$

This fixed-metric iteration contracts exact energy error by $1 - a$. Each true product is one $O(M)$ original-label pass; each inverse application has the separately supplied sparse price.

Finite arithmetic and a residual certificate. Use the preceding range and projection bounds to choose known polynomial bounds $Q \geq \max\{1, \max_e W_e\}$, $w_- \leq \min_e W_e$, and a known margin $\sigma \leq \min_e\{\widehat{\mathbf{f}}_e, \boldsymbol{\upsilon}_e - \widehat{\mathbf{f}}_e\}$. Put $\mathbf{K}_0 = 2MQ$, $C_1 = 1 + \sum_e |\mathbf{c}_e|$. Choose a positive dyadic $\eta = N^{-O(1)}$ satisfying

$$(1 + 2M)\sqrt{Q}\,\eta \leq \min\{\sigma/4, 1/(64MC_1)\}. \tag{E1}$$

These bounds reserve the box and duality-gap margins. Choose the dyadic scales by squared rational comparisons or predetermined polynomial bounds, and the smaller iteration tolerance $\boldsymbol{\nu}$ so that

$$0 < \nu_{\mathrm{it}} \leq \min\{\eta/8, \eta\sqrt{w_-}/(16N\sqrt{\mathbf{K}_0})\}.$$

For each product round the fixed-depth rational $W_e(v_{\mathrm{head}} - v_{\mathrm{tail}})$ once to a common grid of mesh δ_{mult}, then add opposite endpoint integers. If d_v is the incidence count, the error at v is at most $\epsilon_v = d_v\delta_{\mathrm{mult}}/2$; the full error vector has Euclidean norm at most $M\delta_{\mathrm{mult}}$. Choose

$$NM\delta_{\mathrm{mult}}/\sqrt{w_-} \leq a\nu_{\mathrm{it}}/32, \qquad M\delta_{\mathrm{mult}} \leq \eta\sqrt{w_-}/(32N).$$

Round the sparse inverse application and stored addition with another local energy error at most $a\nu_{\mathrm{it}}/32$. The actual iteration has

$$E_{j+1} \leq (1 - a)E_j + a\nu_{\mathrm{it}}/16, \qquad E_j = \|v_j - \mathbf{L}_g^{-1}\rho_g\|_{\mathbf{L}_g}.$$

The polynomial initial energy and inverse-polynomial $\boldsymbol{\nu}$ imply $E_j \leq \boldsymbol{\nu}$ after a predetermined $O(\log N)$ steps. Every dense edge uses a fixed number of word operations and one rounding; unrelated rational denominators are never summed.

One final dense pass computes a rounded value $\widetilde{r}$ of $\rho_g - \mathcal{L}_g v$ and the exact common-grid sum

$$R_2 = \sum_{v \neq h} (|\widetilde{r}_v| + \epsilon_v)^2.$$

Accept the solve only if

$$N^2 R_2 \leq w_-\eta^2/4.$$

Because $\mathbf{L}_g \succeq w_- N^{-2} I$, this certifies energy error at most $\eta/2$, independent of whether the sparse primitive was successful. On its success event,

$$\sqrt{R_2} \leq \sqrt{\mathbf{K}_0}\boldsymbol{\nu} + 2M\delta_{\mathrm{mult}} \leq \eta\sqrt{w_-}/(8N),$$

so the check passes with reserve. If the check fails, report terminal failure; do not repeat the dense solve.

Exact correction on a tree. Form $\mathbf{f} - WAv$ by one dense pass and round to a common terminal primal grid with coordinate error at most $\sqrt{Q}\,\eta/2$. The result $\widetilde{\mathbf{f}}$ satisfies

$$\|\widetilde{\mathbf{f}} - \widehat{\mathbf{f}}\|_\infty \le \sqrt{Q}\,\eta, \qquad \|\mathbf{B}^T\widetilde{\mathbf{f}} - b\|_1 \le 2M\sqrt{Q}\,\eta.$$

For $q = b - \mathbf{B}^T\widetilde{\mathbf{f}}$, take any spanning tree and send each rooted subtree total $q(T_v)$ through its parent edge, with the appropriate sign. This constructs d with

$$\mathbf{B}^T d = q, \qquad \|d\|_\infty \le \|q\|_1/2$$

in $O(M+N)$ exact common-grid work. Set $x = \widetilde{\mathbf{f}} + d$. Equation (E1) preserves the box and gives

$$\|x - \widehat{\mathbf{f}}\|_\infty \le (1+2M)\sqrt{Q}\,\eta, \qquad 0 \le \mathbf{c}^T x - \mathcal{D}(\mathbf{y}) < 1/(8M).$$

Balances, bounds, and this gap are checked exactly. The complete terminal projection costs $O(M \log N)$ plus a strictly subquadratic sparse row.

10.5 Integral rounding and an integer dual

Cost-nonincreasing rounding preserves the endpoint's gap bound: $\mathbf{c}^T\mathbf{f} - \mathcal{D}(\mathbf{y}) \le \mathbf{c}^T x - \mathcal{D}(\mathbf{y}) < 1/(8M)$. This meets the $1/(4M)$ hypothesis of Lemma 10.3, which supplies the integer dual.

Lemma 10.2. *From an exactly feasible common-grid flow x with integer demands and bounds, an integral feasible flow $\mathbf{f}$ with $\mathbf{c}^T\mathbf{f} \le \mathbf{c}^T x$ can be found in $O((M+N)\log N)$ word operations. Each coordinate remains between its initial floor and ceiling.*

Proof. The fractional-edge subgraph cannot be a nonempty forest: a leaf would have exactly one fractional incident coordinate, contradicting its integer balance. Along a fractional undirected cycle choose the circulation direction of nonpositive cost, and move until a coordinate hits its initial floor or ceiling. This preserves feasibility and cost nonincrease while permanently removing a fractional coordinate. The finite argument also proves integral optimum existence from any real optimum in the compact box.

Insert fractional edges into a dynamic forest. A within-tree insertion exposes a cycle; signed path-cost sums choose its direction and directional path minima its step. Store the signed cost and minimum distances to initial floors and ceilings. Augmentation subtracts from one distance and adds to the other; reversal negates cost and swaps the aggregates. These are constant-size dynamic-tree lazy updates. Cut every newly integral forest edge. To handle ties, after a cut recurse on the two surviving endpoint subpaths, discarding a subpath with positive minimum. Each removal creates at most two further queries; every edge is removed permanently. If the inserted edge remains fractional, at least one forest edge has become integral. Cut all tied zero-distance edges before linking the inserted edge; these cuts break the old path. Each fractional edge is inserted at most once and each forest edge is cut at most once, giving $O(M)$ dynamic-tree operations. All flows stay on the initial common grid. □

The dynamic-tree operations above cost $O(\log N)$ amortized time [ST83, Sections 2–4]; see also Kang–Payor, *Flow Rounding*, Sections 2–4. This is one terminal $O(M \log N)$ call, storing each capacity once.

Lemma 10.3. *If $\mathbf{f}$ is integral feasible and $0 \le \mathbf{c}^T\mathbf{f} - \mathcal{D}(\mathbf{y}) < 1/(4M)$, an integer potential p satisfying all residual reduced-cost inequalities can be found in $O(M + N\log N)$ word operations.*

Proof. Let $g_{\text{gap}} = \mathbf{c}^T\mathbf{f} - \mathcal{D}(\mathbf{y})$. Every residual arc $a \to b$ of integer cost w_{ab} has positive integer residual capacity, so (35) gives $\mathbf{y}_b - \mathbf{y}_a \le w_{ab} + g_{\text{gap}}$. For $0 \le \alpha < 1$, put $p_v(\alpha) = \lfloor \mathbf{y}_v + \alpha \rfloor$. If $\mathbf{y}_b - \mathbf{y}_a \le w_{ab}$, its residual inequality holds for every shift. Otherwise the violating shifts occupy measure $\mathbf{y}_b - \mathbf{y}_a - w_{ab} \le g_{\text{gap}}$. There are at most $2M$ residual arcs, so their forbidden union has measure below $1/2$.

Initialize $p = \lfloor \mathbf{y} \rfloor$ and count violations in one scan. If zero, return immediately. Otherwise sort the distinct events $1 - \{\mathbf{y}_v\}$ strictly between zero and one. At each event increment all tied vertices, updating the violation count through their incident residual arcs. Inspect only after the complete tie group. These tests cover every interval between events, so the measure bound guarantees a zero count. Each residual arc is updated at most twice.

The returned inequalities are

$$\mathbf{c}_e - (\mathbf{B}p)_e \ge 0 \quad (\mathbf{f}_e < \boldsymbol{\upsilon}_e), \qquad \mathbf{c}_e - (\mathbf{B}p)_e \le 0 \quad (\mathbf{f}_e > 0).$$

They make every summand of (35) vanish, so $\mathbf{c}^T\mathbf{f} = \mathcal{D}(p)$. This exactly certifies optimality. □

All fractional-part comparisons use represented grid integers. The final primal and dual scans are exact. If the augmented optimum has positive hub flow, use Lemma 10.1; otherwise restrict the flow and potential and restore lower bounds, fixed arcs, and optimized loops. For the original lower-bounded instance the checked dual value is

$$\mathbf{b}^T p + \sum_e \left[l_e \max\{0, \mathbf{c}_e - (\mathbf{B}p)_e\} + \boldsymbol{\upsilon}_e \min\{0, \mathbf{c}_e - (\mathbf{B}p)_e\} \right].$$

Thus either output has a certificate checkable in $O(m+n)$ work.

10.6 Matching and the bounded-size branch

The proof of Corollary 1.2 in Section 16 gives the matching reduction. Below the fixed asymptotic cutoff, enumerate the augmented integer box and choose a feasible minimum-cost assignment; the midpoint and cycle argument guarantee existence. Bellman–Ford from a zero-distance source in its optimal residual graph gives integer potentials. Apply the hub alternative or restrict and check the solution. For fixed cutoff and input exponent this finite work is absorbed by the constant. Empty and one-vertex inputs are handled directly.

10.7 Confidence and one-time work caps

Lemma 10.4. *For any fixed input-range and confidence exponents, the cold and terminal procedures can each be run once with a deterministic cap*

$$\ddot{O}(M \log N) + N^{19/16+o(1)},$$

whose probability of firing on a valid endpoint input is at most any prescribed inverse polynomial. The degree in the dense iterated-logarithm factor is absolute. A cap or failed certificate returns failure and does not restart either dense endpoint.

Proof. Cold contraction uses a predetermined number $\ddot{O}(1)$ of iterations, each with $J = C \log n$ lanes and one all-label squared-distance scan. Including the separate final installed-graph scan, these dense operations cost $\ddot{O}(M \log n)$ on every tape. Terminal refinement uses a predetermined number $O(\log n)$ of true matrix products; its final exact residual test passes or reports failure.

Tree correction, integral cycle rounding, integer dual extraction, and exact output checks cost $O(M \log n + N \log n)$. Sparse-inverse operand guards apply on every tape.

All random preparatory work lies in $S(N) = N^{19/16+o(1)} n^{o(1)}$, including the cold calls, amplitude buckets, and terminal factor applications. Make the stated child, replay, residual-phase, sparse-copy, factor-entry, and sample-count caps literal. For a Bernoulli or categorical-position batch in the sampling reference with public conditional mean count at most V_*, scalar Chernoff gives a cap $C(V_* + (c_0 + A_{\mathrm{st}} + 10) \log n)$ with failure probability at most $N^{-c_0 - A_{\mathrm{st}} - 10}$, where $N^{A_{\mathrm{st}}}$ is the fixed call horizon. This also counts selected root positions that are subsequently discarded. The structural finite-geometric primitive of Appendix A.5.1 has its separately allocated coupling failure and $O(\log^2 n)$ work per selected or exhausted position; stream starts are included. Apply structural work bounds only before the first failed structural premise. On their simultaneous correctness event the stated scan caps cannot be reached, so no conditioning on future success is needed.

After these caps there are at most a deterministic $K_* = O(S(N))$ logical uniform-integer and eligible-pivot choices. Integer rejection succeeds with conditional probability at least $1/2$; combining it with the low-degree eligibility test gives at least $1/4$. Their conditional geometric moment bound implies total elementary trials $O(K_* + (c_0 + 10) \log n)$ except with probability $N^{-c_0 - 10}$. The conditional moment bound applies to adaptive histories.

For geometric gaps, use the finite endpoint construction, with prefix constant $\mathbf{B}$, power mesh ζ, and remaining length d. Each completed gap has its ideal value on the stated coupling, costs $O(\log^2 n)$ deterministically, and has conditional failure probability at most $C(d 2^{-B \log n} + d^2 \zeta)$ for $d \le N^{A_{\mathrm{st}}}$. The chosen precision makes the union over at most $N^{A_{\mathrm{st}}}$ calls smaller than the allocated inverse-polynomial reserve. Report failure at the first ambiguous comparison, without redrawing a gap or restarting the dense endpoint.

Every gap belongs to sparse proxy or preparation work. With the preceding stream-start and sample-count caps, there are $O(S(N))$ gaps: each stream uses at most its raw activation count plus one. Their $O(\log^2 n)$ price is therefore absorbed in $S(N)$'s unspecified logarithmic factor. Dense label scans use no geometric gaps. The scalar table of Lemma 12.1 has one-time deterministic $N^{1/8} n^{o(1)}$ work.

Union conditional failures over the public horizon, fixing each call at its valid input before fresh coins, including the adaptive terminal checkpoint. Confidence changes constants, not dense iteration counts. The sparse row is $o(N^2)$; the bounded-input branch covers the exceptions. □

11 Clean phases and restarts

Here n, m, r again denote the working graph's vertex count, label count, and rank, written $N, M, N-1$ in Section 10.

A *frame* comprises a physical update and its reference/metadata service. A *clean checkpoint* saves the physical state and refreshed references, with zero new observation debt. An *attempt* tries to complete T frames from that checkpoint. Success installs a new checkpoint and completes a *phase*; failure restores the old one.

Each frame draws one independent original-label mask. A failed residual, source-normalizer, or work test aborts the attempt without redraw, so the reference laws of Section 2 and 3 apply at every valid pre-mask state.

11.1 The four implementation contracts

In addition to local drift, use the following contracts, verified for this implementation by Lemma 14.1.

1. *Potential certificates.* At a returned valid state, representatives and class counts give computable nonnegative certificates U_R, U_Z satisfying

$$\Psi_R \le U_R \le 2\Psi_R, \qquad \Psi_Z \le U_Z \le 2\Psi_Z.$$

   Charge their preparation without an all-edge leverage calculation.

2. *Clean checkpoints.* At a returned state in the declared outer ranges, public refresh and rebuild create a clean checkpoint for $\ddot{O}(m) + S_{\mathrm{cp}}(n)$ work, where $S_{\mathrm{cp}}(n) = o(m)$. Refresh every original-label observer reference before zeroing new debt accounts, and save the physical state and certified references. Restore all mutable state consistently at this same price.

3. *Expected work.* For an attempt stopped at its first invalid primitive premise, valid completed outer exit, or prospective genuine-work ledger crossing, its uncapped reference obeys the nonnegative accounting bound

$$\mathbb{E}W \le C\ddot{O}(1)\left[m + \mathbb{E}\sum_t J_t B_{*,t}\right] + S(n), \tag{36}$$

   where J_t says that frame t begins, and $S(n) \le n^{2-\epsilon+o(1)}$ for fixed $\epsilon > 0$. Include original-label proposal evaluations, selected updates, observer credits and payments, scalar arithmetic, inverse/field preparations, list starts, and rejected-endpoint refreshes. Freeze old-owner prices before generating each batch. Observer payments may use earlier frames' credits. The prospective genuine ledger has the same expected bound.

4. *Finite execution and controlled failures.* A fresh construction has inverse-polynomial failure probability given the valid history just before its tapes. A public observer epoch has an inverse-polynomial probability of unreported residual exit given its valid zero-target birth, for adaptive later fields and stopped at an earlier primitive failure. This bound is not conditional on intermediate misses. Union at call starts and epoch births under the predeclared polynomial horizon, from each clean checkpoint's full epoch reset. Domain/operand guards, iterative caps, and actual work counting ensure finite execution on every outcome. Hidden matrix or projection failures need not be detectable; couple to the valid reference only until its first failure.

11.2 Stopped potential and work budgets

Recall from Section 2 that $\beta_s = \gamma\beta_f$ and $B_* = \beta_f m/\sqrt{r}$. Let ϵ_{tot} bound the sum of the local recurrences' additive ϵ_{num} terms over one attempt; the parameter hierarchy makes this reserve inverse polynomial.

Fix inner constants $C_{R,\mathrm{in}}, C_{Z,\mathrm{in}}$ and start with

$$U_R \le C_{R,\mathrm{in}} m, \qquad U_Z \le C_{Z,\mathrm{in}} P^2 m.$$

Choose a fixed phase constant C_T and length

$$T = C_T\sqrt{r}/\beta_s = O\left(n^{3/4+3/p}P^{16}\right).$$

Constants may depend on the fixed γ. The attempt length is fixed independently of the potential.

We use two proof processes. The *killed process* assigns zero potential after an abort or primitive failure and controls endpoint potentials. The *outer-exit process* instead retains the value at a valid

outer crossing, so the crossing remains visible in its expectation. These processes are analytical devices; neither changes the algorithm.

If a stop occurs inside a begun frame, complete that frame only in the uncapped mathematical reference before killing its potential, or assign zero if a primitive failure occurs first. Start from the nonnegative matrix-killed tracking candidate and all-mask central candidate; later observer, epoch, or operational failures only decrease them. The local tracking bound already includes its same-frame signed-error reserve. A pre-mask preparation abort also sets the candidate to zero. The fixed small-step choices give $a_R, a_Z \in [0,1]$, so the affine drift bounds have nonnegative right-hand sides and remain valid at that abort. An abort during source preparation still charges its header work to that frame's B_* and sparse row.

Lemma 11.1. *For any attempt stopped by the preceding rules, before imposing its actual-work cap,*

$$\sum_t J_t B_{*,t} \le C_\gamma C_T m.$$

Thus the work contract gives $\mathbb{E}W = \ddot{O}(m) + S(n)$.

Proof. The frame-start indicator J_t is predictable. The first bound is deterministic: $TB_* = O(C_T\sqrt{r}/(\gamma\beta_f))(\beta_f m/\sqrt{r}) = O_\gamma(C_T m)$. Apply (36). □

11.3 An attempt and its probability of success

Use larger fixed outer constants $C_{R,\mathrm{out}}, C_{Z,\mathrm{out}}$. After each completed valid frame, abort if

$$U_R > C_{R,\mathrm{out}} m \quad \text{or} \quad U_Z > C_{Z,\mathrm{out}} P^2 m.$$

Also abort at a failed computable residual or Γ test, a numerical/domain guard, a prospective genuine-ledger cap, or the actual work cap. At a surviving length-T endpoint, perform the paid public refresh before testing the inner certificates. Accept the checkpoint only if both inner tests pass. A failed endpoint still pays for its refresh.

Lemma 11.2. *The outer constants can be chosen so that an outer exit has probability at most* $1/20$*, in addition to the separately budgeted primitive first-failure probability.*

Proof. Set $V_R = \Psi_R/m$, $V_Z = \Psi_Z/(P^2 m)$, and $V = V_R + V_Z$. The two local drifts bound its conditional expected increase by $C_\gamma \beta_f/\sqrt{r}$ per valid frame. Summing through T gives a fixed constant depending on γ, C_T.

Kill the outer-exit process at a primitive first failure before any undefined barrier evaluation. Freeze its first *valid completed outer crossing* at the crossing value. An earlier operational abort also freezes it, but receives no exit witness. For a midframe abort, complete the mathematical frame before freezing, or assign zero if a primitive failure occurs first. Freezing adds no drift. Once a valid crossing is frozen, a later virtual-continuation failure cannot erase its witness. Killing a nonnegative candidate adds no positive drift, with the local same-frame signed-error reserve included. This remains true when a public-epoch failure has large conditional probability at the current prefix.

A crossing $U_R > C_{R,\mathrm{out}} m$ retains the witness $V \ge V_R > C_{R,\mathrm{out}}/2$; the central crossing is analogous. The expected final outer-exit value is bounded by $V(0) + C_\gamma T\beta_f/\sqrt{r} + \epsilon_{\mathrm{tot}}$. Choose both outer constants sufficiently large and apply Markov to this retained crossing value. □

Lemma 11.3. *The inner constants and C_T can be chosen so that the probability of a surviving endpoint failing either inner test is at most* $1/20$.

Proof. Use the killed process at every earlier abort. Its nonnegative potentials obey the affine recurrences, since killing only decreases them. For tracking,

$$\mathbb{E}\Psi_R(T) \le e^{-a_R T}\Psi_R(0) + Cm + \epsilon_{\mathrm{tot}}, \qquad a_R T = \Theta(C_T/\gamma).$$

For centrality,

$$\mathbb{E}\Psi_Z(T) \le e^{-a_Z T}\Psi_Z(0) + C_\gamma P m + \epsilon_{\mathrm{tot}}, \qquad a_Z T = \Theta(C_T P).$$

Choose $C_{R,\mathrm{in}}$ large relative to the first equilibrium constant, and $C_{Z,\mathrm{in}}$ large relative to the second constant divided by the uniform lower bound $P \ge 512$. Then choose C_T so both contraction factors are small enough. For example both expected endpoint potentials can be made less than one hundredth of their respective inner thresholds.

The refresh preserves physical state and true potentials. At a surviving endpoint, a rejected upper certificate puts the true potential above half its inner threshold. Markov at that threshold proves the claim without conditioning on survival or changing the within-attempt law. □

Last-frame service. Prepare scalar ranges, field enclosures, and final metadata service on a fixed one-step expansion of the outer range. The tracking estimate gives $\delta'_e \ge (1-\epsilon_0)\boldsymbol{\delta}_e$ on a good frame, where ϵ_0 is a sufficiently small fixed constant; therefore an expanded tracking cap is $(1-\epsilon_0)^{-p}C_{R,\mathrm{out}}m$. For centrality, the wall and source caps give, with a fixed small movement constant c_{mov},

$$\frac{|\boldsymbol{\varsigma}'_e - \boldsymbol{\varsigma}_e|}{1-|\boldsymbol{\varsigma}_e|} \le c_{\mathrm{mov}}/P^2, \qquad \psi_P(\boldsymbol{\varsigma}'_e) \le (1-c_{\mathrm{mov}}/P^2)^{-P}\psi_P(\boldsymbol{\varsigma}_e) \le e^{2c_{\mathrm{mov}}/P}\psi_P(\boldsymbol{\varsigma}_e).$$

This expansion is uniform in growing P. A frame that crosses an outer certificate is serviced within these expanded ranges and is never used to begin another valid frame.

11.4 Operational caps and their reference runs

Prospective ledger and token caps. Lemma 11.1 and the work contract bound the expected genuine ledger by $\ddot{O}(m)$ per attempt. Use the changed batch's prices frozen before the mask and test the prospective cumulative ledger before pruning, retirement, or carrying. Include the crossing batch in the stopped expectation; Markov then gives crossing probability at most $1/100$ for a sufficiently large fixed cap.

Choose a separate sharp token cap $N_{\mathrm{sharp}} = \ddot{O}(m)$ after proving the weighted-rate expectation under the coarse admission cap. The current moments proved under that coarse cap price the batch. Here N counts admission tokens and Z is the genuine-event ledger at normalized frozen old prices. Lemma 5.3 gives $N \le N_{\mathrm{init}} + 81Z$. Before commit and copying, test $N_{\mathrm{init}} + 81Z_{\mathrm{pros}} \le N_{\mathrm{sharp}}$, retaining the per-token guard without needing future export counts. The coarse reference includes the whole proposed crossing batch. Choose the cap constant after bounding $\mathbb{E}(N_{\mathrm{init}} + 81Z)$; Markov makes a failed test a small constant-probability attempt failure.

With the owner's threshold $T_0 \asymp m/\log^{1024} n$, only after imposing this sharp cap use the deterministic bound

$$1 + \log\left(1 + \frac{N_{\mathrm{sharp}}}{T_0}\right) = \ddot{O}(1),$$

since $P = O(\log\log n)$. This pathwise depth pays copying; $\mathbb{E}N = \ddot{O}(m)$ alone would not bound $\mathbb{E}[N\log(1+N/T_0)]$.

The actual-work cap. Fix the mean-work constants before the cap $\ddot{O}(m)$. The following coupling includes exact random sampling and bounds exhaustion by $1/100$. A cap may interrupt a call but never redraws an individual query.

For the coupling proof, use a coarse reference $\mathcal{R}_0$ with only its predeclared polynomial horizon and ordinary mathematical stops, and a reference $\mathcal{R}_1$ that additionally applies the sharp ledger and token guards but omits the actual-work cap. Their primitive-failure event $\mathcal{B}$ includes a first failure in the same-tape virtual continuation of $\mathcal{R}_0$, even after the real execution has stopped. Its probability is bounded at valid call starts and epoch births, not conditional on an expensive partial stream. If $\mathcal{A}_W$ is real work-cap exhaustion at w_*, instruction-by-instruction coupling gives

$$\mathcal{A}_W \cap \mathcal{B}^c \subseteq \{W_{\mathcal{R}_1} \geq w_*\}, \qquad \Pr(\mathcal{A}_W) \leq \mathbb{E}W_{\mathcal{R}_1}/w_* + \Pr(\mathcal{B}).$$

This includes a stop during an exact gap call; the reference continues the operation on the same logical tape.

Independent virtual tails. Give attempts and calls independent logical tapes, revealed as used. After a midcall stop, complete only the virtual history with an independent auxiliary suffix, reserving the unread physical stream for the next real attempt. This preserves the reference law and the next attempt's independence from that history. The suffix never controls the algorithm and requires no executed computation or random access: the machine requests fresh independent words and discards abandoned cursors.

Random source pads across restarts. Draw exactly cancelling padding independently of this attempt's uncapped mathematical history. In the reference without the actual-work cap, mathematical stop tests and the prescribed source-finalization schedule do not inspect padding cost. Integrate the conditional carry bound over that history. Each new attempt conditions on its checkpoint and draws fresh padding; independence from the entire restarted trajectory is unnecessary.

11.5 Amplification and running time

Theorem 11.4. *Assume the local drift bounds and the four contracts above. Assume a clean inner initialization and a terminal extraction at $\mu \leq \mu_{\mathrm{end}}$ with polynomial scale ratio μ_0/μ_{end}. Then the clean-phase algorithm makes the required progress with any prescribed polynomial confidence in*

$$O\big(m \log n \ddot{O}(1)\big) + n^{2-\epsilon+o(1)} = O\big(n^2 \log n \ddot{O}(1)\big)$$

work, for fixed positive ϵ and $m = \Theta(n^2)$.

Proof. Allocate a small fixed probability to the two potential failures of Lemmas 11.2 and 11.3, the genuine-ledger, sharp-token and actual-work caps, and all primitive failures. For the last group, condition on the clean checkpoint and sum fresh-call failure bounds at their call times and no-hit failure bounds at their adaptively chosen epoch births. Each valid birth starts at zero residual target, and there are only polynomially many births. The epoch supermartingale bound already covers its entire adaptive future, so no per-frame conditional no-hit bound is used. By the preceding lemmas, each attempt has success probability at least a fixed $q_0 > 0$, conditional on any valid clean checkpoint and the entire preceding history. Every successful attempt makes between two fixed positive constants of progress in $-\log \mu$, because it contains T frames and each actual ratio obeys $1 - \theta = \Theta(\beta_s/\sqrt{r})$.

Restore physical state exactly with its saved certified references. Rebuild owners and mutable metadata, discarding failed queues, due flags, private caches, buffered transactions, and tape cursors. Each restored observer epoch has its refreshed reference and zero new debt; the next attempt uses fresh mask and source tapes. The checkpoint contract pays this restoration.

A polynomial initial-to-terminal scale ratio requires $O(\log n)$ successful attempts. Conditional success probability at least q_0 gives geometric domination; a Chernoff or elementary exponential-supermartingale calculation bounds the number of attempts by $O_{c,q_0}(\log n)$ except with probability n^{-c}. Apply this domination to indicators of actual accepted valid progress before the first primitive failure, set to one thereafter in the analytical continuation. Before failure the checkpoint is valid; afterward the artificial success probability is one. The lower-tail bound therefore needs no conditioning on eventual absence of failure. On the no-failure coupling event the indicators agree with actual progress up to termination; add the first-failure probability separately. Set that number as a deterministic outer cap. No independence between successive checkpoints is assumed. Multiplying the per-attempt actual-work cap, restoration price, and sparse budget by this attempt count proves the stated bound. Fixed powers of $\log n, P$ in the sparse rows are absorbed by the strict polynomial exponent gap; the dense rows retain only $\ddot{O}(1)$ per attempt.

Contract 4 makes the capped execution finite on every tape. The coupling event gives progress and the work bound; the terminal extraction contract guarantees correctness of every released answer.

Run the terminal phase to its full endpoint. Choose the polynomial-range lower reserve only after C_T is fixed:

$$\mu_{\mathrm{low}} = e^{-C_\mu C_T - 1}\mu_{\mathrm{end}},$$

where C_μ bounds a full phase's logarithmic decrement. A last phase begins above μ_{end} and remains above μ_{low}, even if its endpoint lies well below $\mu_{\mathrm{end}}/2$. The fixed-factor extension preserves all polynomial ranges. At a successful endpoint, $\mu \le \mu_{\mathrm{end}}$ is the only terminal-scale condition required by the stated extraction contract. □

12 Choosing constants and finite precision

Choose margins and physical ranges first, confidence and arithmetic grids second, and sharp work caps last. We check this order, including the growing powers on occupied classes and the cache grid reused each frame.

Write $N = 2 + n$. The input magnitude exponent and the requested final failure exponent c_0 are fixed constants. In this section a polynomial range means $[N^{-A}, N^A]$ for a fixed A; its exponent may depend on the fixed inputs and algorithm, but not on n or P. Bit bounds describe stored operands rather than unused high bits of the machine word.

Finite arithmetic convention. Fix polynomial operand and address ranges, common rounding grids, and iteration and operation caps in the order below. Fixed-depth rational expressions use a constant number of words. Before each operation, check its domain and range without overflowing; a failed guard or cap reports failure under that routine's prescribed rule. These checks enforce finite representation on every tape, including after a hidden primitive error. Accuracy and work estimates are coupled to the reference only until its first primitive failure. Failure probabilities are allocated at valid call starts or epoch births as specified in Section 11, without conditioning on eventual success.

12.1 Scalar enclosures and parameter order

Lemma 12.1. *For every fixed precision exponent D, a one-time $N^{1/8}n^{o(1)}$ preprocessing permits certified enclosures of logarithms of positive polynomial-range rationals to absolute error N^{-D} in $\ddot{O}(1)$ word operations per query. The constants in the polynomial ranges and in D are fixed first.*

Proof. Put $v = \max\{1, \lfloor \log n/8 \rfloor\}$ and tabulate certified values of $\log(1 + j2^{-v})$, $0 \le j \le 2^v$, and $\log 2$ on a common inverse-polynomial grid. On $[1, 2]$ the expansion

$$\log x = 2 \sum_{i \ge 0} \frac{((x-1)/(x+1))^{2i+1}}{2i+1}$$

has ratio at most $1/9$. Its explicit geometric tail and outward rounding give each table entry in $n^{o(1)}$ work. Reserve enough fixed powers of N in the table accuracy for later arithmetic.

For a query, binary search among bit positions finds t with $z = 2^{-t}x \in [1, 2)$ in $\ddot{O}(1)$ operations. Exact quotient selects $j = \lfloor 2^v(z-1) \rfloor$. With $z_j = 1 + j2^{-v}$, write $z = z_j(1+u)$, where $0 \le u \le 2^{-v}$, and use

$$\log x = t \log 2 + \log z_j + \sum_{i=1}^{k} (-1)^{i+1} u^i / i.$$

A fixed k, chosen after D, controls the alternating tail. Outward rounding uses one grid and a fixed number of words. For bounded N, use the original series. Including interval endpoints and range reduction, the query costs $\ddot{O}(1)$. □

Lemma 12.2. *Assume the valid-state inequalities, stopped expected-work bounds, and finite sparse-inverse guarantees established by the preceding constructions. The constants and grids in their implementation can be fixed in an acyclic order, with $P = O(\log \log n)$ and $P \ge 512$, so that:*

1. *every ordinary stored operand uses $O(\log N)$ bits, and all ordinary addresses and deadline keys have fixed polynomial ranges;*
2. *the direction and point-storage defects, source/product errors, and clock debts have any prescribed inverse-polynomial total reserve;*
3. *the exact stream sampler has polynomial memory and its stated unconditional expected cost, even though its reference running time is not bounded on every random tape;*
4. *the dense work multipliers are fixed polynomials in $(P, \log \log n)$, and the displayed sparse rows retain fixed exponent gaps below two.*

Empty headers must be skipped before scalar power evaluation, and all invalid-state arithmetic must be guarded as specified below.

Proof. First fix $p = 512$, ε, κ, local response reserves, $\gamma, a_{\mathrm{jump}}, c_{\mathrm{sol}}, c_Z, \eta, c_{\mathrm{cen}}$, and the sampling lower bounds from the scalar and drift inequalities, without using an outer inverse-margin cap. After the potential allowances below, choose dyadic $M, h, \gamma, \widehat{B}_*$ and the initial/terminal μ scales, including overshoot, before state or source precision. Subsequent paragraphs fix the coarse horizon, confidence, arithmetic grids, word width, and sharp caps in that order. Later sampling multipliers are enlarged only where earlier estimates hold uniformly.

Constants and potential allowances. The tracking drift has equilibrium $C_R m$, and the central drift has equilibrium $C_Z P m$, with C_R, C_Z independent of the outer potential allowances. Choose the inner allowances to dominate these constants using only $P \geq 512$, then choose C_T for the two contraction factors, and then the outer allowances for the frozen-exit estimate. The constant C_* in

$$B_\nu = C_* P \widehat{B}_* \sqrt{1 + \beta_f^2/(\vartheta \widehat{B}_*)}$$

may now depend on the outer central allowance. Increasing it lowers the curvature-variance bound without increasing the drift error. Its work multiplier is fixed before the sharp caps. The same monotonicity permits choosing the final confidence multiplier K after the horizon and killed tracking linear-term bound: the local arguments hold for every K above their original threshold, so the outer allowances need not change.

The outer cap gives $\delta_{\min} \geq cN^{-2/p}$ and $d_{Z,\min} \geq cN^{-2/p}$, uniformly for $P \geq p$. Consequently

$$h/\delta_{\min} \leq C c_h N^{-1/p} P^{-16}, \qquad h/\rho_{\min} \leq C c_h N^{-1/p} P^{-7}.$$

These inequalities permit one fixed sufficiently small c_h after the outer constants have been fixed. Since a phase has $T = O(\sqrt{r}/\beta_s)$ frames, decreasing c_h changes phase length and primitive counts by constants and leaves $TB_* = O_\gamma(C_T m)$ unchanged. The terminal lower scale is a fixed factor $\exp(-C_\mu C_T - 1)$ times μ_{end}. It remains inverse polynomial.

Physical bounds precede accuracy. Section 10.3, with $\mu_{\text{low}} \leq \mu \leq \mu_0$, gives polynomial cost-slack, capacity-margin, grounded-potential, and condition-number bounds on the expanded valid range. The three owner coefficients and physical coefficients are fixed-depth positive expressions in those quantities, $w \in [3s, C]$, and $k \in [1 - \varepsilon, 3 + \varepsilon]$. Hence their inverse norms, directions, and required derivatives are bounded before any mesh is chosen. Contractive refinement and the substochastic triangular-factor bound use one absolute grid for all $O(\log n)$ sparse inverse iterations.

Choose a sufficiently large fixed coarse admission exponent A_N. Its operational stop bounds the records, lists, and headers polynomially. The vector and coefficient caps also bound public source and mark position counts by a polynomial, with exponent A_{pos}. For a coarse frame/attempt exponent A_T, full enumeration has a fixed exponent $A_{\text{call}} > A_N + A_{\text{pos}} + A_T + 10$. This gives a combinatorial call horizon before bounding random sampling work. For the same-frame matrix signed correction bounded by N^{A_E}, assigning failure at most $N^{-c_0 - A_E - A_{\text{call}} - 10}$ makes its entire sum negligible. Enlarge the earlier concentration multipliers to achieve this exponent. Apply the call-start and epoch-birth convention above. Later observation failures kill nonnegative candidates and need no signed correction.

Growing central powers. Put $B = CP^2 m$ for the expanded central allowance. For every occupied central header its representative comparisons give $d_Z^{-P} \leq CB$. For every fixed derivative order j,

$$|\psi_P^{(j)}(b)| \leq 2(P + j)^j (CB)^{1 + j/P} \leq \ddot{O}(N^{2 + 2j/p}).$$

For every fixed $0 \leq j \leq p$, power means also give

$$\sum_e d_{Z,e}^{-j} \leq m(CP^2)^{j/P} = O_j(m).$$

The constants are uniform in $P \geq 512$ because $P^{2j/P}$ is bounded. No step calls for the moment order $j = P$ in a fixed-order density or weighted-price estimate.

An unused center may have $d_Z = N^{-a}$ and potential N^{aP}. Test multiplicity first: zero-mass headers skip powers, derivatives, breakpoints, and reciprocals; occupied headers use the finite arithmetic guards. Valid-range inputs never trigger them.

For an occupied header, first enclose the inverse bases $(1-b)^{-1}$ and $(1+b)^{-1}$, and then use rounded repeated squaring on those bases. Only exponents up to $2(P+j)$ are needed. Its exact positive intermediates lie between $2^{-2(P+j)}$ and a fixed polynomial in N. Since $P = O(\log\log N)$, the reciprocal lower bound is a fixed power of $\log n$. Relative enclosure error u per operation accumulates to at most $\exp(CPu) - 1$. One absolute grid with $u \le N^{-D}/(CP)$, refined by a fixed polynomial factor, therefore gives the required absolute accuracy. This avoids the artificial $N^{C\log P}$ loss from separate Lipschitz bounds at each squaring. At a nonzero center, $|\psi_P'(b)| \ge 2P(P+1)|\mathbf{b}|$ and its dyadic description has a polynomial nonzero gap; thus relative G^c, J^c enclosures are also available. No exact denominator of length $P\log n$ is stored.

A fixed-depth precision ladder. If the direction perturbation bounds are N^{A_E} over at most N^{A_T} frames, take $\Delta_D \le N^{-A_E-A_T-c_0-20}$, also refining it for the stated minimum central radius and source norm. This uses only physical magnitudes. Next choose the point/state mesh so its polynomial sensitivity times its mesh is at most a fixed small multiple of Δ_D, and so its accumulated residual and potential error fits the numerical reserve. The shifted weight storage $\boldsymbol{\lambda} = s + q_{\mathrm{st}}\Delta_{\mathrm{state}}$ preserves the floor exactly. Store μ on a mesh much smaller than $\mu_{\mathrm{low}}\beta_s/\sqrt{r}$ and use the actual ratio $\theta = \mu'/\mu$. No product of old ratio denominators is retained.

The logical state bit bound $W_s = A_s \log n$ is now fixed. Coefficients and old–new differences are fixed-depth rational expressions, so every nonzero $\boldsymbol{\lambda}|\mathbf{b}|$ is at least 2^{-CW_s}. Use this separation to choose the upward normalizer grid. The central rate uses Lemma 3.2: its exact zero test and outward rational enclosure are chosen after these logical operand bits, without altering the analytical noise. Neither enclosure changes the direction grid or requires resolving state cancellations below this separation. The exact-zero branch and fixed $d_c \ge \Delta_D$ absorb point rounding in $|b^{\mathrm{sel}}| \le Cd_c/p$. The ordinary metric remains exact.

Source and packet grids follow this separation bound and the public position counts. Source Richardson iterations reuse one absolute grid because their error convolution is contractive. The observer's relative filter evaluation has only $k = 16$ layers. First fix its positive dyadic action-error tolerance $\nu \le c_{\mathrm{num}} d_0^{-1/4}$, also below the ordinary numerical reserve. Since $R_0 \ge 1$, this gives $\nu^2 \le CR_0^2/\sqrt{d_0}$; the choice uses only the already fixed polynomial-range filter parameter d_0. At each layer the previous product grid provides a nonzero input lower bound; choose the next grid from it and the fixed spectral bounds. A recursion $A_{j+1} \le CA_j + C'$ for 16 stages has a fixed final exponent. Such a recursion over $\Theta(P)$ or $\Theta(\log n)$ layers would not justify the same claim and is not used here.

Choose the signed clock mesh Δ_C after the horizon so that

$$2T_{\max}a_{\max}\Delta_C \le \rho_{\min}/1000, \qquad CmT_{\max}a_{\max}\Delta_C/\rho_{\min} \le N^{-c_0}.$$

Round signed vertex increments once to this mesh and the common positive increment upward. Every debt translation is then exact integer arithmetic. With magnitude bound N^{A_C} and mesh N^{-D_C}, a shifted key needs at most $b = (A_C + D_C + 2)\log_2 N$ bits. Fix that key width first. A reserved deadline block has size $2^{O(b)}$; polynomially many fresh instances still require only polynomial addresses. The arithmetic word may be enlarged afterward to hold those addresses. Its extra bits do not enlarge the deadline universe or change a scalar grid.

Proxy caches on a fixed grid. First fix a to bound the physical base and observer coefficients, their relevant positive eigenvalues, and rate multipliers by $[N^{-a}, N^a]$. The observer coefficient has

a fixed positive metadata grid chosen from these bounds; more accurate fine reads leave that grid unchanged. Also fix the coarse proxy copy-count guard N^q from the sparse count bounds, before proxy precision. On a valid old proxy, spectral comparison with the current base gives $c_\alpha \le N^a$ and $\|T_{\mathrm{old}}\| \le N^a$. Thus for every assignment of the atom signs,

$$\|\mathbf{x}_j^{\mathrm{rand}}\|_2 \le N^{q+a+1}, \qquad \|T_{\mathrm{old}}\mathbf{x}_j^{\mathrm{rand}}\|_\infty \le N^{q+2a+2}.$$

Use the coordinate guard N^{H_c}, where $H_c = q_{\exp} + 2a + 4$, with the numerical reserve. Choose the observer score floors γ, γ_Q from the physical bounds in (CF1)–(CF4), with $N^{-d_c} \le \gamma, \gamma_Q \le 1$ for a fixed d_c. Every non-aborted score evaluation then satisfies

$$N^{-d_c} \le \mathcal{A}_e, z_v \le N^{K_c}, \qquad K_c = a + 2H_c + 6.$$

The field guard and positive score floor enforce this representation bound even on a bad spectral tape.

The cache, residual-M, and residual-Q packet weights are respectively

$$\frac{\mu_e}{C\log n\mathcal{A}_e}, \qquad \frac{\omega_e}{CR_0\epsilon^{-2}\log n\mathcal{A}_e}, \qquad \frac{\omega_e}{CR_0\epsilon^{-2}\log n\mu_e z_v}.$$

Consequently each positive packet lies in $[N^{-k_c}, N^{a+d_c+2}]$ for $k_c = 2a + K_c + 2$, before its output grid is chosen. The lower bound uses the score *upper* bound, independently of previous inverse precision or score-floor denominator. The score floor bounds packets above even at zero empirical distance. Choose once $\sigma = 2^{-g\log_2 N}$ with $g_{\mathrm{key}} \ge k_c + a + 10$, also fine enough to divide the fixed deterministic core grids. It preserves positivity and gives the required relative rounding reserve. Every new cache returns to this same packet mesh. The cold static cache and the exact core atoms fit these bounds after one fixed enlargement of a, q, k_c.

Only afterward choose the field accuracy, for example coordinate error at most $N^{-d_c-a-H_c-8}$ for (CF1), and the sparse inverse's pivot and scratch grids from that accuracy and input mesh σ. The next cache returns scores to γ and packets to σ without inheriting the scratch grids. The Q_0 coordinate fields use its explicit current base core and residual diagonal. A component with positive residual diagonal has a polynomial grounded eigenvalue bound from its base edges and one edge to ground; zero-residual components use their quotient inverse. Thus the same choice applies with γ_Q.

In the observer filter, a nonzero layer product on grids σ and Δ_j has $\|P_M z\|_{Q^{-1}} \ge \sigma\Delta_j/\sqrt{\lambda_{\max}(Q)}$. This determines the next action grid once for each of the fixed 16 layers. Each frame restarts this ladder; its output never becomes a cache coefficient. Public categorical counts are fixed from these guarded magnitudes and envelope bounds. They may exceed the coarse copy-count guard: crossing that guard aborts, rather than changing the position law. Range or count failures abort without redrawing fields or packets. Spectral and numerical claims stop at the first primitive failure; representation bounds hold on every executing tape. Expectations remain unconditional on eventual guard success.

Exact stream inputs and endpoint precision. After representing all actual rational probabilities, fix A to bound their numerator and denominator bit lengths by $A\log n$, their list lengths by N^A, and the total number of calls by N^A. The exact sampler of Section 7 then has polynomial memory and uniform conditional expected cost. For the finite endpoint sampler, choose $B > 2A + c_0 + 20$ and $D > 3A + c_0 + 20$ in its prefix length $B\log n$ and mesh $2^{-D\log n}$. Its total coupling failure is inverse polynomial. Choose scalar-table accuracy from the required numerical

reserves and then fix the final arithmetic word. This does not change the stream inputs or their exponent A. Only the capped execution needs an actual-work counter; the exact reference may continue on its original tape with bounded memory.

Final caps and work exponents. Use Section 11's prospective token guard after its coarse-stop expectation and before copying. Its sharp cap is eventually below the coarse cap and, for $T_0 \asymp m/\log^{1024} n$, gives pathwise depth $\ddot{O}(1)$. Include copying, metadata, sources, and exact sampling in the complete expectation before choosing the actual-work cap. The coarse horizon bounds call starts; sampling work uses the uniform conditional expectation proved above.

Every explicit growing-P factor is fixed degree: h^{-1} contributes P^{16}, the central floor contributes P^{12}, the normalizer contributes P, the central potential allowance contributes P^2, the central radius has inverse at most $CP^9 d_Z^{-1}$, and the center count is $\ddot{O}(\log n)$. Products of their fixed number of uses remain fixed polynomials. In particular the principal sparse exponents are

$$\frac{1001}{512} < 2, \qquad \frac{1011}{512} < 2, \qquad \frac{15}{8} + o(1) < 2.$$

Any fixed powers of $\log n, P$ can be absorbed in these strict gaps. The dense row retains only $\ddot{O}(1)$, and $O(\log n)$ total attempts introduce its one explicit factor $\log n$. All constants can be increased to cover the finitely many inputs below the fixed asymptotic cutoff, using the separately specified finite exact fallback for bounded input size. This completes the parameter-order implication. □

13 The complete algorithm

The following state and execution rules satisfy the four phase contracts; Section 14 verifies them.

13.1 Persistent state and the three coefficient owners

Store original-label arrays $\mathbf{f}, \boldsymbol{\lambda}, \mathbf{c}, \boldsymbol{\upsilon}^{\max}$, vertex potential $\mathbf{y}$, and common parameter μ, with $\mathbf{f}, \mathbf{y}, \mu$ on fixed common grids and $\boldsymbol{\lambda} = s + q\Delta_{\text{state}}$ with an integer $q \geq 0$. Unselected labels retain their stored $\mathbf{f}, \boldsymbol{\lambda}$, so the structural metrics

$$\boldsymbol{\varrho}_e^2 = \frac{\mathbf{z}_e^2}{\mathbf{w}_e^2}, \qquad \frac{\mathbf{z}_e^2}{\mathbf{w}_e}, \qquad \frac{\boldsymbol{\varrho}_e^2}{\boldsymbol{\lambda}_e} \tag{37}$$

change only on the selected list. Their fixed-factor dyadic envelopes support, respectively, the ordinary inverse, physical solve, and observer sampling. The latter two graphs are maintained bases: the graph with coefficient $\mathbf{z}^2/w$ is spectrally comparable to $\mu\mathbf{L}$, while the observer's margin-dependent coefficient ω remains metadata. Thus changes to $\mathbf{y}, \mu$ and reference metadata need no dense structural update.

Use the static construction and original-label pruner of Theorem 4.1 in each simple layer and amplitude bucket. The hybrid counter of Section 5 uses $\tau_0 \asymp n^{9/8}$, $R \asymp n^{1/16}$, $T_0 \asymp m/\log^{1024} n$. A core's current cut certificate is in its immutable actual birth-degree reference. Its original labels occur in compact live arrays with direct handles; private deleted ghosts are absent from scalar and sampling arrays.

Use a stored field coordinate as each fixed-reference weighted median, shared at both endpoints. The energy bound (15) holds for every field, including constant and past-dependent fields. Start under the coarse polynomial token stop, charging lower-capacity records at their full worst deficit price; impose the sharper cap below.

Use the joint lists, exact-count guards, and saved cache contributions of Section 8.3, further refined by coefficient bucket and required public precision bin. Their $n^{o(1)}$ headers per core vertex belong to field preparation; touched labels update only a fixed number of handles. The physical metric needs the narrow actual-$\boldsymbol{\lambda}$ bin to control $w^{-1/2}$ in addition to $\mathbf{z}^2/w$. Lemma 8.3 handles $\mathbf{y}, \mu$ changes without member scans. Changing header waterfills or feedback causes no label migration.

13.2 One physical frame in execution order

Begin with all previous requests served and tracking cells, central windows, coefficients, lists, and owners valid. Execute these steps in order.

1. *Freeze the inputs.* Form the class masses and potential upper certificates. Compute the rational represented source $\mathbf{F}$ and class feedback $\boldsymbol{\varphi}$, including its lower-wall guard. Choose the next stored μ' on its fixed grid so that its actual ratio $\theta = \mu'/\mu$ has the required decrease. Freeze all source branches, actual coefficients, and old read and deletion prices.

2. *Prepare the sparse operators.* Use the physical-base owner to construct the sparse physical proxy, sampled right-hand side, and direction $\mathbf{x}^{\mathrm{dir}}$ of Lemma 6.3. Prepare or retain a fixed inverse of the old ordinary graph with the scaling required for the source iterations. Construct the base top-list observer filter using fresh proxy tapes. This filter is fixed before the physical mask.

3. *Fix the original-label rates.* Use the fixed lazy direction $\boldsymbol{\Upsilon}_c$ of Section 3, evaluated from the old data. Use $\mathbf{d}_c = |\boldsymbol{\Upsilon}_c| + |\boldsymbol{\varphi}|$ and the corresponding $\boldsymbol{\nu}_c$ in every physical, quadratic, wall, and curvature rate. Use certified rational upward enclosures as the actual component probabilities. For the central rate, Lemma 3.2 supplies the exact-zero branch and a relative enclosure with no numerical background rate. For the tracking wall, the literal target is $C_\delta \mathbf{d}_c/(a_{\mathrm{jump}}\overline{\boldsymbol{\delta}})$; the certified comparison $\overline{\boldsymbol{\delta}} \le C_\delta \delta$ makes it dominate $\mathbf{d}_c/(a_{\mathrm{jump}}\delta)$ without reading the true current score. The represented class source fixes the source-cap set. All target probabilities depend only on this frozen data.

4. *Draw one product mask on original labels.* Homogeneous Bernoulli streams use Theorem 7.1. Thin endpoint unions only on each original label's first proposal, using the actual incidence-union denominator. Mark a label processed even when its thinning coin rejects it. Distinct labels use independent coins. Evaluate the full union probability p_e on every selected label and use the actual inverse probability $1/p_e$.

5. *Test the candidate.* Compute and buffer the selected finite candidate points. Unselected points are literal no-ops. Form the exact candidate residual from the maintained n-entry residual vector and the selected stored endpoint changes, using Lemma 6.4. Its norm test applies the retained old physical inverse with a fixed gap that reserves the new-metric comparison. Form the upward common-grid source normalizer Γ of (19) from actual rational coefficient changes. A failed residual or normalizer test aborts the entire attempt.

6. *Finish the old-graph work.* With the old graph and its owner still installed, construct all conserved change sources, perform their original-product Richardson iterations, and complete the observer filter evaluations and old-view clock increments. Generate residual marks with their public categorical-position count. Finish the signed central-clock increments too. Selected labels, positive-account requests, residual marks, and central deadlines are buffered and coalesced by original label. Use the operators from Step 2.

7. *Commit the batch.* After the prospective ledger and token checks below pass, commit the actual selected $\mathbf{f}, \boldsymbol{\lambda}$ and the common $\mathbf{y}, \mu$. Detach the whole genuine coefficient-change batch at its frozen old prices before pruning, retirement, or carries. Pruner exports and carries move live ownership and translate their stored debts and due flags; they are not new external deletion sources. Update actual-$\boldsymbol{\lambda}$ bins before preparing final-graph fine fields. No new physical query occurs during this commit.

8. *Refresh the observations.* Read each queued fine label once from the shared final fields at its old frozen margin price. Close its old public observation epoch before replacing its reference metadata and starting new accounts. Serve central-only requests without a fine score read. Install all final joint classes and guards and evaluate the returned state's potential upper certificates. Check phase/attempt caps before another frame begins.

Deduplicate by generation stamps, then service each coalesced list in original-label order; fix source-packet order before signs and pads. For IDs in $\{0, \dots, m-1\}$, $m = O(n^2)$, a fixed number of stable radix passes with power-of-two radix in $[n, 2n]$ costs $O(\mathbf{z} + n)$ on $\mathbf{z}$ labels. The fixed number of lists charges $O(\mathbf{z})$ to enumerated labels and $O(n)$ array initialization to sparse frame preparation. Thus private proposal order does not control public commit order.

An abort discards the pending transaction, attempt-local caches, and tapes, then restores the last clean checkpoint; a retained frame never rewinds a consumed random cursor.

13.3 Finite arithmetic

Use the fixed direction mesh, bounded-denominator enclosures, and coefficient-separation and clock grids of Section 12. Paired packets share one rounded scalar for exact conservation. Round signed fields before clock increments, positive clocks upward, and transfer debts exactly in stored units with upward credit for unit conversion. The parameter-order proof bounds total clock error and avoids cross-label denominator products.

13.4 Work ledger and operational caps

Before each mask freeze $R_e = 1 + W_e + \bar{\delta}_{e,\mathrm{old}}^{-3}$, the price used for raw proposals and selected updates. Lemma 5.2 and the sparse-component bounds in Section 5 cover every rate in Step 3, using its implemented direction envelope.

After Lemma 14.1 bounds the genuine old-price ledger Z under the coarse stop, set $N_{\mathrm{sharp}} = \ddot{O}(m)$. Before each structural commit, test the cumulative certificate $N_{\mathrm{init}} + 81 Z_{\mathrm{pros}} \le N_{\mathrm{sharp}}$ in the owner's normalized prices, retaining a per-token insertion guard. This uses the enumerated old-price batch before its exports are known and bounds historical copy depth by $\ddot{O}(1)$.

Central inspections, joint-list moves, feedback cache changes, and signed-index operations cost $\ddot{O}(1)$ per touched label. An occurrence move unable to preserve its signed debts queues a cheap central inspection, without a fine score read. Header enumeration is sparse field preparation.

Lemma 14.1 bounds stopped-reference expected work, and Section 15 gives its explicit dense factor. Choose the actual work cap afterward as in Section 11; exhaustion aborts the whole attempt, without redrawing an individual mask, source, or proxy.

13.5 Clean starts, amplification, and exact output

Initialize once by Section 10. At each later clean checkpoint perform the observer's full public refresh using current packed fields before zeroing debts, and save the physical arrays and certified

references. A retry literally restores the saved checkpoint state and rebuilds owners and caches as specified in the proof of Theorem 9.1.

Run complete clean phases under the total attempt cap of Theorem 11.4. When a successful phase ends with $\mu \le \mu_{\mathrm{end}}$, perform the one-time exact extraction of Section 10. The common-shift dual or hub cut certifies the original answer; matching returns original edges and an original-cost integer dual. External exact flow is used only in sparse structural subproblems.

14 Why the phase theorem applies

We verify the four contracts of Section 11 for the algorithm in Section 13, starting with its state invariant.

Lemma 14.1. *Fix the input exponent and a polynomial confidence exponent. With the constants and grids ordered as in Lemma 12.2, the implementation in Section 13 supplies all four contracts of Section 11. There is a polynomial $Q_{\mathrm{led}}(P, \log\log n)$ of absolute degree, independent of the input-range and confidence exponents (which may change its multiplicative constant), such that one clean attempt, before its actual-work cap and stopped as specified there, has expected work at most*

$$Cm Q_{\mathrm{led}}(P, \log\log n) + n^{15/8+o(1)} + n^{1011/512+o(1)}.$$

The initialization and terminal procedures of Section 10 supply the remaining endpoint premises of Theorem 11.4.

Proof. Work on the reference stopped at its first primitive failure, before imposing the actual-work cap, without conditioning on eventual success.

State invariant and clean start. At each returned state, $\mathbf{f}, \boldsymbol{\lambda}, \mathbf{y}, \mu$ satisfy both tubes and their exact incidence residual meets the physical-metric bound. Tracking cells and central representatives are valid, all preceding selected labels are serviced, and no request or due flag is pending. Each metric has one live owner per label, no public ghosts, and immutable birth-reference core certificates. Clocks retain debt exactly. The preceding-frame sparse base cache is constant-comparable to the current base by the one-frame coefficient bound.

The cold midpoint, deterministic core upper scores, contraction scans, and final-graph reference scan establish this invariant. Warm checkpoints refresh references before zeroing accounts.

Constructing a frame. The owner energy bound (15) applies to every field, including past-dependent fields. Lemma 8.3 constructs tracking feedback, and the waterfill in Section 8 constructs central sources, without member scans. Center counts and outward enclosures certify potentials, skipping empty headers as prescribed.

Lemma 6.2 supplies the leverage bounds for the true-metric direction of Lemma 6.3. Step 2 fixes the filter and old inverse before the mask, without fine reads: physical upper scores use birth references and observer coefficients use metadata. The lazy direction is therefore fixed on every label.

The exact-stream multiplier covers all endpoint and direct occurrences. Thus Steps 3–5 use independent original-label bits and the actual union denominator, meeting the scalar identities and Theorems 2.5 and 2.6. Candidate tests use the rounded point and abort without local retry.

Restoring valid observations. Step 6 supplies the raw-moment premises: source products and capture use the old matrix and actual coefficient changes, with the filter fixed before those changes are sampled.

An unrequested label is unselected, so $\boldsymbol{\varrho}_e, \boldsymbol{\lambda}_e$, its epoch price, and threshold stay fixed. Positive-account ceilings, residual no-hit bounds, capture, and inverse order protect its window up to first exit. The one-frame margin bound protects requested labels until their final shared reads, using the same old certified references.

The tracking tube $|\log(\boldsymbol{\lambda}/(s+\boldsymbol{\tau}))| < \kappa$ supplies both comparisons required by the observer: $\boldsymbol{\tau} \le C\boldsymbol{\lambda}$ and $\boldsymbol{\lambda} \le C(s+\boldsymbol{\tau})$. Its reference $\overline{\boldsymbol{\varpi}}$ stays comparable to $s+\boldsymbol{\tau}$ throughout the epoch. On the same-frame safety event the actual candidate has $\boldsymbol{\tau}'_e \le C\boldsymbol{\lambda}'_e$ and $\boldsymbol{\lambda}'_e \le C(s+\boldsymbol{\tau}'_e)$, and its true tracking margin is comparable to the old one. Hence the packed read's additive error $c_\star\overline{\boldsymbol{\delta}}\boldsymbol{\lambda}$ is a relative error $O(c_\star\overline{\boldsymbol{\delta}})$ in $s+\boldsymbol{\tau}$. Choosing $c_\star$ within the tracking-cell reserve gives the required logarithmic accuracy.

The updated independent actual-$\boldsymbol{\lambda}$ index supplies the final-graph fine proxy and lanes before joint tracking classes change. Step 8 then uses the simultaneous prefix guarantee for adaptively requested labels and precisions at their old margin prices. New interior windows and central-only service restore the invariant.

The commit rule in Step 7 retains due flags and signed and positive debt at frozen old prices; due positive debt uses old credit. Representation changes pay any needed central inspection, without resetting tracking epochs or creating genuine deletion sources.

Contract 1: potential certificates. The restored windows and occupied-header arithmetic give $\Psi_R \le U_R \le 2\Psi_R$ and $\Psi_Z \le U_Z \le 2\Psi_Z$.

Contract 2: clean checkpoints. The packed all-label checkpoint read costs $\ddot{O}(\sum \overline{\boldsymbol{\delta}}^{-3}) = \ddot{O}(m)$. Rebuilding owners/caches and copying physical/reference arrays cost $\ddot{O}(m)$. Retry discards failed mutable state, restores saved arrays, references, and point enclosures literally, and rebuilds graph/proxy caches with fresh tapes. The saved refresh, central inspections, and baselines justify zero new debt without cold lane scans.

Contract 3: expected work before the actual-work cap. Under the provisional token stop, Lemma 5.1 gives pathwise moments before any event-count estimate. Lemma 5.2 prices all raw proposals, including rejects and duplicates, at frozen old weight $1+W_e+\overline{\boldsymbol{\delta}}_e^{-3}$. Its curvature estimate cancels $(1+\zeta_*)^{1/2}$ and $(1+\zeta_*)^{-1/2}$ pointwise. Thus the dense conditional mean is $\ddot{O}(B_*)$ without a cumulative curvature assumption; other physical components retain their strict polynomial reserves.

Lemma 11.1 gives $\sum J_t B_* \le Cm$ deterministically, so genuine events cost $\ddot{O}(m)$ in expectation. Pathwise pruning gives $N \le C(m+Z)$ for admissions and the old-price genuine ledger Z. Impose Lemma 5.3's sharp token stop before using historical depth $\ddot{O}(1)$. This pays copies, exports, and debt-preserving migrations; static births cost $n^{15/8+o(1)}$. Copying uses the stopped prefix, not a convex function of $\mathbb{E}N$.

The selected-rate bound supplies Theorem 9.1's selected-read premise. Lifetime positive credit and residual intensities pay autonomous reads at $\ddot{O}(1)\mathbb{E}\sum J_t B_*$, plus selected reads and sparse preparation; releases may use earlier frames' credit. Lemma 8.2 pays central reports by signed displacement, paid visits, and migrations. Neither kind of report changes structural coefficients or creates genuine events.

Packed assembly costs $\ddot{O}(1)$ expected work per selected packet plus sparse destination setup. Use the one-uncapped-attempt padding analysis and same-tape exact sampling of Section 11, be-

fore the actual-work cap and without redraw. Inverse, field, header, timer-header, and proxy preparations have exponent at most $1011/512 < 2$. These bounds also pay the prospective genuine ledger. Markov bounds genuine-ledger, sharp-token, and actual-work cap crossings on their stopped references by small fixed probabilities before a primitive failure; add the primitive-failure budget afterward.

Contract 4: failure accounting and finite execution. Each structural, proxy, solve, source, or fine-field call is fixed before its independent tapes, permitting the call-horizon union bound. Residual no-hit bounds cover adaptive futures from valid zero-target epoch births; union over births conditional on the checkpoint, not intermediate misses.

The finite arithmetic convention of Section 12 enforces finite execution and first-failure coupling, with ranges fixed before sharp work caps. Later failures kill nonnegative drift candidates; only the same-frame tracking matrix event needs signed correction. The separate frozen outer-crossing process retains its witness, as required in the phase proof.

Endpoint hypotheses. Section 10 supplies the connected dense strictly feasible midpoint, inner start, polynomial data and scale ratio, precision bounds, and terminal feasible-gap premise. Lemma 10.1's hub alternative and the terminal primal/dual checks preserve and certify the original answer. Lemma 10.4 pays cold construction and $O(m \log n)$ terminal work once, separately from interior work. □

15 The explicit bound for dense work

We expand Lemma 14.1's dense polynomial: weighted event counts pay copies, and each copy pays scalar and handle work. Write $p = 512$; the exponent $P \geq 512$ with $P = O(\log\log n)$ is the integer fixed in Section 1.1. All constants may depend on the fixed input-range and confidence exponents. The powers of $P, \log\log n$ displayed here do not.

Lemma 15.1. *The dense expected work of the stopped clean attempt in Lemma 14.1 is at most*

$$O\left(mP^{14}(\log\log n)^{19}\right) = \ddot{O}(m).$$

The looser polynomial $(P+1)^{16}(\log\log n + 1)^{24}$ *is therefore an admissible choice for its* $Q_{\mathrm{led}}(P, \log\log n)$. *The separately specified sparse rows, one-time cold construction, and one-time terminal construction are unchanged.*

Proof. Weighted event counts. Put $R_e = 1 + W_e + \overline{\boldsymbol{\delta}}_e^{-3}$, as in the weighted raw-rate lemma. The current-density calculation and the fixed tracking moment give

$$\sum R_e^{12} = O(m(\log\log n)^{132}), \qquad \sum_{e,o} \phi_{eo}^{-12} = O(m(\log\log n)^{60}).$$

For every fixed order j used here, the central potential allowance gives $\sum d_Z^{-j} \leq m(CP^2)^{j/P} = O_j(m)$ uniformly in $P \geq 512$. Thus Hölder gives, explicitly,

$$\sum R^4\phi^{-4}d_Z^{-2} = O(m(\log\log n)^{64}), \qquad \sum R^2\phi^{-2}d_Z^{-2} = O(m(\log\log n)^{32}),$$

where the second formula holds for either certified margin d. These are population moments of the current owners, including their low-capacity records; they are not historical update counts.

In the notation of the weighted raw-rate proof, the two factors in the curvature bound are

$$S_R \le C_\gamma P^2 (\log\log n)^{32} (m/s)(1+\zeta_*), \qquad B_\nu/V \le C_\gamma P (1+\zeta_*)^{-1/2}.$$

Pairing $\sqrt{S_R}$ with B_ν/V cancels ζ_*. Since $k_P = CP^{12}$, the resulting curvature price is at most $CP^{14}(\log\log n)^{16} B_*$. The tracking-wall, central-wall, and forced source-wall rows are respectively bounded by $C(\log\log n)^{16} B_*$, $CP^{12}(\log\log n)^{16} B_*$, and $CP^2(\log\log n)^{11} B_*$. These prices cover rejected and duplicate raw hits, deletions, and fine reads at their old margins. As $\sum J_t B_* \le C_\gamma m$ in an attempt, the resulting dense weighted count is

$$\mathbb{E} \sum_{\text{raw}} R_e = O(mP^{14}(\log\log n)^{16}), \tag{38}$$

apart from the already separated sparse physical components. The P^{16} in the number of frames does not multiply this quantity: it has canceled against β_f in TB_*.

Cost of a visited label. A changed label updates a fixed number of handles, cache summands, keys, and accounts. Signed-key access and exponent comparisons cost $O(\log\log n)$. Targets are fixed-depth rational or quadratic-field expressions; the separation bounds of Section 3 and the scalar-update routine of Section 2 give square-root and scalar-root evaluation in $O(\log\log n)$ stages, at most $O(\log\log n)$ work each. Thus cheap point operations cost $O((\log\log n)^2)$. The normalized root derivative bound is uniform at capacity walls. Growing-P powers use $O(\log P)$ rounded multiplications with polynomial values and precision by Section 12. Only sparse preparation scans all headers; touched labels use their own handles.

Packed fine reads cost $O((\log\log n)^2 \overline{\boldsymbol{\delta}}_e^{-3})$ despite their $\Theta(\log n \overline{\boldsymbol{\delta}}_e^{-2})$ confidence lanes. Source coefficients have $O(\log\log n)$ mantissa bits, fitting the $O((\log\log n)^2)$ square comparison allowance; their fixed number of packed lane blocks costs $O(\log\log n)$ conditional expected carry work. Setup and transposition stay in the sparse $nT\log n n^{o(1)}$ row. Exact gaps cost $O((\log\log n)^2)$ expected work per start and raw success, by Section 7.

Copies and central reports. Equation (38) also bounds the expected genuine deletion ledger Z. Thus the sharp prospective stop may use $N_{\text{sharp}} = C_{C,c_0} m P^{14} (\log\log n)^{16}$ with a sufficiently large fixed multiplier. The pathwise certificate $N \le N_{\text{init}} + 81Z$ and the mixed-radix argument give

$$1 + \log(1 + N_{\text{sharp}}/T_0) = O(\log\log n), \qquad \mathbb{E} N_{\text{copies/moves}} = O(mP^{14}(\log\log n)^{17}).$$

Charging $O((\log\log n)^2)$ for the attached scalar and handle service therefore costs

$$O(mP^{14}(\log\log n)^{19}).$$

The unweighted fixed-birth pruning work is already in W_e; there is no further dynamic-tree factor on that dense row.

The central scheduler's radius satisfies $\rho(\overline{\boldsymbol{\varsigma}})^{-2} \le CP^{18} d_Z^{-2}$. Its actual endpoint variation is therefore at most $CP^9(\log\log n)^5 B_*$, by its displayed fourth-moment estimate; the common positive-clock contribution is at most $CP^9 B_*$. Its lifetime inspection count is bounded by

$$C\left(m + P^9(\log\log n)^5 \sum J_t B_* + N_{\text{paid}} + N_{\text{move}}\right).$$

Paid visits and moves enter additively, without the $P^9(\log\log n)^5$ multiplier. An inspection changes list membership without creating another occurrence-representation move. Its $O((\log\log n)^2)$ service is dominated by the copy bound.

Observer and checkpoint rows. The autonomous observer supplies conditional new-credit and raw-mark expectation $O(B_*)$. Its residual multiplier obeys

$$\frac{\log n A_L R_0^3}{\sqrt{d_0}} \le n^{-8/p+o(1)} = O(1)$$

for the observer's fixed sufficiently large logarithmic exponent D_0 and the asymptotic cutoff. The sharp true-energy trace estimate, not the crude R_0 comparison, pays remaining positive credit. Autonomous reads therefore cost $O(m(\log\log n)^2)$ per attempt, including initial zero-credit conventions. At a checkpoint, $\sum \overline{\boldsymbol{\delta}}^{-3} = O(m)$ gives $O(m(\log\log n)^2)$ for complete packed refresh. Saving physical/reference arrays costs $O(m)$; restoration and bulk owner/metadata construction cost $O(m(\log\log n)^2)$, charging only records actually built. The cold lane scan remains a separate one-time cost.

Dense ledger. The resulting dense ledger is

raw physical proposals, selected scalar work and fine reads	$O(mP^{14}(\log\log n)^{18})$,
conserved source preparation and expected packed additions	$O(mP^{14}(\log\log n)^{18})$,
genuine coefficient changes and unweighted pruning	$O(mP^{14}(\log\log n)^{16})$,
copies, migrations, and attached scalar/debt service	$O(mP^{14}(\log\log n)^{19})$,
autonomous central reports	$O(mP^{9}(\log\log n)^{7})$,
autonomous observer reads and complete checkpoint service	$O(m(\log\log n)^2)$.

Central inspections paid by a prior visit or move are included in that visit or move's scalar service. All charges use their frozen old read prices. Exact stream work is included in the conditional expected scalar price. This proves the lemma. □

Where the full logarithm enters. The terms $n^{15/8+o(1)}$ and $n^{1011/512+o(1)}$ cover the full header population, solves, proxy/field lanes, timer-header scans, and sparse exact-flow backend. Their exponent gaps absorb fixed powers of $\log n$; this does not apply to dense $m\log^a n$ terms. Cold initialization costs $O(m\log n\log\log n + m(\log\log n)^3)$ and is paid once, with its own confidence budget. Terminal correction, integral rounding, and the exact dual construction have dense cost $O((m+n)\log n)$ and are also paid once. Finally, $O(\log n)$ attempts multiply the per-attempt work. Since $P = O(\log\log n)$, the looser polynomial's two exponents add to 40, a conservative bound for all dense rows; in the notation of Section 1.1 every dense row is $\ddot{O}(m)$ and the total is $\ddot{O}(m\log n) = \ddot{O}(n^2\log n)$.

16 Proof of the main result

Proof of Theorem 1.1. Apply the exact preprocessing and augmentation of Section 10 and Lemma 10.1. They preserve the original answer and give a connected working host with $N = \Theta(n)$ vertices, $m = \Theta(n^2)$ labels, and $r = \Theta(n)$, represented in a fixed number of simple layers. Cold construction starts from its exactly feasible midpoint and gives an inner checkpoint with zero residual. The data and μ_0/μ_{end} have polynomial ranges; the dense price is $O(m\log n\log\log n) + mn^{o(1)}$, paid once with polynomial confidence by Lemma 10.4.

Fix constants and meshes by Lemma 12.2. Lemma 14.1 supplies all four phase contracts and the endpoint premises. Theorem 11.4 therefore gives $O(\log n)$ total attempts, with the sharp-token, actual-work, and attempt-count caps specified there and separately budgeted first-primitive failures.

Lemma 15.1 permits $Q_{\mathrm{led}}(P, \log\log n) = (P+1)^{16}(\log\log n + 1)^{24}$. Since $P = O(\log\log n)$, the interior work is $O(m \log n (\log\log n)^{40}) = \ddot{O}(m \log n)$. The two sparse rows in the closure lemma remain $o(n^2)$ after multiplication by $O(\log n)$, by their strict exponent gaps. The one-time endpoint prices also fit this bound; input range and confidence affect constants, not the degree 40. The bounded-input branch covers the finitely many smaller dimensions.

At the terminal checkpoint, Section 10 supplies exact tree correction, cost-nonincreasing integral rounding, and an integer dual, or the verified hub cut when optimal hub flow is positive. Its dense cost is $O(m \log n + n \log n)$, paid once under Lemma 10.4.

Finally verify the proposed answer on the original input in $O(m+n)$ exact word operations. For a flow, verify integrality, both capacity bounds, and all vertex balances. For the supplied integer potential $\mathbf{y}$, put $\mathbf{c}_e^y = \mathbf{c}_e - (Ay)_e$ and verify

$$\mathbf{f}_e < \boldsymbol{\upsilon}_e \implies \mathbf{c}_e^y \geq 0, \qquad \mathbf{f}_e > l_e \implies \mathbf{c}_e^y \leq 0.$$

These are exactly zero complementary slackness on every interval $[l_e, \boldsymbol{\upsilon}_e]$. Together with primal feasibility they prove optimality by weak duality. Self-loops and zero-width intervals satisfy the same direct tests. Alternatively verify the stated bound, demand, or capacity-cut inconsistency directly. Reject any failed check with FAIL. These checks require only polynomial-size integer sums and a fixed number of words per operand.

Allocate the original confidence exponent with a fixed reserve among cold initialization, fresh primitives and epoch events, attempt-count exhaustion, and terminal construction. Their union has probability at most n^{-c_0}. A global word-operation counter enforces the displayed bound on every tape, including all malformed intermediate and rare exact-arithmetic branches. All good executions finish below that cap. Exact final verification makes every answer actually released correct, also on a tape outside the good event. □

Proof of Corollary 1.2. Orient every edge left to right, put lower capacity zero and upper capacity two on each edge, and give each left vertex demand -1 and each right vertex demand 1. A feasible integral flow has exactly one unit edge at each vertex, and hence is a perfect matching. Every edge of this integral flow is strictly below its upper bound. The verified flow dual therefore has $\mathbf{y}_j - \mathbf{y}_i \leq c_{ij}$ on every edge, with equality on the selected edges. Set $\alpha_i = -\mathbf{y}_i$ on the left and $\beta_j = \mathbf{y}_j$ on the right. Then $\alpha_i + \beta_j \leq c_{ij}$ and equality holds on the output matching, giving the stated integer assignment dual. □

References

[BF13] Karl Bringmann and Tobias Friedrich. Exact and efficient generation of geometric random variates and random graphs. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013*, volume 7965 of *Lecture Notes in Computer Science*, pages 267–278. Springer, 2013.

[KS16] Rasmus Kyng and Sushant Sachdeva. Approximate Gaussian elimination for Laplacians – fast, sparse, and simple. In *57th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2016*, pages 573–582, 2016.

[Pin94] Iosif Pinelis. Optimum bounds for the distributions of martingales in Banach spaces. *The Annals of Probability*, 22(4):1679–1706, 1994.

[ST83] Daniel D. Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences*, 26(3):362–391, 1983.

[Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[vdBCP+23] Jan van den Brand, Li Chen, Richard Peng, Rasmus Kyng, Yang P. Liu, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford. A deterministic almost-linear time algorithm for minimum-cost flow. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*, pages 503–514. IEEE, 2023. Available at: https://arxiv.org/abs/2309.16629.

A Construction and pruning of original-edge cores

We prove Theorem 4.1, using the graph, reference, and label conventions of the structural interface.

A.1 A child guard and its fresh sparse sample

The guard estimates degrees by probes and scans a full original row only when removing its vertex. Freeze root incidence arrays with exact positive degrees w, volume $W_0 = 2q_0$, and minimum degree $d \geq 1$. Each child A, given by a vertex list and stamped membership, is fixed before its fresh tapes; its failure budget is $\delta_c = n^{-O(1)}$. Public removal excludes a row from the current original vertex set and scans its root row; private removal changes auxiliary ghost-flow state. These operations need not coincide.

Lemma A.1. *For inverse-dyadic $e \leq 1/4096$, use $k = 8\log_2(8n/\delta_c)/e^2$ independent uniform root-neighbor membership probes at each listed vertex. Let $\widetilde{d}_u$ be w_u times the empirical membership fraction. Repeatedly remove a vertex with $\widetilde{d}_u < w_u/3$, scan its root row, and decrement every still-child neighbor's estimate by one. Except with probability δ_c, on the surviving child C the known integers*

$$r_u = \widetilde{d}_u + 3ew_u, \qquad d_C(u) \leq r_u \leq 2d_C(u), \qquad w_u/4 \leq r_u \leq 2w_u \tag{39}$$

are valid. Every removed vertex had current true child degree below $3w_u/8$. Original degree estimates keep the same additive error under all subsequent exact public row removals.

Proof. Hoeffding and a row union bound give $|\widetilde{d}_u - d_A(u)| \leq ew_u$. Every decrement changes the estimate and truth by the same amount; adaptively chosen removals do not change this error. A survivor has $d_C \geq (1/3 - e)w > w/4$. Also $d_C + 2ew \leq r \leq d_C + 4ew \leq d_C + 16ed_C \leq 2d_C$, using positive integral d_C and $16e \leq 1$. The remaining comparisons follow, and at removal $d_A < (1/3 + e)w < 3w/8$. Only listed rows are initialized. Uniform positions use exact integer rejection; its exponentially decaying extra work is capped globally. □

An empty guarded child terminates immediately. Otherwise write $v = |C|$, choose a known $d_* \leq \min d_C$ and $\rho \geq \max\{1, v/d_*\}$, and take a fresh sample after the guard. For an inverse-dyadic probability choose

$$\min\{1, C_0 \log n/(\epsilon^2 d_*)\} \leq p \leq \min\{1, 2C_0 \log n/(\epsilon^2 d_*)\}. \tag{40}$$

Run Bernoulli streams of rate p on root rows. Inspect every selected position, retaining only internal edges in their smaller-endpoint orientation, with weight $1/p$. Each internal edge thus has one independent eligible position; never search until an internal edge appears.

The game uses the completed Laplacian comparison. The later repair uses the raw adjacency bound to control degree error after adaptive removals.

Lemma A.2. *For a sufficiently large C_0, with the allocated confidence,*

$$\|D_{d_C}^{-1/2}(\widehat{A}-A_C)D_{d_C}^{-1/2}\|\le\epsilon/4,\qquad |\widehat{d}_u-d_C(u)|\le(\epsilon/4)d_C(u). \tag{41}$$

If $\epsilon\le e/4$, complete both adjacencies to known row sums r:

$$A^*=A_C+\textsc{Diag}(r-d_C),\quad \widehat{A}^*=\widehat{A}+\textsc{Diag}(r-\widehat{d}),\quad D=\textsc{Diag}(r),\quad V=r(C),\quad J_r=D-rr^T/V.$$

Both loop completions are nonnegative, and

$$-\epsilon J_r\preceq\widehat{L}-L_C\preceq\epsilon J_r,\qquad \|D^{-1/2}(\widehat{A}-A_C)D^{-1/2}\|\le\epsilon. \tag{42}$$

Sample generation costs $O((v+p\,w(C)+\log n)n^{o(1)})$ operations and that many or fewer original position accesses, including every discarded position and exhausted stream.

Proof. Each independent normalized adjacency summand has norm at most $2/(pd_*)$, and its total variance parameter is $O(1/(pd_*))$. The dimension in matrix Bernstein is at most n. Scalar Bernstein and a union over rows give the degree assertion. If $p=1$ both are exact. The sample degree error is smaller than the guard's positive margin $r-d_C\ge 2ew$, so loops are nonnegative. Adjacency error plus the degree diagonal gives completed error at most ϵ in D norm. Both Laplacians annihilate constants: centering a test vector in r replaces D by J_r. A Chernoff bound caps selected positions by $O(pw(C)+\log n)$. Finite geometric skipping with the same distribution up to a globally budgeted coupling error is specified in Appendix A.5.1. A negative computed loop is reported as failure, rather than used as an edge. □

A.2 A low-trace filter and sparse smoothing

The filter approximately preserves low-energy cuts and reduces the initial potential to $O(\rho)$, controlling the number of game rounds.

Let P subtract the r-mean. Define

$$\widehat{W}=D^{-1}\widehat{A}^*,\quad c_0=1+\epsilon,\quad B_0=(I+\widehat{W}/c_0)/2.$$

Set $h=512\log_2(64n)$, computed by integer bit-length comparisons. It is a multiple of 16, satisfies $h\ge 512\ln(64n)$, and is $O(\log n)$. Put

$$p_h(t)=\sum_{j=13h/16}^{h}\binom{h}{j}t^j(1-t)^{h-j},\qquad F=Pp_h(B_0)P. \tag{43}$$

Lemma A.3. *For the above guard and sample parameters, F is a centered, D-self-adjoint contraction satisfying*

$$\|F\|_{\mathrm{HS},r}^2\le 16\rho+1=:P_{\max},\qquad \|(I-F)x\|_r\le\left(\frac{1}{16}+\sqrt{8\lambda_C(x)}\right)\|x\|_r \tag{44}$$

for every centered x, where $\lambda_C(x)=x^TL_Cx/\|x\|_r^2$. One application takes $O(M\log^2 n)$ sparse arithmetic operations at inverse-polynomial output accuracy, where M counts stored sample atoms and vertices.

Proof. The binomial tail is at most $\xi = e^{-h/512}$ for sampled adjacency eigenvalues at most $1/2$, and at least $1-\xi$ for eigenvalues at least $3/4$. Write the true normalized completed adjacency as $U + \textsc{Diag}(\ell)$ with $U = D^{-1/2} A_C D^{-1/2}$. Simplicity gives

$$\textsc{Tr}(U^2) = \sum_u \sum_{v:uv\in E(C)} \frac{1}{r_u r_v} \le v/d_*, \qquad 0 \le \ell_u \le 16e + 1/d_*.$$

If $d_* \ge 8$, loop norm plus sample perturbation is below $1/4$. At most $16v/d_*$ sampled eigenvalues can exceed $1/2$, by the trace bound on true U and eigenvalue monotonicity under a norm perturbation. If $d_* < 8$, dimension $v < 8v/d_*$ suffices. Compression by P cannot increase Hilbert–Schmidt norm, and $v\xi^2 < 1$ proves the first claim. In particular we do not square all the completion loops.

On the centered subspace, let E_0 project to true normalized Laplacian eigenvalues at most $1/8$, and let Q_0 project to sampled adjacency eigenvalues below $3/4$. True completed adjacency eigenvalues on E_0 are at least $7/8$. The Sylvester equation for Q_0E_0 has spectral separation $1/8$. Its integral solution bounds $\|Q_0E_0\| \le 8\epsilon$. Thus $\|(I-F)E_0\| \le \xi + 8\epsilon \le 1/16$. The complementary component of x has norm at most $\sqrt{8\lambda_C(x)}\|x\|_r$. This proves the second claim. De Casteljau evaluation of the Bernstein polynomial uses $O(h^2)$ sparse products. The amplification $2^{O(h)}$ is polynomial in n, permitting the common $O(\log n)$-bit grid of Appendix A.5. □

To turn sample error into one-sided control of original cuts, join a hub to each original vertex u with weight ϵr_u. Eliminating it gives

$$L_+ = \widehat{L} + \epsilon J_r, \qquad L_C \preceq L_+ \preceq L_C + 2\epsilon J_r. \tag{45}$$

For any active $A \subseteq C$, let $D_A = \textsc{Diag}(r|_A)$ and eliminate the hub and inactive originals:

$$S_A = \textsc{Schur}_A(L_+), \quad \gamma = (1+2\epsilon)^{-1}, \quad K_A = I - \gamma D_A^{-1} S_A, \quad T_A = (K_A + K_A^2)/2.$$

Route on the explicit sample plus hub; apply Schur operators by sparse solves without storing them.

Lemma A.4. *K_A is a reversible nonnegative Markov kernel and a contraction. If $X = T_A Z$, then*

$$\textsc{Tr}(X^T S_A X) \le (\|Z\|_r^2 - \|X\|_r^2)/(3\gamma). \tag{46}$$

For every full original set $U \subseteq C$, extending kernels by the identity on inactive coordinates gives

$$\|(T_A - I)\mathbb{1}_U\|_{r,1} \le 3[\partial_C(U) + 2\epsilon \min\{r(U), V - r(U)\}]. \tag{47}$$

Each action uses two grounded solves on a sparse graph of size $O(M)$.

Proof. Schur signs and zero row sums, with $(S_A)_{uu} \le (1+2\epsilon)r_u$, prove the Markov assertions. Its spectrum lies in $[-1,1]$. Equation (46) follows by diagonalization from $3(1-x)((x+x^2)/2)^2 \le 1 - ((x+x^2)/2)^2$ on that interval. For an indicator, its K displacement is twice γ times its Schur cut. The full original indicator is an admissible extension in the Schur variational minimum, so this cut is bounded by (45). A second K step has at most twice the first displacement by weighted ℓ_1 contraction. Averaging gives the factor three. This argument applies to all original sets, including ones chosen only at the end of the execution.

Here L_{XY} are blocks of the explicit sample-plus-hub Laplacian. A K action solves on eliminated vertices B, including the hub, with right-hand side $-L_{BA}f$. Replace B–A edges by ground edges preserving the diagonal and retain their active values in this right-hand side. The block is positive definite for nonempty A; Appendix A.5 supplies certified solves. □

A.3 One projected candidate and a deficient exact flow

Represent the embedding by actions on signs, retaining active rows and original-child columns. A projected field selects separated source and target demands; flow supplies endpoint mixing and, when deficient, seed removal, the two loss terms below. Equation (56) records smoothing, mixing, seed deletion, and recentering.

Put $V_A = r(A)$. Freeze a centered current embedding Z, set $X = T_A Z$, draw fresh Rademacher column signs g, and compute a centered approximation y to Xg. Let a be a weighted median of y; change the signs of both y and the analytical field Xg if necessary so that $a \geq 0$. All projection events and quadratic estimates are invariant under this common sign change.

The eligible sets are $y < 0$ and $y \geq 3a/2$, omitting zero coordinates. At least $2/11$ of squared energy lies in their union: excluded energy is at most $(9/4)r(A)a^2$, while zero mean and the median force negative energy at least $r(A)a^2/2$. Choose an eligible set containing at least $1/11$ of all energy and take its largest-$|y|$ prefix of reference mass at most $r(A)/8$, splitting its final vertex if necessary. This retains at least $1/8$ of that eligible energy. Use this prefix as source demand l, and the opposite median half as target demand t. Then

$$0 \leq l, t \leq r, \quad \textsc{Supp}(l) \cap \textsc{Supp}(t) = \varnothing, \quad l(A) \leq V_A/8, \quad t(A) = V_A/2, \quad \sum l_u y_u^2 \geq \|y\|_r^2/128. \tag{48}$$

Every source differs from every target by at least $|y_u|/3$. Demands have denominators dividing 8.

Write ∂_+ for cut weight in the explicit sample-plus-hub graph. Give augmented routing edges capacity β times their weight, where β is a dyadic integer, and terminal arcs capacities l, t. Compute an exact maximum flow after common dyadic scaling. If all sources are saturated, put $U = S = \varnothing$. Otherwise let U be the residual reachable routing vertices from the super-source, and $S = U \cap A$. In the latter case, the minimum cut gives

$$\beta \partial_+(U) + t(U) < l(U), \qquad r(S) < 5V_A/8, \qquad \partial_+(U) < r(S)/\beta. \tag{49}$$

Indeed $t(U) \geq r(S) - V_A/2$ and $l(U) \leq V_A/8$.

Cancel opposite flow and cycles, and retain source–target records whose source lies outside U. These sources are saturated and their paths avoid U: an entering positive arc would have a residual reverse arc leaving U. For retained amounts A_{uv} set $R_{uv} = A_{uv} + A_{vu}$ off-diagonal and complete row sums to r with nonnegative loops, using disjoint demand supports and (48). Then $M_R = (I + D_A^{-1} R)/2$ is reversible and fixes seed rows, and $E_R(X) = \textsc{Tr}(X^T (D_A - R) X)$ equals routed endpoint squared energy.

Lemma A.5. *If Δ_X is squared-norm loss after M_R, deletion of S, and recentring, then*

$$\Delta_X \geq \tfrac{1}{2} E_R(X) + \sum_{u \in S} r_u \|X_u\|^2, \tag{50}$$

$$E_R(y) + \sum_{u \in S} r_u y_u^2 \geq 2^{-11} \|y\|_r^2. \tag{51}$$

For every full original $W \subseteq C$,

$$\|(M_R - I)\mathbb{1}_W\|_{r,1} \leq \beta[\partial_C(W) + 2\epsilon \min\{r(W), V - r(W)\}]. \tag{52}$$

The endpoint records can be constructed in $O(M \log(2M))$ work and $O(M)$ records.

Proof. For a reversible $P_R = D_A^{-1} R$ with eigenvalue $x \in [-1, 1]$, $1 - ((1+x)/2)^2 \geq (1-x)/2$. Seed rows are unchanged by M_R, and centering only reduces norm. This proves (50). Each source is

either in S, whose energy is paid by deletion, or routed in full to separated targets. Thus the left side of (51) is at least $\sum l_u y_u^2/9 \geq \|y\|_r^2/1152$. For (52), place the hub on the larger-reference side of W. The resulting augmented cut has weight at most the bracket by (42); every endpoint-crossing route crosses it and has congestion at most β. The lazy displacement equals the endpoint cut weight; the bound uses the explicit augmented cut, not a Schur cut.

Compressed routes. Maintain a forest with at most one selected positive outgoing arc per vertex. The super-source has only outgoing arcs and the super-sink only incoming arcs. Extend from the current root, purging zero arcs on both relevant root paths before root/cycle tests. Link distinct trees; on a cycle subtract its minimum flow. At the sink subtract the path bottleneck and record endpoints and amount.

After source flow is exhausted, scan a fixed arc list to cancel circulations. Query linked values from the dynamic tree. Process each positive arc's tail until its outgoing flow vanishes, so that arc has zero flow before the cursor advances. Flows only decrease: every positive bottleneck zeros an arc, every linked arc is cut once, and a renewed closing-arc attempt has a new zero-arc charge. Current outgoing pointers therefore give $O(M)$ dynamic-tree events. Check initial conservation once; path and cycle subtractions preserve it. Finish with a sparse zero-residual scan.

For endpoint queries, synchronize two more forests with the same links and cuts. Give source arcs $s \to u$ tag u in one, sink arcs $w \to t$ tag w in the other, and all other arcs a larger sentinel. Each s–t path has exactly one tag of each kind. Its two path minima return the endpoints in $O(\log(2M))$. Only the flow forest receives path additions. These are precisely the operations of [ST83, Section 2]; no represented path is traversed. □

A.4 A uniform projection event and finite game chronology

The flow response may depend on the projected field, so progress must hold for every valid response to the same signs. Put $b_\rho = \log_2(2+\rho)$, $\vartheta = 2^{-60}/b_\rho$, and $B_c = 64 b_\rho$.

Lemma A.6. *If $V_A\|y - Xg\|_\infty^2 \leq 2^{-20}\|Z\|_r^2$, with probability at least $1/16$ over fresh signs, every candidate response allowed above loses at least $\vartheta\|Z\|_r^2$, including smoothing loss.*

Proof. Unmarked norms of fields and embeddings are r-weighted; norms of individual row vectors are Euclidean. If smoothing already loses at least $\vartheta\Phi$, where $\Phi = \|Z\|_r^2$, the claim holds. Otherwise use the harmonic extension of X through the fixed augmented graph. As proof-only randomness, take the first $G = \min\{v, \rho\ln(1/\vartheta)\}$ vertices of a uniform permutation of the original child. At each active row, let c_u be the harmonic-extension value at the first bank vertex adjacent to it in C, or zero if none. For each fixed active row, the probability of being uncovered is at most ϑ; conditional on coverage, the chosen neighbor is uniform among its neighbors in C. Since $r_u/d_C(u) \leq 2$, the harmonic energy bound (46) and $\gamma^{-1} < 3/2$ give

$$\mathbb{E}\delta_X \leq 2(\|Z\|_r^2 - \|X\|_r^2) + \vartheta\Phi \leq 3\vartheta\Phi, \qquad \delta_X = \sum r_u\|X_u - c_u\|^2.$$

The factor two counts both orientations of each original edge; original energy is bounded by augmented harmonic energy. The bank has at most 2^{4b_ρ} centers including zero. Rademacher tails and a union bound give simultaneously $|\langle c_\star - c'_\star, g\rangle|^2 \leq B_c\|c - c'_\star\|^2$ for all pairs, except with probability below $1/1000$. Markov gives

$$\delta_X \leq \Phi/(2^{24}B_c), \qquad \delta_g := \sum r_u\langle X_u - c_u, g\rangle^2 \leq \Phi/2^{24}$$

outside total probability below $1/48$, including the pair event: the two Markov failure bounds are at most $3\vartheta 2^{24}B_c$ and $3\vartheta 2^{24}$. Meanwhile $\mathbb{E}\|Xg\|^4 \le 3(\mathbb{E}\|Xg\|^2)^2$, so Paley–Zygmund gives $\Pr\{\|Xg\|^2 \ge \|X\|^2/2\} \ge 1/12$. The intersection has probability greater than $1/16$.

For any response define the positive form $\mathcal{F}(Y) = E_R(Y) + \sum_S r_u\|Y_u\|^2$; its normalized norm is at most three. Expand $X = c_\star + (X - c_\star)$ twice. Pairwise center distortion, the inclusion of zero as a center, and (50) give, simultaneously for every R, S,

$$\mathcal{F}(Xg) \le 4B_c\mathcal{F}(X) + 12B_c\delta_X + 6\delta_g \le 8B_c\Delta_X + 12B_c\delta_X + 6\delta_g. \tag{53}$$

Writing $y = Xg + e$ and using (51) gives

$$\mathcal{F}(Xg) \ge 2^{-13}\|Xg\|^2 - (3 + 2^{-12})\|e\|_r^2.$$

On the displayed events this is more than $2^{-15}\Phi$ before the two center-error terms; those total at most $18\Phi/2^{24}$. Thus $\Delta_X \ge \Phi/(2^{20}B_c) > \vartheta\Phi$. Marginalizing the proof-only bank proves the stated sign probability; the bank is never constructed and makes no original access. □

Committed rounds and validation. Fix cut conductance χ and a volume-loss reserve $S \ge 1$. Set

$$\tau_0 = 2^{-12}, \qquad t = 16\vartheta^{-1}\log_2(P_{\max}/\tau_0) + 1, \qquad \beta = 2^{\log_2(2/\chi)}, \tag{54}$$

$$K = t(\beta + 3), \qquad \phi = 2^{-\log_2 \max\{2048, 16K\}}, \qquad \eta = \beta\phi/(128S), \quad \epsilon \le 1/(4096K). \tag{55}$$

Let the diagonal frame be $F_0(u,u) = 2^{-\log_4 r_u}$, so $1/(4r_u) \le F_0(u,u)^2 \le 1/r_u$. Keep implicitly

$$Z = P_A Q F F_0, \tag{56}$$

Here P_A subtracts the r-weighted mean on active rows, and Q contains committed smoothers, couplings and deletion masks, retaining all original columns. Intermediate centerings may be omitted: Markov factors preserve the active constant, deletion maps it to the next constant, and final P_A kills it. Thus (56) equals centered replay, Q is nonnegative substochastic, and contraction gives $\Phi \le P_{\max}$.

Candidate signs find progress; independent validation signs measure it before a commit. Permit at most t committed rounds and one additional terminal validation block. At each of these at most $t + 1$ blocks perform the following operations in order.

1. Freeze the active state and $D_0 = C \setminus A$. Generate $s = 16\log_2(16(t+1)/\delta_c)$ independent candidates from this same state.

2. If some seed has $r(D_0 \cup S_i) \ge \eta V$, return that particular union as a cut. Never accumulate rejected seeds.

3. Otherwise freeze every candidate map. Independently draw a shared validation bank in the smallest odd number J of groups of 128 signs satisfying $J > \log_2(16(t+1)(s+1)/\delta_c) + 1$. Median group means estimate Φ and each ideal candidate loss Δ_i.

4. Put $a_{\mathrm{stat}} = \vartheta\tau_0/16$, $U = 2(\widehat{\Phi} + a_{\mathrm{stat}})$, and $L_i = \max\{0, (2/3)(\widehat{\Delta}_i - a_{\mathrm{stat}})\}$. Stop if $U \le 4\tau_0$. Otherwise report failure before any further commit if t rounds have already been committed. If fewer have been committed, select the largest L_i, reporting failure if it is below $\vartheta U/16$, and commit only that candidate.

Frozen candidates are contractions, so each ideal loss is a nonnegative quadratic form with Rademacher variance at most twice its squared mean. Chebyshev gives $[1/2, 3/2]$ accuracy with failure at most $1/16$ per 128-sign group, and the median fails with probability below 2^{-J}. Including Appendix A.5's arithmetic error gives

$$\Phi \le U \le 3\Phi + 4a_{\mathrm{stat}}, \qquad \Delta_i/3 - 4a_{\mathrm{stat}}/3 \le L_i \le \Delta_i. \tag{57}$$

When $\Phi > \tau_0$, all s candidates miss the uniform progress event with probability at most $e^{-s/16}$. If a candidate has that progress property, it passes the selection threshold, and the committed round loses at least $\vartheta\Phi/16$. The choice of t forces $\Phi \le \tau_0$, where the stop test succeeds. No relative accuracy is required in this small-potential case.

Lemma A.7. *Except with the allocated failure probability, the game returns either a cut of true reference conductance at most χ, whose two sides have reference volume at least ηV, or an active near-expander A_1 with*

$$r(A_1) > (1-\eta)V, \quad \partial_C(A_1) \le r(C \setminus A_1)/\beta, \quad \partial_C(U) \ge \phi r(U) \quad (r(U) \le r(A_1)/2). \tag{58}$$

It uses no original row scans. Its sparse cost is $M^{1+o(1)}$ when $\chi^{-1}, S, \epsilon^{-1}, \log(2+\rho)$ are bounded by fixed powers of $\log n$, the ambient quantities v, V, ρ have the prescribed polynomial range, and $\log(1/\delta_c) = O(\log n)$. In particular, ρ itself need not be polylogarithmic.

Proof. For a failed-flow seed use the full original set $U \cap C$ in (49). Hub elimination and (45) show its true cut is at most the augmented cut. Inactive originals already belong to D_0. Submodularity therefore increases the boundary of $D_0 \cup S$ by at most $r(S)/\beta$. At a triggering union its volume is between ηV and $(5/8 + 3\eta/8)V$; the complement is at least $3(1-\eta)V/8$. Since $\eta \le 1/9$, this complement is at least ηV, and the volume ratio is at most two. Conductance is at most $2/\beta \le \chi$.

At potential stopping, put $r_0 = r(C \setminus A_1) < \eta V$. Boundary is at most r_0/β, hence $b := \partial_C(A_1)/r(A_1) \le \phi/64$. For a final set $U \subseteq A_1$, $s_0 = r(U) \le r(A_1)/2$, let $x = \mathbb{1}_U - (s_0/r(A_1))\mathbb{1}_{A_1}$. It is centered on the full child, $\|x\|_r^2 \ge s_0/2$, and $\lambda_C(x) \le 2\partial_C(U)/s_0 + b$ (expand its three edge classes). If the cut ratio is at most $1/2048$, Lemma A.3 and $\Phi \le 4\tau_0$ give

$$\|P_{A_1} Q x\|_r \le (2\sqrt{\Phi} + 1/16 + \sqrt{8\lambda_C(x)})\|x\|_r < \|x\|_r/2.$$

Here the frame gives $\|P_{A_1} Q F\|_{\mathrm{HS},r} \le 2\sqrt{\Phi}$. For this comparison, view Q on all rows of C, padding deleted output rows with zeros. Since Q is substochastic and x takes values in an interval containing zero and of length one, $|(Q-I)x| \le 1$. Orthogonal centering on A_1 and $P_{A_1}x = x$ imply $\|(Q - I)x\|_r \ge \|x\|_r/2$, hence $\|(Q-I)x\|_{r,1} \ge s_0/8$. Every deletion mask fixes this final-support vector. Suffix telescoping of $Q - I$, followed by weighted ℓ_1 contraction, tests the same full original sets U, A_1 in every earlier factor. Equations (47) and (52) give

$$s_0/8 \le K[\partial_C(U) + (s_0/r(A_1))\partial_C(A_1) + 4\epsilon s_0].$$

Thus the ratio is at least $\min\{1/2048, 1/(8K) - b - 4\epsilon\} \ge \phi$. This proves near-expansion. For the work bound below, $\vartheta^{-1} = O(\log(2+\rho))$ and

$$t = O\big(\log(2+\rho)\log(P_{\max}/\tau_0)\big) = O\big(\log^2(2+\rho)\big).$$

The proof-only bank and scalar $P_{\max}$ require no list of ρ objects. Other reciprocal parameters are polylogarithmic; vertex and volume factors affect only $O(\log n)$-bit precision. The static construction chooses admissible dyadic accuracies. □

A.5 Finite arithmetic, sparse solvers, and work

A.5.1 Finite Bernoulli skipping

For a row suffix of length $N \le n$, invert the geometric CDF $\Pr(T > k) = (1-p)^k$ by binary search on $1, \dots, N+1$, where $N+1$ means exhaustion, using a uniform b-bit dyadic U. Repeated squaring on mesh ζ evaluates each power with error $O(n\zeta)$, by telescoping products in $[0,1]$. Coupling U to a continuous uniform succeeds unless it lies within $2^{-b} + O(n\zeta)$ of one of the N CDF boundaries. Thus per-call error is $O(n2^{-b} + n^2\zeta)$; choose $b = O(\log n)$ and $\zeta = 2^{-O(\log n)}$ for the polynomial horizon. Each selected or exhausted position costs $O(\log^2 n)$ arithmetic and reads no unselected label. The coupling transfers the ideal Chernoff work cap to this finite sampler.

A.5.2 Uniform replay error and centering

Every ideal replay of $F_0 g$ has D-norm at most $\sqrt{v}$. Choose final coordinate error

$$a_* \le \frac{\vartheta \tau_0}{2^{16} v (V+1)}. \tag{59}$$

Budget $a_*/(4t+8)$ in D norm per factor. Contractions and exact centerings add these errors; $r_u \ge 1$ gives coordinate error at most a_*, meeting Lemma A.6 for $\Phi \ge \tau_0$. Squared-norm error is at most $2\sqrt{v}\sqrt{V}a_* + Va_*^2$, small enough for a_{stat} even when subtracting two norms.

After each sparse action round to a fixed dyadic grid, then center exactly. This introduces only the integer volume denominator V_A, discarded at the next rounding. Round *each coefficient–coordinate product before summing*, so denominators never accumulate. Polynomial magnitudes, $2^{O(h)}$ filter amplification and polynomially many summands permit a common $O(\log n)$-bit grid fine enough for all errors.

A.5.3 A finite grounded inverse

Use Lemma 6.1 and its grounded-block construction with ambient size $n+2$, including the child, hub, and ground vertex. The sample weights $1/p$, dyadic hub weights ϵr_u, and ground-edge sums share an $O(\log n)$-bit dyadic mesh and polynomial positive ranges. Together with the sparse count cap and the preceding common-grid rule for replay right-hand sides, these satisfy the lemma's input hypotheses.

The lemma supplies fixed symmetric analytical inverse maps and finite evaluations in $M^{1+o(1)}$ work with $O(\log n)$-bit words. Choose its accuracy for the per-action allowance $a_*/(4t+8)$ in (59), and allocate factorization and random-work failure budgets over the predetermined structural call horizon. Its zero-pivot rule, operand guards, and work caps apply with the same first-failure convention. The lemma and its grounded-block reduction are proved independently of core construction.

For the actual game block L_{BB}, a path to ground gives $\lambda_{\min}(L_{BB}) \ge w_{\min}/(v+1)^2$. Its sparse exact residual certifies energy error using this bound. If $C_0 = -L_{AB}$, its row sums are at most $\gamma^{-1} r_A$. Let $D_{B,A}$ be the diagonal matrix of total augmented edge weights from each vertex of B to A. Weighted Cauchy gives

$$C_0^T D_A^{-1} C_0 \preceq \gamma^{-1} D_{B,A} \preceq \gamma^{-1} L_{BB}.$$

Thus squared K_A output error is at most γ times grounded solve energy error. Sparse residual checks certify replay precision.

A.5.4 Exact flow on small sparse children

Common-scaled capacities have $O(\log n)$ bits. Give all network arcs cost zero and add a sink-to-source return arc of cost -1 and capacity Q, the sum of source-arc capacities, with zero demands. Zero is feasible; adding or removing the return arc identifies minimum cost with negative maximum-flow value. Apply [vdBCP+23, Theorem 1.1] with U the maximum of two and all capacities, including Q, and cost bound $C = 2$. Its $M^{1+o(1)} \log U \log C$ bound includes capacity scaling, so $\log U = O(\log n)$ gives $M^{1+o(1)}$ work even for tiny children. Check capacity, conservation, integrality and the final residual cut by sparse scans; handle empty and constant-size cases directly.

After local renumbering the input has $z = O(M \log n)$ bits. Keep input/output integers in $O(\log n)$ bits and simulate the $O(\log^{O(1)} z)$-bit fixed-point words of [vdBCP+23, Section 2] with $n^{o(1)}$ ordinary words and work per operation. Store any wide addresses as multiword keys in a balanced dictionary pointing to ordinary records. There are at most as many keys as operations, so lookups cost $n^{o(1)}$ and allocated memory is polynomial. This preserves the work bound and the exact integral interface of [vdBCP+23, Theorem 1.1 and Lemma 4.1]. Lemma A.5 supplies route decomposition.

A.5.5 Total game work and failure probability

Keep at most $t + 1$ active-set factors for prefix replay. Candidates in each block share the current factor; their final seed deletion needs no new factor. Every rejected candidate pays one replay, smoothing, sort, exact flow, and compressed route decomposition. Independent validation replays the prefix and each frozen suffix. A sufficient total bound is

$$O((s + 128J)(t + 1)^2 M^{1+o(1)} + s(t + 1)M^{1+o(1)}). \tag{60}$$

This includes filter application and every factor construction. Predeclare a polynomial call horizon and abort before any work cap. Allocate failure among probes, sample concentration and skipping coupling, factors, candidate progress, validation, and random-integer tails. New tapes follow a fixed input state; bound the first failed operation conditionally on its past and union over the capped horizon. No randomness is conditioned on future successful execution.

A.6 A weighted fixed-birth pruner with atom-count work

Sample atoms can have large weight $1/p$, so work must depend on their count. We first establish the mass and cut bounds, then implement them with a dynamic forest. Within a pruning theorem and its proof, V denotes its birth vertex set.

Theorem A.8. *Let a weighted undirected birth graph have m stored atoms with positive integer weights, positive integer birth degrees b, volume B, and reference conductance $\alpha \le 1/2$, an inverse power of two. An arbitrary adaptive sequence of distinct nonloop birth atoms may be fed as genuine deletions. Each feed deletes that entire atom at its full birth weight, at most once; partial decrements of one atom are outside this interface. With fed weight k, a deterministic structure removes birth volume at most $16k/\alpha$. Its surviving live graph satisfies*

$$\partial_{\text{live}}(X) \ge (\alpha/4)b(X) \quad \text{for } X \subseteq A,\ b(X) \le b(A)/2.$$

For

$$H = 2 + \ln(2B)/\ln(1 + \alpha/8),$$

total work is $O((m+|V|)H \log(2m+2|V|))$, independently of numeric weights except their bit length and H. No fed-weight budget is required for correctness or the volume bound.

Implement height by $2 + 16/\alpha \log_2(2B)$, with binary logarithms computed by bit lengths and comparison. For the density-height bound below at $\theta = \phi/8$, use $2/\theta$ and $\log_2 \max\{1, M/(\theta d^2)\}$.

A.6.1 State and the mass proof

Put $c_\star = 2/\alpha$. Initially $A = V$, $f = 0$, and all levels are zero. Private vertices are the survivors A; private edges include every birth edge induced by A, including deleted ghosts. Keep antisymmetric integral flow $|f(u,v)| \le cw_{uv}$, and

$$\sigma_A(u) = c_\star\, w(E(u, V\setminus A)) + 2c_\star \deg_D(u), \quad m(u) = \sigma_A(u) + \sum_{v\in A} f(v,u), \quad e(u) = \max\{m(u) - b_u, 0\}. \tag{61}$$

Here $\deg_D$ counts all atoms fed so far at their birth weights, including atoms whose other endpoint is no longer private. Feeding weight w adds $2cw$ at each endpoint still private; charge $4cw$ even if one is gone. Thus charged injection is $J = 4ck = 8k/\alpha$. Maintain

$$m \ge 0, \quad l(u) > 0 \Longrightarrow m(u) \ge b_u, \quad cw_{uv} - f(u,v) > 0 \Longrightarrow l(u) \le l(v) + 1, \quad 0 \le l \le H. \tag{62}$$

Once absorbed, sink mass never leaves a surviving vertex.

Let $S_i = \{l \ge i\}$ and let z_i be crossing weight between levels i and $i-1$. If $S_H \ne \varnothing$, some nonempty S_i has $z_i \le (\alpha/8) b(S_i)$: otherwise successive volumes grow by $1 + \alpha/8$ from at least one, exceeding B.

Remove this upper set simultaneously, replacing crossing edge $u \to v$ by source cw at survivor v. Its mass gain is $cw - f(u,v) \ge 0$, zero across a level gap of at least two by residual saturation. Adjacent-level edges create at most $2cw$ each, totaling at most $b(S)/2$, while removed positive-level sinks destroy at least $b(S)$. For cumulative creation C and removed birth volume R, every completed prefix satisfies

$$C \le R/2, \qquad R \le J + C, \qquad R \le 2J = 16k/\alpha. \tag{63}$$

At a zero-excess checkpoint, for smaller-reference $X \subseteq A$ let x be its private cut, y its birth boundary outside A, and $z = \sum_{u\in X} \deg_D(u)$. Birth expansion and feasible mass give

$$x + y \ge \alpha b(X), \qquad cy + 2cz \le b(X) + cx.$$

Since $1/c_\star = \alpha/2$, substitution gives $x - z \ge \alpha b(X)/4$. Deleted crossing weight is at most z, proving the fixed-reference live guarantee.

A.6.2 The dynamic forest algorithm

Batch pushes along a child-to-parent forest. Each arc satisfies $l(u) = l(v) + 1$ and stores residual capacity $cw_{uv} - f(u,v)$. Lazy zeros may remain linked. Before cutting, query and materialize its flow; unlinked flows are explicit. Row queries test linkage in *either* direction and read linked values from the tree, never a stale explicit copy.

Each vertex has a row cursor reset only on relabel, incoming-child list, and one active queue handle. Validate private membership and excess before use; clear stale handles at excess-interval end or removal.

Choose any active v. Purge zero residual arcs on its root path, repeating root finding after cuts, and let w be its current root. Process v only while it has excess.

1. If $m(w) < b_w$, push $a = \min\{e(v), b_w - m(w), \text{path residual minimum}\}$ along the root path, subtracting a lazily from residuals. Source and root masses change by $-a, +a$; internal masses stay fixed and the root gains no excess. The path is nonempty since v has excess.

2. For a full root sink, scan its current row for positive residual $w \to z$ with $l(w) = l(z) + 1$. If found, link it, leave the cursor there, and repeat, purging extended-path zeros. The trees are distinct: a z–w path would force $l(z) > l(w)$.

3. If no admissible arc remains, raise $l(w)$ by one, reset its cursor, and materialize and cut *all* incoming arcs. Their children now have the parent's level; advance their rejected slots. Every residual neighbor had level at least old $l(w)$, so raising preserves (62). This applies even to a full root with zero excess.

4. On reaching height H, immediately prune a qualifying upper set and restart the queue; otherwise continue with v.

For pruning, mark the whole set and collect crossing atoms. Query, materialize and cut *all linked crossing atoms before any survivor source update.* Save their flows, then update sources, remove vertices and internal links, and update handles. No push interrupts this batch; only its completion is a checkpoint. Strict descent prevents a surviving child from pointing into the removed upper set.

A.6.3 Cursor validity and the atom count

At a fixed sender level, a rejected positive slot's neighbor is already as high and only rises. Saturated slots reopen only by a reverse push, which also puts the neighbor too high; parent relabels have the same effect. Removal is permanent. Thus each directed slot is scanned and selected $O(1)$ times per level: scans, links and value queries total $O(mH)$, relabels $O(|V|H)$, and each link is cut once.

Every positive push saturates a link, first fills a root sink, or ends its source's excess interval; break ties arbitrarily. A saturated link is purged before reuse. Pushes leave internal masses unchanged and create no destination excess, so only feeds and prune-boundary additions start excess intervals. Each atom is fed and crosses a prune boundary at most once. Hence there are $O(m+|V|)$ intervals and sink fills, and $O(mH + |V|)$ pushes. Charge root queries to the next push, scan, link, cut, relabel or interval completion.

Search levels downward to the first successful test, scanning only ultimately removed rows: total row work is $O(m+|V|)$ and empty-level work $O(|V|H)$. Materialize each crossing atom once. Dynamic trees add one logarithm per operation, proving Theorem A.8 and termination even behind full zero-excess roots.

A.7 Adaptive degree-bias repair and the local induced core

Return to the sampled child C, with $V = r(C)$. Repair uses the global raw-adjacency bound to convert a sampled induced expander into an original induced core. We first prove this repair step; Lemma A.10 then obtains its input by trimming the game's near-expander.

Lemma A.9. *Let $0 < \alpha \le 1/2$ be an inverse power of two, with α^{-1} bounded by a fixed power of $\log n$. Let C be simple, $d_C \le r \le 2d_C$, and let a sample satisfy the raw error in (42) and $\widehat{d}_C \le r$. A sample-adaptive A_0 has $r(A_0) \ge 3V/4$, and its sampled induced graph completed to degree $2r$ is an α-expander. Suppose original degree estimates have error at most $\delta_0 r_u/4$ and are maintained by exact public row decrements. If*

$$0 < \delta_0 \le \alpha/100, \quad \log n \ge \log_2(2v), \quad \epsilon \le \alpha\delta_0/(256S\log n), \quad S \ge 1,$$

then sparse work $M^{1+o(1)}$ plus scans of first publicly removed original rows produces an original induced $\alpha/12$-core R, removing reference volume at most $576\epsilon^2 V/(\alpha\delta_0^2)$.

Proof. Give Theorem A.8 the fixed birth graph on A_0, with sample edges and loops $2r_u - \widehat{d}_{A_0}(u) \geq r_u$. Initialize public $R = A_0$ and keep it inside the private survivor set. Removing a row from R scans its original root row, independently of private ghost-flow removal.

A public row is BAD if $|\widehat{d}_R(u) - \widetilde{d}_R(u)| > \delta_0 r_u/2$. Both degree counters exclude the reference-completion loops. Choose one BAD row and feed every not-yet-fed birth nonloop edge incident to it. Complete the pruner batch, then publicly remove that row and all privately pruned rows. A forced BAD row may remain private, with its state and loops untouched. Finish all public updates before making the next BAD decision.

Let V_B be reference volume of chosen BAD rows and $V_P = r(A_0 \setminus R)$. Fed weight $k \leq V_B$ since birth nonloop degree is at most r. Private pruning removes r-volume at most $8k/\alpha$, so

$$V_P \leq (8/\alpha + 1)V_B \leq 9V_B/\alpha. \tag{64}$$

This holds for every completed prefix with no small-budget premise.

For the analysis of V_B, put $E = D^{-1/2}(\widehat{A} - A_C)D^{-1/2}$ and order $P = A_0 \setminus R$ by completed batches. A chosen BAD row has true bias greater than $\delta_0 r_u/4$, whose normalized value is its coordinate of $ED^{1/2}\mathbb{1}_{A_0} - ED^{1/2}\mathbb{1}_{P_{\text{earlier}}}$; P_{earlier} contains only rows removed before its decision. The second term on chosen rows is therefore a restriction of the strict block-lower part $T(E)D^{1/2}\mathbb{1}_P$. Neither E nor $T(E)$ is constructed. For every block order,

$$\|T(E)\| \leq \log n \|E\|. \tag{65}$$

To prove this, recursively bisect the ordered blocks. At each depth, lower off-diagonal rectangles have disjoint rows and columns, so their direct sum has norm at most $\|E\|$. Sum over at most $\log n$ depths. This is valid for an order chosen after seeing E.

Consequently

$$(\delta_0^2/16)V_B \leq 2\epsilon^2 V + 2\log^2 n \epsilon^2 V_P.$$

Combine (64) and $\epsilon^2 \leq \alpha\delta_0^2/(576\log^2 n)$ to absorb the last term:

$$V_B \leq 64\epsilon^2 V/\delta_0^2, \qquad V_P \leq 576\epsilon^2 V/(\alpha\delta_0^2). \tag{66}$$

At termination every edge from R to a BAD row still private has been fed. For $U \subseteq R$, $r(U) \leq r(R)/2$, both sides of the private cut have b-volume at least $2r(U)$, where $b = 2r$. Hence the sampled live public cut is at least $(\alpha/2)r(U)$; use only $(\alpha/6)r(U)$. Final degree bias is $|\widehat{d}_R - d_R| \leq \delta_0 r$, and the raw norm gives $|\mathbb{1}_U^T(\widehat{A} - A_C)\mathbb{1}_U| \leq \epsilon r(U)$ for every adaptive U. Degree sum minus internal adjacency yields $|\widehat{\partial}_R(U) - \partial_R(U)| \leq (\delta_0 + \epsilon)r(U)$. Since $\delta_0 + \epsilon \leq \alpha/12$, the original bound follows.

Sparse degree counters and a BAD queue update only affected rows. Each sample edge is fed once. Only first public removal scans an original row; later private removal cannot scan it again. Since C is simple and $r \leq 2d_C$, the weighted birth volume is at most $2V \leq 4n^2$. Its height is therefore $n^{o(1)}$, and Theorem A.8 costs $M^{1+o(1)}$. □

Lemma A.10. *In the near branch of Lemma A.7, set*

$$\begin{gathered}\alpha = \phi/8, \quad \delta_0 = \alpha^2/256, \quad e \leq \min\{1/(4096\log n), \delta_0/16\}, \\ \epsilon \leq \min\{e/4, \alpha\delta_0/(256S\log n), 1/(4096K)\},\end{gathered} \tag{67}$$

all inverse powers of two where appropriate. An original induced core of reference conductance at least $\alpha/12$ remains after removing less than $V/(16S)$ reference volume. Removed root volume, and hence original row work, is less than $w(C)/(2S)$. Other work is $M^{1+o(1)}$.

Proof. Complete the original-vertex rows of the sample-plus-hub graph to $2r$. To trim A_1, let b_u now denote its augmented boundary supply and put $b(S) = \sum_{u\in S} b_u$. Use sparse maximum flow with sink capacity $\phi r_u/2$ and routing-edge capacities equal to their sampled weights. Route only on sampled edges induced by original vertices of A_1; the hub and excluded originals are boundary sources, never transit. Since $B_1 = \partial_C(A_1) \le r(C \setminus A_1)/\beta$, total source is at most $B_1 + 2\epsilon V$. For the residual set S_1 reachable from unsaturated sources (empty if none), the mincut gives

$$\widehat{\partial}_{A_1}(S_1) + (\phi/2)r(S_1) \le b(S_1), \qquad r(S_1) \le 2(B_1 + 2\epsilon V)/\phi.$$

Remove S_1 and replace its saturated outward arcs by boundary supply of the same weights at their surviving endpoints. The survivors' old sources were saturated, so the remaining flow routes all new boundary supply to the same sinks. Put $A_0 = A_1 \setminus S_1$.

For smaller-r $U \subseteq A_0$, let x be its sampled internal cut and y its full augmented boundary outside A_0. Near-expansion and augmented domination give $x + y \ge \phi r(U)$; the retained flow gives $y \le (\phi/2)r(U) + x$. Hence $x \ge (\phi/4)r(U)$, giving conductance $\alpha = \phi/8$ after completion to $2r$.

Loss so far divided by V is at most

$$\frac{\beta\phi}{128S} + \frac{1}{64S} + \frac{\delta_0}{512\log nS},$$

so $r(A_0) \ge 3V/4$. Discard the hub and use the integral sample with completion loops $2r - \widehat{d}_{A_0} \ge r$. The guard estimate error is $ew \le 4er \le \delta_0 r/4$, and exact public decrements preserve it.

Initialize A_0 estimates by scanning committed removed root rows in $C \setminus A_0$ once and decrementing retained neighbors via membership flags. The displayed loss pays game removals and trimming; rejected seeds trigger no scan. Stamp scanned rows to prevent scanning these original rows again. Lemma A.9 removes at most

$$576\epsilon^2 V/(\alpha\delta_0^2) \le 9\alpha V/(1024\log^2 nS^2).$$

The sum is less than $V/(16S)$. Each removed row costs $w_u \le 4r_u$; $V \le 2w(C)$ gives the root-volume and row-work bound. □

A.8 Delayed original-label classification and static cover

The recursion first finds vertex sets. Assigning original labels only after the tree is fixed keeps its depth out of the dense scan cost.

For $q_j > \tau$, put $S_j = AD_j^2$ and $t_j = q_j/(nS_j)$. Peel current degrees below t_j, exporting at most $nt_j = q_j/S_j$ labels in one original scan. The nonempty remainder has exact root degrees w, minimum $d \ge \max\{1, t_j\}$, and frozen arrays. Let q_0 be its edge count and $W_0 = 2q_0$ its volume. Set common root parameters

$$\begin{gathered} d_* = \max\{1, d/4\}, \quad \rho = \max\{1, n/d_*\}, \quad J_0 = 1 + \log_2 \max\{1, 8W_0/d^2\}, \\ \chi = 2^{-\log_2(256 S_j J_0)}, \qquad S = 32S_j. \end{gathered} \tag{68}$$

Choose each accuracy as the largest inverse power of two satisfying all its upper bounds, hence within a factor two of their minimum. Before child tapes, choose (68), then $b_\rho, \vartheta, P_{\max}$ and (54)–(55), then $\alpha = \phi/8$, $\delta_0 = \alpha^2/256$, and finally e and ϵ in that order from (67). Thus ϕ is independent of the later ϵ, and reciprocal accuracies have fixed polylogarithmic exponents. Use these parameters and the same sampling probability at every node of the binary refinement tree. Stop before calls with fewer than $d/8$ vertices; otherwise guard, stopping if empty, then run the game and bootstrap

its near branch. CUT recurses on both surviving sides; CORE exports removed vertices and terminates. Fresh tapes follow the fixed child, and the guard supplies the common d_*.

By (39), each accepted cut has both root-volume sides at least $\eta w(C)/8$, and true boundary at most twice χ times its smaller root-volume side. Enforce the stated volume-balance test even on a bad-sample execution.

Lemma A.11. *The next explicit residual has $q_{j+1} \le 9q_j/(8S_j) \le q_j/2$ for sufficiently large A. This phase uses $O(q_j + n)$ full original scan work and $n^{1+o(1)}$ probes and sample positions. Its sparse work is $n^{1+o(1)}$.*

Proof. Let P_g be all guard-removed vertices, F_g the distinct labels whose other endpoint is still in the current node when their first guard endpoint is removed, and B_g all accepted-cut labels. Every label incident to P_g belongs to $F_g \cup B_g$: a terminal CORE cannot be an ancestor of a later recursive node. The removal bound in Lemma A.1 and endpoint counting give

$$|F_g| \le 3w(P_g)/8, \quad w(P_g) \le 2(|F_g| + |B_g|), \quad w(P_g) \le 8|B_g|, \quad |F_g| \le 3|B_g|.$$

Charge each cut to its smaller root-volume side. Each vertex's next charged parent has at most half the preceding volume. Split parents have volume at least $d^2/8$, so the sum of smaller-side volumes is at most J_0W_0 and $|B_g| \le 2\chi J_0W_0$. For a small leaf U, simplicity gives $e(U) \le w(U)/16$ and root boundary at least $7w(U)/8$. Summing those boundaries counts $B_g \cup F_g$ at most twice, so adding small-leaf internal labels multiplies that charge by at most $8/7$. Terminal CORE input sets are disjoint; their root-volume losses total at most $W_0/(64S_j)$. The full post-peel residual is therefore at most

$$\frac{4|B_g|}{1 - 1/8} + \frac{W_0}{64S_j} \le \left(\frac{8\chi J_0}{1 - 1/8} + \frac{1}{64S_j}\right) W_0 < q_0/(8S_j).$$

Including first-peel exports proves the size recurrence.

Uniform balance bounds depth by $H_0 = 2 + (8/\eta)\log_2(2W_0/d) = n^{o(1)}$, even with guards before cuts. Consequently $\sum_i |A_i| \le nH_0$ and $\sum_i w(A_i) \le W_0H_0$. Probes cost $k\sum_i |A_i|$. Common sampling probability satisfies $p \le C \log n/(\epsilon^2 d_*)$ even when clamped to one, and $d_* \ge q_j/(4nS_j)$. Together with the sample-cost bound in Lemma A.2, this gives

$$p\sum_i w(C_i) \le CnS_j \log nH_0/\epsilon^2, \qquad \sum_i M_i + \text{all probe and stream work} \le n^{1+o(1)}.$$

This includes discarded sample positions and exhausted streams, not only accepted sample edges. Stamped local IDs avoid an ambient n clearing per call. Sum (60) and the sparse bootstrap cost as $(\sum_i M_i)^{1+o(1)}$.

Guard scans occur once and terminal export allowances are disjoint. After fixing the tree, classify each root label *once*: equal retained-core endpoint flags install it, otherwise it enters the residual. This builds degrees and direct handles in $O(q_j + n)$, without routing labels through tree depth. □

Repeat on the residual until size τ. Lemma A.11 gives $\sum_j q_j \le 2q$ and $O(\log(2 + q/\tau))$ phases. Cores are vertex-disjoint within phases and globally edge-disjoint, induced in their phase residual; total vertex population is $n^{1+o(1)}$. The $n^{1+o(1)}$ probe and sample touches per phase are $O(q_j)$ for $q_j > n^{9/8}$, by asymptotic absorption of any fixed logarithmic power.

From $\rho \le 4n^2S_j/q_j$ and (68), $J_0, \log(2 + \rho) = O(D_j)$. The game has $t = O(\log^2(2 + \rho))$, hence

$$h_j^{-1} = O(\phi^{-1}) = O(S_jJ_0 \log^2(2 + \rho)) = O(D_j^5).$$

The guard and first peel give retained reference minimum $q_j/(4nS_j)$. Appendix A.5's first-failure bound covers adaptive refinement and residual phases, completing the static portion of Theorem 4.1.

A.9 Unweighted fixed-birth pruning at a density-dependent height

Simplicity and minimum degree give installed cores a smaller pruning height and a charge per genuine deletion. In the next lemma, $\theta > 0$ is the level-growth threshold; the pruner uses $\theta = \phi/8$.

Lemma A.12. *In a simple private graph whose degrees are bounded by positive integer references r, with $\sum r \le M$ and $\min r \ge d$, give vertices levels $0, \dots, H$. Put $S_i = \{l \ge i\}$ and let z_i count edges between levels i and $i-1$. If $S_H \ne \varnothing$ and*

$$H \ge 2 + \ln \max\{1, M/(\theta d^2)\} / \ln(1+\theta),$$

some nonempty S_i, $i \ge 1$, has $z_i \le \theta r(S_i)$. It can be found by scanning only rows ultimately removed.

Proof. If all tests fail, simplicity gives $\theta d|S_H| \le \theta r(S_H) < z_H \le |S_H|\,|\{l = H-1\}|$. Thus the next level has more than θd vertices and $r(S_{H-1}) > \theta d^2$. Each further failed level has $r(S_{i-1}) > (1+\theta) r(S_i)$, contradicting total M at this height. Scan levels downward and stop at the first success; only rows in its upper set have been read. Empty-level bookkeeping costs $O(H)$. The d^2 argument requires simplicity; Appendix A.6 uses a different work proof for weighted parallel atoms. □

Theorem A.13. *Let a simple unweighted birth graph have actual degrees $r > 0$, volume $M = 2q_b$, minimum degree d, and conductance $\phi \le 1/2$, an inverse power of two. Define*

$$\Delta = 1 + \ln\left(1 + \frac{M}{(\phi/8)d^2}\right), \qquad H = O(\Delta/\phi)$$

using Lemma A.12. For any adaptive sequence of genuine deletions there is a deterministic structure with live reference conductance at least $\phi/4$, total work

$$O(q_b + |V| + H + U\Delta/\phi^2),$$

and $O(U/\phi)$ exported live labels. All prices and references are fixed at birth. Retire at the first $U \ge \phi M/128$, exporting every remaining live label.

Proof. Use the ghost source and mass state (61) with $b = r$, unit edge weights, and $c_\star = 2/\phi$. A genuine deletion removes public handles but retains private flow and adds $2c_\star$ at each endpoint; its injection is $4c_\star = 8/\phi$. Preserve (62). Choose the minimum-level vertex with excess and level below H. On an admissible arc $u \to v$ push $a = \min\{e(u), c_\star - f(u,v), r_v\}$. Minimum-level choice makes the recipient nonexcessive, so its new excess is at most r_v. Sending only excess preserves filled sinks. When no admissible arc remains, relabel and reset the cursor; the weighted forest's rejection argument applies unchanged.

When only height-H excess remains, apply Lemma A.12 with $\theta = \phi/8$ and mark its entire upper set before removal. Replace each crossing private edge by source $c_\star$ at its surviving endpoint. The source-and-mass calculation leading to (63) gives, for injection $J = 8U/\phi$, cumulative creation C, and removed birth volume R, $C \le R/2$, $R \le J + C$, and $R \le 2J$. Scan each removed birth row once. At most R live incident labels are exported, each once; historical ghosts are never exported. The feasible-cut calculation following (63) proves the $\phi/4$ live guarantee without subtracting deleted incidences twice.

Put $B = J + C \le 2J$. Each ever-active vertex absorbs r_u until removal, so $\sum_{\text{ever active}} r_u \le B$ and scans/relabels cost $O(HB)$. A push of a lowers $\Psi = \sum e(u)l(u) \ge 0$ by at least a, since

its nonexcessive recipient is one level lower. For cumulative direct injection/creation I_u, we have $e(u) \le r_u + I_u$: received pushes start at zero excess and add at most r_u. All relabels together increase Ψ by at most $H \sum_{\text{ever active}} r_u + H \sum I_u \le 2HB$; direct additions raise Ψ by at most HB, and removals only lower it apart from charged creation. Thus pushed integral mass and push count are at most $3HB$.

Use per-level active lists. Pushes lower the minimum by at most one; positive integral additions lower it by at most H, charged to their amount. Pair upward scans with these decreases and initialization: queue work is $O(HB + H)$. At most $R \le 2J$ upper-set removals cost $O(HJ)$ in empty levels. Total update work is $O(HJ + H)$, pathwise under adaptive deletions.

At retirement $q_b = M/2 \le 64U/\phi$, paying the final scan and exports. Charge a frozen batch in full as genuine input at old birth prices, even if it triggers retirement; its exports are not new deletions in this birth. □

A.10 Installing the static cores and closing the interface

The construction reference was chosen before the core was known. Installation replaces it by actual original birth degrees and freezes those degrees for pruning.

In a phase-j core, each retained reference is at least $m_j = q_j/(4nS_j)$. In a core with at least two vertices, singleton cuts have reference at least m_j on both sides, so actual birth minimum degree is at least $h_j m_j$. Actual degree is no larger than the static reference, so actual-degree conductance is at least h_j. Omit singleton cores with no original label. Round ϕ down within a factor two of h_j and freeze it. Birth volume is at most $2q_j$; hence

$$\frac{M}{(\phi/8)d^2} \le C\frac{n^2 S_j^2}{h_j^3 q_j}, \qquad \Delta = O(D_j).$$

Theorem A.13 now gives Theorem 4.1's deletion and export prices.

B Table of symbols

Table 1: Symbols, in order of first appearance.

Symbol	Meaning
n, m, r	Vertices, labels, and rank $r = n - 1$ of the working graph
P	Central potential exponent, the integer fixed in Section 1.1
C, c_0	Magnitude exponent of the input data and confidence exponent
$\mathbf{B}$	Incidence matrix
s	Density $s = r/m$
$\mathbf{f}$	Stored flow
$\boldsymbol{\lambda}$	Tracking weights
$\mathbf{y}$	Vertex potential
μ	Barrier parameter
$\mathbf{c}, \boldsymbol{\upsilon}^{\max}$	Cost vector and upper capacity vector
$\mathbf{z}$	Local barrier scale
$\mathbf{w}$	Regularized weight $\mathbf{w} = 2s + \boldsymbol{\lambda}$
$\boldsymbol{\varrho}$	Row scale $\boldsymbol{\varrho} = \mathbf{z}/\mathbf{w}$
$\overline{\boldsymbol{\Xi}}$	Barrier gradient scale $\mathbf{z}(\mathbf{f}^{-1} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-1})$
$\mathbf{D}$	Barrier curvature scale $\mathbf{z}^3(\mathbf{f}^{-3} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-3})$
$\mathbf{H}$	Ordinary coefficient Laplacian $\mathbf{B}^T \textsc{Diag}(\boldsymbol{\varrho}^2)\mathbf{B}$
$\boldsymbol{\tau}$	Edge leverage scores of $\mathbf{H}$
$\mathbf{U}, \boldsymbol{\Pi}$	Scaled incidence $\textsc{Diag}(\boldsymbol{\varrho})\mathbf{B}\mathbf{H}^{-1/2}$ and the projection $\mathbf{U}\mathbf{U}^T$
$\mathbf{L}$	Physical metric for the direction solve and residual norm
$\boldsymbol{\varpi}$	Logarithmic tracking error
$\boldsymbol{\delta}$	Margin of the tracking error to the boundary
$\overline{\boldsymbol{\delta}}$	Stored margin
$\overline{\boldsymbol{\varpi}}$	Positive stored reference scale comparable to $s + \boldsymbol{\tau}$
$\boldsymbol{\varsigma}$	Central error
$\overline{\boldsymbol{\varsigma}}$	Stored representative of the central error
$\boldsymbol{\sigma}$	Relative change of the squared row scale in a frame, $(\boldsymbol{\varrho}'_e/\boldsymbol{\varrho}_e)^2 - 1$
$\mathbf{b}$	Demand vector
d_Z	Margin $1 - \lvert\overline{\boldsymbol{\varsigma}}\rvert$ of the stored central representative
p	Tracking potential exponent, $p = 512$
ψ_q	Potential kernel $(1 - x)^{-q} + (1 + x)^{-q}$
Ψ_R	Tracking potential $\sum_e \psi_p(\boldsymbol{\varpi}_e/\kappa)$
Ψ_Z	Centrality potential $\sum_e \psi_P(\boldsymbol{\varsigma}_e)$
M	Frame scale, $10^5 n^{1/4} \le M \le 2 \cdot 10^5 n^{1/4}$
h	Central step size, within a factor two of $c_h n^{-3/p} P^{-16}$
β_f	Flow step scale h/M
β_s	Sampled step scale $\gamma\beta_f$
B_*	Per-frame variation scale $\beta_f\sqrt{m/s}$
W_e	Frozen structural update price of an edge
a_{jump}	Fixed scalar jump allowance of the owner rates
k_P	Coefficient of the central correction
$G_{Z,e}$	Certified central gradient bound of an edge

Symbol	Meaning
c_Z	Constant of the central correction cap
$\mathbf{F}$	Central correction
$\mathbf{K}_0$	Root-update coefficient $(C+q)/(\theta(1+b_{\mathbf{w}}))$
T_0	Target of the root update
$C_{R,\mathrm{out}}$	Threshold of the tracking potential at a valid pre-state
Δ_D	Direction mesh
μ_{low}	Supplied lower bound on μ over an attempt range
c_{cen}	Rational constant of the central scheduler
$\rho_{\min}$	Dyadic floor of the scheduler radii
τ_0	Residual cutoff of the owner, $\tau_0 \asymp n^{9/8}$
ϕ_{eo}	Owner gate weight of a label occurrence
N_{init}	Number of initial tokens, $O(m)$
N_{sharp}	Sharp token cap $C_0 m B(P, \log\log n)$
T_P, T_H	Approximate inverses of P and of $\mathbf{H}$
$\boldsymbol{\Xi}$	Central ratio $\mathbf{z}(\mathbf{c}-\mathbf{B}\mathbf{y})/(\mu\mathbf{w}) = \overline{\boldsymbol{\Xi}} + \varepsilon\boldsymbol{\varsigma}$
$\boldsymbol{\Xi}_0, \mathbf{D}_0$	Central ratio and curvature scale at the initial point
$\mathbf{R}$	Reduced-cost ratio $(\mathbf{c}-\mathbf{B}\mathbf{y})/(\mu\mathbf{w})$
$\mathbf{S}$	Third-order barrier scale $\mathbf{f}^{-3} - (\boldsymbol{\upsilon}^{\max} - \mathbf{f})^{-3}$
$\mathbf{Z}$	Squared barrier scale $\mathbf{z}^2$
$\mathbf{Y}$	Weight ratio $3\boldsymbol{\lambda}/\mathbf{w}$
$\mathbf{k}$	Step denominator $\mathbf{D}\boldsymbol{\Xi} + \mathbf{Y}(1 - \mathbf{D}\overline{\boldsymbol{\Xi}})$
$\mathbf{g}, \mathbf{g}_0$	Step right-hand side and its leading term
$\mathbf{q}$	Potential part of the step, $\mathbf{z}\mathbf{B}\mathbf{x}^{\mathrm{dir}}/(\mu\mathbf{w})$
$\mathbf{x}^{\mathrm{dir}}$	Potential direction of the solve
$\boldsymbol{\Upsilon}, \boldsymbol{\Lambda}$	Implemented and exact flow directions
$\mathbf{d}$	Movement scale $\lvert\boldsymbol{\Upsilon}\rvert + \lvert\boldsymbol{\varphi}\rvert$
$\boldsymbol{\nu}$	Central noise coefficient $(\mathbf{g}_0 - \mathbf{q})/(\varepsilon\theta)$
$\boldsymbol{\varphi}$	Source of the central correction
$\boldsymbol{\iota}$	Leverage correction $1 - s/\boldsymbol{\lambda} - \boldsymbol{\tau}/\widehat{\boldsymbol{\tau}}$
$\mathbf{T}$	Selection scaling, $\mathbf{T}_e = I_e/p_e$
$\boldsymbol{\xi}$	Selection noise $(\mathbf{T}-1)\boldsymbol{\nu}$
$\mathbf{G}$	Gradient of the tracking potential
$\mathbf{E}$	Scaled change $\mathbf{U}^T \textsc{Diag}(\boldsymbol{\sigma})\mathbf{U}$
$\mathbf{Q}, \mathbf{K}_0$	Root-update operators of the local step
$\mathbf{X}$	Per-edge variation scale
$\mathbf{I}$	Identity
$\mathbf{M}, \mathbf{Q}_\omega, \mathbf{Q}_\mu$	Observer matrices $\mathbf{B}^T \textsc{Diag}(\omega)\mathbf{B}$, $L_C(\omega) + D_{R_E}(\omega)$, $L_C(\mu) + D_{R_E}(\mu)$
$\mathbf{W}$	Price diagonal $\textsc{Diag}(W_e)$ of the exact output
$\mathbf{x}$	Generic vector; iterates $\mathbf{x}_i$ of the observer solves
$\varkappa$	Scalar step argument of Φ and $\mathbf{R}$
ϵ_0	Prefix threshold of the upper-stopped drift argument
Δ_C	Common-grid gauge mesh of a core occurrence
U_R, U_Z	Computable upper bounds of the two potentials
d_0	Endpoint list depth $n^{64/p}\log^{D_0} n$
L_C, D_{R_E}, A_{R_E}	Weighted Laplacian of the endpoint union and residual degree and adjacency

Symbol	Meaning
Q_0	Reference operator $L_C(\omega) + D_{R_E}(\omega)$
$\mathbf{x}_j^{\mathrm{rand}}$	Random vertex vectors of the observer
P_M	Proxy $(1 \pm \epsilon)M$ of the observer
$\mu_0, \mu_{\mathrm{end}}$	Initial and final barrier parameters
C_T	Phase constant of the attempt length
T	Attempt length $C_T\sqrt{r}/\beta_s$
W_0	Root volume $2q_0$ of the core construction
d_C, d_*	Degree in the core C and a known lower bound on $\min d_C$
r_u	Guard reference degree of a vertex
V_A	Reference volume $r(A)$
A_0, A_1	Sample-adaptive set and active near-expander of the pruning game
δ_0	Degree-estimate error parameter of the repair lemma
S_1	Residual set reachable from unsaturated sources