Computability of noise, mixing and correlation diagrams

An uncomputable disorder line in a zero-field Ising crystal

Seth Douglas Nidhal Mghirbi

Version 4.0.0, November 2026

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

Archive: doi:10.5281/zenodo.24147268

Abstract

Which temperature sets can one finite magnet realize through its correlations? On $[87/4, 45/2]$, a zero-field Ising model with exchanges in $\{-1, +1\}$ has a sampled spin correlation that changes, at a right-computably enumerable inverse temperature of halting degree, from positive strictly decreasing decay to an asymptotically damped cosine of period three in units of a fixed vector d. Its nonplanar degree-three graph has 32 vertices per geometric cell, about 4.7×10^{22} spins per signed repeat, and $|d| \approx 1.8 \times 10^{10}$ lattice spacings. One Gibbs state, exponential weak mixing and analytic pressure persist throughout the window. Every fixed-lag correlation is computable, but no computable lag or signal guarantee follows from a prescribed distance above the cut, and no sampled correlation in the window is negative before lag 10^{329}. More generally, for every relatively effectively open subset W of a positive rational window, a periodic ± 1 sign pattern on this fixed geometry, obtained effectively from an enumeration of W, makes W the oscillatory set of a sampled spin correlation; the converse assumes uniqueness and the all-lag sign promise. For finite rational magnets, negative point-field response somewhere has exactly the relatively effectively open diagrams under uniqueness. These results complement a noise diagram on the line, a mixing/order diagram on the square lattice, effective openness of mixing sets, their sequential realization, and exact realization of explicit semialgebraic diagrams. Realizing every relatively effectively open mixing set by one magnet with opposite order on its complement remains open.

Contents

1 Introduction **2**

2 Models and convergence **9**

3 The local rule **10**

4 Innovation estimates **15**

5 Organization of the simulation **21**

6 The dynamical alternatives **82**

7 The square-lattice interaction **87**

8 Effective openness of planar mixing sets **92**

9 Sequential realization of effectively open sets **94**

10 Exact explicit diagrams by line decorations **99**

11 A single magnet with a prescribed response-sign set **107**

12 Finite hidden Markov rows force exponential forgetting **115**

13 Prescribing sampled decay shape on one Ising geometry **119**

14 An explicit crystal and its halting-degree disorder line **124**

15 Discussion **133**

How this work was done **134**

A Numerical choices **135**

B Case tables for the numerical bounds **137**

C Dependencies of the lemmas **139**

D Software checks: scope and results **142**

E Reproducing the finite objects **143**

References **144**

1 Introduction

A finite interaction can have an uncomputable boundary in the shape of its correlations while its Gibbs state stays unique, exponentially mixing and analytic. We construct a zero-field Ising example with unit-magnitude exchanges. At one noncomputable inverse temperature, a specified spin correlation changes from positive strictly decreasing decay to an asymptotic damped cosine, with period three in units of one specified vector d. This is a disorder line in Stephenson's correlation sense [46]. Fisher and Widom studied the corresponding monotone/oscillatory distinction for linear systems [47].

The observable requires an infinite sequence of separations. Every fixed-lag correlation is computable, but lag and positive signal bounds cannot depend computably only on a promised distance above the cut. The graph is nonplanar, with a 32-vertex geometric cell and degree three; its signs repeat only after about 4.7×10^{22} spins, and one displayed lag spans about 1.8×10^{10} geometric lattice spacings. The conclusion concerns an exactly specified infinite-volume correlation.

This example leads to a realization question: which inverse-temperature sets can be oscillatory sets on one fixed Ising geometry? For sampled correlations every relatively effectively open set is realized, and the converse holds under uniqueness and an all-lag sign promise. A separate construction answers the negative point-field-response question among finite rational magnets unique on the window. The full mixing-phase realization problem remains open. These questions extend two earlier strands: temporal forgetting on the line and one Gibbs mixing/order diagram on the square lattice.

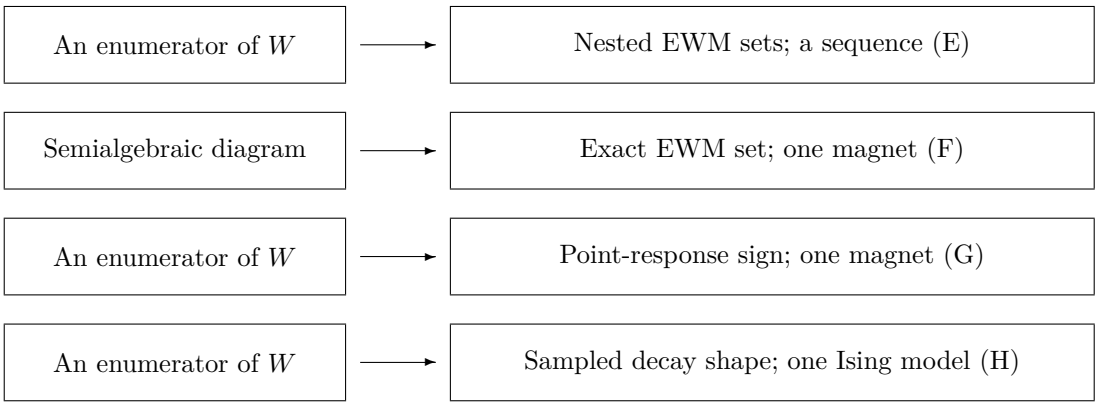

Figure 1: Realization in a positive rational inverse-temperature window. E exhausts W by actual EWM sets and orders outside W; F realizes an explicit diagram with order on its complement. G and H keep one Gibbs state, EWM and analytic pressure throughout, prescribing respectively negative point-field response somewhere and oscillation along d. H is an algorithmic existence construction. The explicit crystal I has its own binary-energy parent. Whether arbitrary relatively effectively open W is one magnet's EWM set with opposite translation-invariant order on its complement remains open.

Decay shape on a fixed Ising geometry. For a sampled spin write $\mathcal{C}_\beta(n) = \langle \sigma_o \sigma_{o+nd} \rangle_\beta$, $n \geq 0$, and $\mathcal{F}_\beta(t) = \sum_{n \geq 0} \mathcal{C}_\beta(n) t^n$; here d is a lattice vector, unrelated to the integer d of Theorem F. "Nearest" refers to the circle of convergence of this lag generating function. For a finite-period interaction, a local spin X, a site o and a vector d, the sampled correlation $\langle X_o X_{o+nd} \rangle_\beta$ satisfies the *all-lag sign promise* on $[a, b]$ with *oscillatory set* $\mathcal{O} \subseteq [a, b]$ if it is positive for every $n \geq 1$ at each $\beta \in [a, b] \setminus \mathcal{O}$, and at each $\beta \in \mathcal{O}$ its sign agrees with that of $\cos(2\pi n/3)$ for all large n. Under the promise, $\mathcal{O}$ is exactly the set of $\beta \in [a, b]$ at which some positive lag is negative.

The geometric graph is fixed; the finite periodic signs and the sampling data are chosen once from the input enumeration.

Theorem H (Ising decay-shape diagrams)**.** Given rational $0 < a < b$ and an enumerator of a relatively effectively open $W \subset [a, b]$, there is an effective construction of one periodic sign pattern $\epsilon_e \in \{-1, +1\}$ on the fixed 32-vertex, three-regular geometric graph of Section 14.6, a site o, and a period vector d. For the zero-field Hamiltonian $\mathcal{H} = -\sum_e \epsilon_e \sigma_u \sigma_v$, only $\beta \in [a, b]$ varies. Throughout this window there is exactly one DLR measure, uniform arbitrary-volume EWM, real-analytic pressure and real-analytic finite marginals. Off W, $\mathcal{C}_\beta(n) > 0$ decreases strictly from $n = 0$. On W, for pointwise constants $A_\beta > 0$ and $0 < \widehat{\varsigma}_\beta < \varsigma_\beta < 1$,

$$\mathcal{C}_\beta(n) = A_\beta \varsigma_\beta^n \cos(2\pi n/3) + O(\widehat{\varsigma}_\beta^n),$$

and $\mathcal{F}_\beta$ has exactly two nearest singularities, both simple poles. Conversely, for a finite-period rational interaction, a rational local spin X, a site o and a sampling vector d, uniqueness throughout the window and the all-lag sign promise with oscillatory set $\mathcal{O}$ imply that $\mathcal{O}$ is relatively effectively open, uniformly in these data and the window.

The period-three law is in units of d. The construction is an algorithmic existence proof; a general compiler from enumerator code is not supplied. A square-lattice variant uses exchanges in $\{-1, 0, +1\}$ on nearest and diagonal pairs. The proof is in Section 13.

The next theorem uses a separately specified parent and a pointwise executable bond printer.

Theorem I (An explicit crystal)**.** On that same nonplanar 32-vertex geometric graph, the periodic signs specified in Section 14 give one zero-field Ising model with every exchange in $\{-1, +1\}$. In the window $\mathcal{I}_{\mathrm{cr}} = [87/4, 45/2]$, there is exactly one DLR measure, uniform arbitrary-volume EWM, real-analytic pressure and real-analytic finite marginals. For

$$o = (552699460, 0, 0), \qquad d = (0, 17686381696, 0),$$

where the third coordinate indexes the geometric basis, the mean-zero spin correlation $\mathcal{C}_\beta(n)$ is positive and strictly decreasing from $n = 0$ for $\beta \le \beta_\star$, and has the two-simple-pole cosine form of H for $\beta > \beta_\star$. The interior cut $\beta_\star$ is right-computably enumerable and has exactly the Turing degree of the halting set. At every $\beta \le \beta_\star$, including equality, the convergence circle has one real simple pole and two non-pole singularities at phases $\pm 2\pi/3$.

The signed repeat contains 46 752 085 813 280 748 601 344 spins; every vertex has degree three. One lag is 17 686 381 696 axial lattice spacings and four parent lags, so period three is again in units of d. The monotone-side oscillatory remainder has no exponential gap below its positive leading rate.

Proposition 1.1 (Finite-data limits for I)**.** *Every fixed-lag correlation of I is computable, uniformly from a computable inverse temperature in $\mathcal{I}_{\mathrm{cr}}$. There is no total computable $N(k)$ bounding a negative sampled lag for all $\beta \ge \beta_\star + 2^{-k}$ in that window, and no total computable $\varepsilon(k) > 0$ guaranteeing any negative lag of magnitude at least $\varepsilon(k)$ there. In its binary parent, no negative value occurs at any parent inverse temperature in $[1/4, 1/2]$ through parent lag 10^{330}. The quoted scale 10^{2536} for a parent onset is an estimate only.*

Section 14.8 proves these statements and keeps parent lags, sampled lags and geometric spacings distinct.

From forgetting to planar order. Independent noise can make a cellular automaton forget its initial configuration. Gács showed that strictly positive transition probabilities on the line can also coexist with indefinite memory [10, 11]. For a deterministic rule perturbed by independent, uniform whole-symbol replacement, the replacement probability is a single scalar. Our first diagram asks whether the full automaton forgets at that rate; the associated space–time diagram asks whether its Gibbs specification mixes.

Fix a universal machine U with inputs $n \in \mathbb{N} = \{0, 1, \ldots\}$. In the following statements, uniform exponential ergodicity refers to the entire observable state at each site. Failure of convergence means failure of weak convergence from at least one initial law.

Theorem A. There exist one finite-alphabet, translation-invariant, finite-range cellular automaton on $\mathbb{Z}$ and a computable sequence of disjoint nonempty rational open intervals $I_n \subset (0, 1)$. At every uniform replacement rate $\varepsilon \in I_n$, the resulting probabilistic cellular automaton is uniformly exponentially ergodic if $U(n)$ halts. If $U(n)$ runs forever, its laws starting from some initial law fail to converge weakly.

Theorem B. There exist one translation-invariant, finite-range interaction Ψ on $\mathbb{Z}^2$ with a finite alphabet $\mathcal{B}$, whose anchored terms take values in $\{0, 1\}$, a fixed-point-free involution σ of $\mathcal{B}$ that leaves Ψ invariant, a one-site observable M with $M \circ \sigma = -M$, and a computable sequence of disjoint nonempty rational open intervals $J_n \subset (0, \infty)$ of inverse temperature. Let $\beta \in J_n$. If $U(n)$ halts, Ψ has exactly one DLR measure, and its specification has exponential weak spatial mixing in arbitrary volumes. If $U(n)$ runs forever, Ψ has two translation-invariant DLR measures, exchanged by σ, under which M has expectation at least $1/4$ and at most $-1/4$, respectively. No assertion is made at inverse temperatures outside the union of the J_n.

Theorem C. There exist one rational replacement rate $\varepsilon_* \in (0,1)$ and a computable family of finite-alphabet radius-one cellular automata f_n on $\mathbb{Z}$. At rate ε_*, the perturbation of f_n is uniformly exponentially ergodic if $U(n)$ halts. If $U(n)$ runs forever, its laws starting from some initial law fail to converge weakly.

Which mixing sets can a finite magnet realize? The square-lattice result leads to a second question: which open sets of inverse temperatures can be mixing sets of finite magnets? Here a *finite rational magnet* is a translation-invariant finite-range interaction on a finite alphabet, specified by finite rational energy tables, together with an invariant fixed-point-free alphabet involution σ and a rational one-site observable M, $|M| \le 1$, with $M \circ \sigma = -M$. Exponential weak spatial mixing (EWM) always means the arbitrary-volume estimate (3). For a rational window $0 < a < b$, write

$$\mathcal{M}_\Phi = \{\beta \in [a,b] : \Phi \text{ has EWM at } \beta\}.$$

A relatively effectively open subset of $[a,b]$ is its intersection with a computably enumerated union of rational open intervals. Magnetic order with gap m_* means two translation-invariant DLR measures exchanged by σ, with M-means at least m_* and at most $-m_*$.

The upper bound applies uniformly to every finite table.

Theorem D (Effective openness)**.** For every finite rational magnet Φ on $\mathbb{Z}^2$ and rational window $0 < a < b$, the set $\mathcal{M}_\Phi$ of inverse temperatures with arbitrary-volume EWM is relatively effectively open, uniformly from the interaction table, its range and a, b.

The converse holds for a sequence whose spatial range is allowed to grow.

Theorem E (Sequential realization)**.** Given rational $0 < a < b$ and an enumerator of a relatively effectively open $W \subset [a,b]$, one can compute a sequence $(\Phi_s)_{s\ge 0}$ of finite rational magnets on $\mathbb{Z}^2$ such that

$$\mathcal{M}_{\Phi_s} \subseteq \mathcal{M}_{\Phi_{s+1}} \subseteq W, \qquad \bigcup_s \mathcal{M}_{\Phi_s} = W, \qquad \overline{\mathcal{M}_{\Phi_s}} \subseteq \operatorname{int}_{[a,b]} \mathcal{M}_{\Phi_{s+1}}.$$

Every stage has magnetic order with gap $3/8$ at every $\beta \in [a,b] \setminus W$. The sequence has a common alphabet, involution, observable, integer energy scale J, anchored energy values $\{0, J\}$, and vertical range one; its finite horizontal range may grow. No classification of a stage is asserted on $W \setminus \mathcal{M}_{\Phi_s}$.

An explicit algebraic description permits exact realization by one magnet.

Theorem F (Explicit semialgebraic diagrams)**.** Given rational $0 < a < b$, an integer $d \ge 1$, and an explicitly specified rational semialgebraic open set $U \subset (0,\infty)$, one can compute one finite rational nearest-neighbor magnet Φ on $\mathbb{Z}^2$ whose arbitrary-volume EWM set in $[a,b]$ is exactly

$$W_U = \{\beta \in [a,b] : e^{-\beta/d} \in U\},$$

and which has magnetic order with gap $1/4$ at every other inverse temperature in the window, including every relative boundary point of W_U in $[a,b]$. Its single-site alphabet is the product of 65536 colors, at most three hidden line states for each of the two directions, and two outgoing edge indices from a finite list. Three line states can be necessary among ferromagnetic 65536-color Potts line decorations with identical horizontal and vertical matrices and exact individual-Perron weights for every periodic color word, even when the rational energy denominator is freely chosen.

A complete response-sign characterization. If a magnet is unique on the window, symmetry gives $\mu_\beta(M_x) = 0$. Its response at x to a dimensionless point field at the origin is

$$R_\beta(x) = \frac{d}{dh}\Big|_{h=0} \frac{\mu_\beta(M_x e^{hM_0})}{\mu_\beta(e^{hM_0})} = \langle M_0 M_x \rangle_\beta.$$

Write $\mathcal{R}_\Phi = \{\beta \in [a,b] : R_\beta(x) < 0 \text{ for some } x \neq 0\}$. For this observable one magnet suffices for every relatively effectively open set.

Theorem G (Response-sign diagrams)**.** Given rational $0 < a < b$ and an enumerator of a relatively effectively open $W \subset [a,b]$, one computes one finite rational magnet Φ on $\mathbb{Z}^2$, with $M = S \in \{-1,1\}$, having the following properties.

(i) At every $\beta \in [a,b]$ it has exactly one DLR measure μ_β, and its specification has arbitrary-volume EWM (3). Its pressure is $p(\beta) = \log g_0(e^{-\beta})$ for one positive-coefficient integer polynomial g_0, hence is real-analytic on the window.

(ii) Its point-field law $\mu_{\beta,h} = e^{hS_0}\mu_\beta/\mu_\beta(e^{hS_0})$ is the unique Gibbs state with that field, and $\mu_{\beta,h}(S_x) = \tanh(h) R_\beta(x)$ for every real h and every $x \neq 0$, where $R_\beta(x) = \langle S_0 S_x \rangle_\beta$. One has $\mathcal{R}_\Phi = W$, and $R_\beta(x) = 0$ off the column $\{(0,n) : n \in \mathbb{Z}\}$.

(iii) Off W, $R_\beta(x) \geq 0$ for every x, and $R_\beta((0,n)) > 0$ for every $|n| \geq N_0$, with N_0 independent of β. On W, $R_\beta((0,n)) < 0$ for every $|n| \geq N(\beta)$. The least such onset $N(\beta)$ is unbounded near every relative boundary point of W.

(iv) Conversely, for every finite rational magnet unique at every $\beta \in [a,b]$, its response-sign set $\mathcal{R}_\Phi$ is relatively effectively open, uniformly in the table, range, observable and rational window.

In the constructed family, some negative response and eventual negative column response coincide throughout the window. Proposition 12.1 explains why retaining its stationary iid hidden row under strictly positive finite decorations freshly sampled from the hidden neighborhood, ignoring their own earlier values, cannot supply the magnetic order required in the open mixing-phase problem.

A finite positive transition matrix contracts total variation after one step, so dimension one is the smallest possible lattice dimension for Theorems A and C. For a finite-alphabet one-dimensional finite-range interaction with finite energies, blocks of length at least the range give a nearest-neighbor transfer matrix with strictly positive entries. The Perron–Frobenius theorem gives convergence of block marginals from every pair of end conditions as the ends recede. The DLR equations then give uniqueness at every finite inverse temperature; see [28, Chapter 3] for the positive-matrix and Markov-specification argument. Theorem B attains the smallest possible dimension for this finite-energy dichotomy.

In Theorem B the alphabet is $\mathcal{B} = \mathcal{A}^T \times \mathcal{A}^T \times \{0,1\}$, where $\mathcal{A}$ is the alphabet of the automaton of Theorem A and T is its clock period, and each site carries $2T$ anchored terms. Merging the terms with a common support gives integer energies at most $2T-2$. The ordered states are the blocked trajectory of a period-T clock phase and its one-step time translate, placed in two independent copies in opposite orders. The symmetry exchanges the copies and flips a free bit. The order is translation invariant, with an exact symmetry and an odd one-site order parameter; it does not come from a ferromagnetic mechanism.

The automaton uses finite groups of neighboring sites to represent simulated sites. Each level repeats the same local program and executes a longer bounded-tape computation as the block size increases. To select the program, the rule measures distances between rare randomly refreshed walls. Within each noise window the distances usually identify its index. The measured index determines the lengths of spatial stripes, which are stable for the duration of the local simulations that read them. A comparison of three stripe samples at each scale locks the simulated input at the intended level. In the family of Theorem C the program

Result	Object and input	Observable and conclusion in the window
A, C	One rule with program windows (A); a family at one fixed rate (C)	Full-state exponential forgetting for halting; nonconvergence from a specified law for nonhalting.
B	One interaction; programs in fixed inverse-temperature windows	EWM on halting windows; opposite order on nonhalting windows; other temperatures unclassified.
D	Any finite rational magnet	Its arbitrary-volume EWM set is relatively effectively open.
E	A sequence; an enumerator of W	Nested actual EWM sets exhaust W; every stage orders outside W.
F	One magnet; an explicit semialgebraic diagram	Exact EWM set and opposite order on its complement.
G	One magnet; an enumerator of W	Some negative point-field response exactly on W; eventual negative sign along the column on W, eventual positive sign there off W; EWM everywhere.
H	One sign pattern on a fixed Ising geometry; an enumerator of W	Sampled oscillatory set exactly W; all-lag positive decrease off W; one DLR state and EWM everywhere.
I	One explicit nonplanar crystal; a fixed rational window	Monotone decay through a halting-degree cut; asymptotic cosine above it, period three in units of d.

Table 1: Known realization results, with each observable and object specified. Theorems D–H use a rational window $0 < a < b$; I uses $\mathcal{I}_{\mathrm{cr}}$.

is specified directly in the local rule. In both constructions a halted computation forces a region of simulated sites to issue the empty symbol. A long physical disagreement would then require a long connected chain of locally bad space–time regions, which has exponentially small probability. When the computation continues, a synchronized initial hierarchy retains a clock signal on every observation scale with high probability.

Related results. Kurdyumov's undecidability of ergodicity for one-dimensional probabilistic cellular automata allows zero transition probabilities [21]. Toom obtained undecidability results for binary cellular automata with zero transition probabilities [20]. For deterministic cellular automata on the line, Bušić, Mairesse and Marcovici established that ergodicity is equivalent to nilpotency [2], which is undecidable by Kari's theorem [13]. Theorem C concerns ergodicity at one fixed positive rate. Other undecidability results with positive noise concern stochastic stability in a vanishing-noise limit [16]. Marcovici and Taati prove that positive-rate PCA with stationary Bernoulli measures are exponentially ergodic in every dimension, and that in dimensions at least two the existence of such a measure is algorithmically undecidable [15]. The fault-tolerant mechanism adapts Gács [10, 11]: the 1986 paper's Sections 5, 7 and 9 supply computation, Purge, Shrink and Grow designs, and Sections 10–16 supply the tile, properness, nested-rectangle, healing and block-integrity, weak/strong computation, separation and end-relief organization. Its Lemma 4.1 supplies the sparsity argument. We specify the modifications and prove the required implications below.

Marsan, Sablik and Törmä constructed a one-dimensional automaton that is ergodic under small and large uniform noise and nonergodic at an intermediate rate [18]. Their work asks how the noise parameter can organize changes of dynamical behavior. Marsan and Sablik characterized the possible uniqueness sets for noisy continuous maps on a Cantor space, including effective realizations [17]. Theorem A concerns an uncomputable diagram of convergence for a cellular automaton. Its halting windows lie in the uniqueness set; no membership claim for that set is made on its nonhalting windows. Its nonhalting witness

separates the temporal phases of a one-site clock. The computable ordering of the windows accumulates at zero. Because the halting set and its complement are both infinite, the two regimes switch infinitely often along this ordering.

The quantum spectral-gap undecidability theorem of Cubitt, Pérez-García and Wolf [3] and the continuous one-parameter quantum Hamiltonian family of Bausch, Cubitt and Watson [1] give ground-state precedents for uncomputable phase diagrams. In the latter, the varied parameter enters the Hamiltonian, and the alternatives have different spectral gaps and ground-state order. Purcell, Li and Cubitt place a single ground-state phase transition at a halting probability in another quantum Hamiltonian family, again varying a coupling parameter [43]. Here the interaction tables are fixed and the varied equilibrium parameter is inverse temperature; Theorem G changes the sign of the point-field response along a column while retaining uniqueness, EWM and analytic pressure. For classical models, Gu, Weedbrook, Perales and Nielsen proved undecidability of macroscopic ground-state properties of infinite periodic Ising lattices, including correlation quantities [22]. Their argument supplies a classical precedent; H and I concern positive-temperature sampled decay shape in a unique mixing Gibbs state. Fixed finite-range models can have chaotic zero-temperature limits, as shown by Chazottes and Hochman [23], Chazottes and Shinoda [24], and Barbieri, Bissacot, Dalle Vedove and Thieullen [25]. Gayral, Sablik and Taati show that uniform finite-range models with rational-valued interactions have closed connected Π^0_2-computable ground-state sets and prove a converse realization theorem up to computable affine homeomorphism. They also prove that zero-temperature chaoticity of a two-dimensional finite-range interaction is undecidable and anticipated that most thermodynamic questions near zero temperature are undecidable in dimensions at least two [26]. The passage from cellular automata to space–time Gibbs measures originates with Domany and Kinzel [4] and was developed by Goldstein, Kuik, Lebowitz and Maes [12] and Lebowitz, Maes and Speer [14]. Miękisz and Taati use noise-resilient automata to obtain ordered low-temperature behavior for finite-range interactions [19].

Our fixed-rate reduction in higher dimensions left the one-dimensional case open [5]. We previously proved infinitely many transitions for one noisy automaton on $\mathbb{Z}^4$ [7] and a halting-indexed diagram on $\mathbb{Z}^3$ with a space–time Gibbs consequence [6]. Theorems A and C give the corresponding reductions on the line. The all-DLR uniqueness implication in Theorem B uses the one-dimensional space–time equivalence [8]. Our earlier two-dimensional positive-temperature diagram used hard constraints and controlled only translation-invariant Gibbs measures [9]; our all-Gibbs diagrams used higher-dimensional lattices [6]. Our spontaneous-magnetization preprint treated finite-energy two-dimensional input-dependent families at one fixed positive temperature [27]. Theorem B gives one finite-energy interaction on $\mathbb{Z}^2$, with the input encoded in temperature windows, uniqueness among all DLR measures in the halting windows and translation-invariant order with an exact symmetry in the nonhalting windows, closing the finite-energy question raised by [9].

Related planar constructions. The finite certificate behind Theorem D is the Dobrushin–Shlosman uniqueness condition; see Section 8. Summing auxiliary spins into a temperature-dependent coupling uses the decoration transformations generalized by Fisher [32]; Syôzi obtained the transition temperature of the Kagomé lattice [31]. Decorated planar Ising-type models (Ising–Heisenberg and spin–electron) already exhibit several re-entrant transitions [33, 34], and magnetic re-entrance in the Ising model on a decorated square lattice is analysed in [44]. The construction for Theorem F shares auxiliary spins along entire rows and columns. It prescribes every explicitly given rational semialgebraic EWM set, selects individual real algebraic conjugates, and controls the full-spin specification under arbitrary exteriors. The positive-coefficient conversion uses Pólya's 1928 positivity theorem and its

quantitative development by Powers and Reznick [35]. The nearest flexibility precedents for Theorem E are the perturbed continuous maps of Marsan and Sablik [17] and the two-transition automaton of Marsan, Sablik and Törmä [18]; their model classes and convergence properties differ from the EWM and magnetic-order assertions here.

Reading guide. Sections 2–7 develop the line automaton, its noise diagram and the planar Gibbs diagram, proving Theorems A–C. The hierarchy in Sections 3–6 can be treated as the computational machinery for those results. Sections 8, 9 and 10 then ask which full EWM sets can be prescribed, proving D, E and F. Only the ordered-state argument for E uses that hierarchy. Section 11 proves G by an independent finite Fourier construction, using Theorem 7.1 for the PCA-to-Gibbs step. Section 12 proves the supporting hidden Markov proposition. Sections 13 and 14 prove H and I. A reader interested in the decay-shape result can start with these statements, the binary-parent return formula in Section 14.1, the thermal reduction in Section 14.5, and the discussion. The hierarchy for A–C can be treated as a separate computational input; neither H nor I uses it.

Changes from version 3.0.0. Theorem G adds exact response-sign diagrams and their uniqueness converse; Proposition 12.1 gives exponential forgetting for positive line PCA with a finite hidden Markov stationary row. References to the lettered theorems now print their letters. Theorems H and I add sampled Ising decay-shape diagrams and an explicit crystal with a halting-degree disorder line, together with finite-data limits, a Boolean circuit identity certified on all inputs, and its exact thermal reduction. Theorems A–F and their proofs are retained. The title changes from "An uncomputable noise diagram on the line and a finite-energy Gibbs diagram on the square lattice" to reflect the new results. The higher-dimensional constructions and their formal appendix remain in version 1.4.0 [6].

2 Models and convergence

A finite-alphabet cellular automaton of radius r on $\mathbb{Z}$ applies one map $f : \mathcal{A}^{[-r,r]\cap\mathbb{Z}} \to \mathcal{A}$ simultaneously at all sites. A probabilistic cellular automaton (PCA) has a probability table $p(a \mid \xi)$ in place of f. Given the previous row, its outputs are conditionally independent at all sites. For uniform replacement with rate ε the table is

$$p_\varepsilon(a \mid \xi) = (1-\varepsilon)\mathbf{1}_{\{a=f(\xi)\}} + \frac{\varepsilon}{|\mathcal{A}|}. \tag{1}$$

One realization at each site and time draws an independent replacement flag of probability ε and an independent uniform symbol. The pair is called an innovation. A flag replaces the entire proposed symbol. We couple trajectories by using the same innovations for all initial configurations.

Write P_ε for the Markov kernel and η for a deterministic initial row. Uniform exponential ergodicity means that there are a stationary law μ and constants $C, c > 0$ such that, for all initial rows η, finite sets $F \subset \mathbb{Z}$ and integer times $t \geq 0$,

$$d_{\mathrm{TV}}\big((\eta P_\varepsilon^t)_F, \mu_F\big) \leq C|F|e^{-ct}. \tag{2}$$

Here d_{TV} is the supremum of the difference of probabilities over events. The bound extends by integration to every initial law. A finite-range kernel on a finite product alphabet is Feller. The Cesàro averages of the row laws from any start have weak limit points by compactness, and the Feller property makes every such limit stationary. The bound implies that the stationary law is unique. For the other alternative we exhibit an initial law and a one-site

event with probabilities that fail to converge. Cylinder events are clopen in the product topology, so the laws fail to converge weakly.

For a finite-range interaction Φ on a finite alphabet on $\mathbb{Z}^2$, let $H_\Lambda(\omega \mid \tau)$ be the sum of interaction terms meeting the finite set Λ in the configuration with interior ω and exterior τ. An interaction may list finitely many translation-invariant families of anchored terms; H_Λ sums every term meeting Λ. Its finite-volume specification γ_Λ^τ at inverse temperature β assigns weights proportional to $e^{-\beta H_\Lambda(\omega|\tau)}$. A probability measure is a DLR measure if its conditional interior law given the exterior is γ_Λ for every finite Λ. This definition allows arbitrary measures, including measures that break time-translation symmetry.

The normalized space–time interaction of a strictly positive table p contributes $-\log p(a \mid \xi)$ on each anchored child and its parents on the preceding row. In such a specification all factors meeting the volume enter, including the factor of an exterior child with an interior parent. In Section 7 we use exponential weak spatial mixing in the arbitrary-volume form

$$d_{\mathrm{TV}}\big((\gamma_\Lambda^\tau)_E, (\gamma_\Lambda^{\tau'})_E\big) \le C \sum_{v \in E} \exp\big(-c \operatorname{dist}(v, \Lambda^c)\big) \tag{3}$$

for finite Λ, $E \subset \Lambda$, and all exteriors τ, τ', with graph distance on $\mathbb{Z}^2$.

3 The local rule

A *rank* is the number of times a site has been encoded in a space–time block. The physical data track has rank zero. The rule acts at rank zero by the native recipe below. At higher ranks the recipe describes the computation performed on packets. It uses the same schedule and finite set of operations at every rank, with word width $\ell_j = 2^j \ell_0$. One tick and one site at rank j have physical sizes T^j and P^j.

3.1 States, bonds and structural updates

The Purge, grades, provisional growth, Shrink and attack/withdrawal design adapts Gács [10, Section 7(1–5), Section 9(M1–M12)]. The three evaluations and readout adapt his Section 5 and Section 9(M3–M5); the precise rule and guards used here are specified below.

Choose a sufficiently large integer w, and set

$$q = w^2, \quad A = w^3, \quad B = 128q, \quad v = w^4, \quad P_0 = 2^w, \quad P = 5P_0, \quad d = 4w + 8.$$

The precise ordering of numerical choices is in Appendix A. The native state is either empty, D, or a live tuple with clock $\tau \in [0, T)$, extended integer address $\pi \in [-11P/10, 21P/10)$, retained letter X, bits Cn and Temp, and the scratch registers below. A letter is either a distinct message-empty symbol or a finite tag with one word of width ℓ_j. Retained letters on live tuples are live. A live tuple can have arbitrary legal contents in the scratch registers.

The bond between an ordered left and right neighbor has grade zero if either is empty, if their clocks disagree, or if their residues modulo P do not increase by one modulo P. Under the clock and residue conditions the bond has grade two when the integer address increases by one, grade one for a nonexact wrap from $\pi > P/2$ to $\pi < P/2$, and grade zero in other cases. The ordinary addresses belong to $[0, P)$. A true endpoint is address zero on the left or $P - 1$ on the right in the site's ordinary home. During the extended phase the true positions of a source are its own addresses zero and $P - 1$. The schedule depends on the old clock. Each tick uses the old row and updates simultaneously. Untargeted fields are retained. Surviving incumbents increment the clock modulo T and normalize their addresses modulo P in every ordinary update. Death clears the whole state.

Every recruitment is made into an old vacancy. A donor's clock must request a recruitment instruction, and its proposed address is its own plus the displacement to the vacancy. The proposed address must be legal. The neighbor opposite the donor must be empty or have the donor's clock. When both proposals qualify the left one is taken. A birth inherits the donor's next age and the proposed address, with ordinary normalization when required. Its retained letter is the corresponding live letter of the neutral packet encoding D. Its scratch is default, except that provisional births set $\mathsf{Temp} = 1$. Old incumbents do not accept proposals.

Instruction	Structural effect
Purge initialization	Set $\mathsf{Cn} = 1$.
Purge propagation	Set $\mathsf{Cn} = 0$ at a left bond below grade two or if the old left $\mathsf{Cn} = 0$. Retain it otherwise.
Purge kill	Kill when the old $\mathsf{Cn} = 0$ and the right bond has grade below two.
Provisional birth	Try both directions. During Grow use the extended legal range. In ordinary work demand an in-home address in $[0, P)$ on every proposal.
Provisional removal	Kill a marked incumbent with a grade-zero face unless that face is its own true outward face.
Grow attack	Try both directions, with proposed address outside the donor's ordinary home and in the legal extended range.
Grow withdrawal	Kill an exposed extended left tip ($\pi < 0$) at a zero left bond, or an exposed extended right tip ($\pi \geq P$) at a zero right bond.
Shrink(i)	Kill at a grade-zero face of depth greater than $3iA$ from its corresponding true endpoint.

At the end of Grow perform the final old-address withdrawal, then normalize every surviving address on entry to ordinary work. The next Purge initializes Cn before a propagation or kill at the new address. A single exact source during Grow has a fixed value of physical child coordinate minus extended address, along with its clock. Normalization preserves residues and can merge labels that had distinct source identities.

3.2 Clock

A handle has one Purge (initialization, q propagation ticks and q killing ticks) and one elementary instruction. Set $H_0 = 2q + 2$. Each Main cycle has $2A + 4$ handles. Its instructions, in order, are a reset of Temp to zero, B simultaneous bilateral provisional attempts, B provisional removals, $2A + 2 - 2B$ idle instructions and the Main instruction. Its duration is $d_0 = (2A + 4)H_0$. Write $d_1 = qd_0$.

Begin with $\lfloor 7P/(5v) \rfloor$ Grow cycles, each of $2v$ Main attacks followed by v Main withdrawals. Next perform $\lfloor 3^w/q \rfloor$ calls of Shrink(0), $2P$ calls of Shrink(1), and $2P$ calls of Shrink(2). Use one call in each stretch of q Main cycles, leaving the other $q - 1$ Main instructions idle. The end of Grow is s_0, the end of the first Shrink stage is s_1, and r_{last} is the start of the final Grow withdrawal. The first Shrink stage accounts for more than $0.85T$ of the full schedule at large w.

Computation follows these three stages. Choose a service date s_c in early computation with more than w^6 fine ticks before the first input transaction. Include three repetitions of the program below, separated by P idle Main cycles, and the verification and final ballot. The time for computation is a fixed constant times P plus a fixed polynomial in w in Main cycles. Within this allowance insert two no-birth retreat-through-padding intervals in the verifier

phase separated by more than $3w^6$ fine ticks, each containing more than $100A + 4(M + 1)$ fine ticks of verification, with $M = 128d$. The second reaches the ballot, which is the last Main instruction of a normal cycle. The next period begins with a fresh Purge. The native handle cadence therefore continues across the vote. Verification runs on separate registers *on fine ticks*, concurrently with Purge and the other scheduled operations. The no-birth portions of the normal cycles have more than enough fine ticks for the specified intervals. Use a fixed schedule satisfying these bounds, with the finite program below padded by idle instructions to the chosen time budget. This fixes T before any helper constants are selected. Since Shrink(0) has duration of order $3^w d_0$ and the other stages use $O(Pq + \mathrm{poly}(w))$ Main cycles, the required fraction is achieved.

3.3 Representation and evaluation

In addition to the structural fields, allocate three input and three output letters, two directional mailbox letters, a carrier Y, two validation lanes each with a letter and a bit, a local invalid bit, a word-tape symbol, three external vector latches and a fixed tuple of parameter latch words. For the verifier allocate an active bit, two lane bits and two counters in $[0, M]$. Default scratch is blank, zero, headless, or message-empty in message fields. A diagnostic read of D supplies canonical defaults and empty carriers. The word-tape symbol is the product of finitely many tape tracks, a finite control/head mark, and finitely many words. In particular the old head can carry two operands and one result. The head reads only its own old cell, writes the old cell's work tracks and requests a stationary move, a left move or a right move. For each cell resolve incoming requests in order stationary, from left, then from right. Impose work-interval barriers and an inert reject state on malformed instructions. These rules define an update even with multiple forged heads.

Serialize a rank-$(j+1)$ tuple in one rank-j packet of P_0 letters. Use a live/dead discriminant and the canonical binary enumerations of all bounded fields in an $O(w)$-length area centered on the packet. Enumerate the registers in the declared order. Use separate scalar-bit tags with zero words and typed slots for the low and high word halves of each word field. Elsewhere require fixed live padding, including at both endpoints. Check tag and padding conventions at the packet level, reject out-of-range scalars and parse every legal tuple literally. In particular any legal mutable tag and payload, including headless registers with noncanonical content in the old state, is carried literally inside the typed word slots. A malformed packet parses as D. Encode D using a playable all-live neutral packet with a dead discriminant and canonical body; the writer for an actual dead result instead emits all message-empty. For a live s, its aligned encoding $E_j(s)$ consists of five repeats of its serialized packet $\delta_j(s)$: the site at block position $a = cP_0 + u$, with $0 \le c < 5$ and $0 \le u < P_0$, has age zero, ordinary address a, retained letter $\delta_j(s)[u]$, $\mathsf{Temp} = 0$, $\mathsf{Cn} = \mathbf{1}_{\{a \ge q\}}$ and all scratch at default. These two structural bits are the values left at the ballot. The last complete Purge before it has a grade-one left wrap and grade-two interior bonds, so it leaves $\mathsf{Cn} = 1$ exactly at $a \ge q$; the last Main cycle clears Temp in the full home; and the ballot retains both bits. Define $E_j(\mathsf{D}) = \mathsf{D}^P$, the all-empty block; it is distinct from the neutral packet that encodes D above. Every field of each endpoint symbol of $E_j(s)$, including the two structural bits, is independent of the particular live s.

Number the fifths of a home $0, 1, 2, 3, 4$. In each of three repetitions use the following ordered program. Reset the private output bank at its start.

1. For each supplier $j' \in \{-1, 0, 1\}$ and fifth $s \in \{1, 2, 3\}$, send its retained letter to bank s in fifth $2 + j'$ of the owner. For a displacement kP_0, inject a fresh message on that lane at each designated source and message-empty everywhere else. Then shift $|k|P_0$ times by unit nearest-neighbor shifts; each step copies the upstream mail only across a positive bond. Otherwise put message-empty. At receive, overwrite every addressed

bank including message-empty. Zero displacement is local assignment, and births have empty mail.

2. Vote at each offset across the three homologous received letters. Any repeated letter wins. In a three-way tie choose the first live letter if at least two are live, and message-empty otherwise. In each supplied fifth independently check local tag, selector and padding constraints, and spread the OR of invalid bits for P_0 steps within that fifth. Shift the two outer packets together with their invalid bits into the middle fifth. Overwrite the whole work tape, including heads, with the initialized interpreter and its three supplied input tracks.
3. At $u = \lfloor P/2 \rfloor, u+1, u+2$ load the actual read-only external words, and at u load the actual parameter. Gather these by two addressed nearest-neighbor transfers. The loads are unconditional on the old mutable tape and precede evaluation. The interpreter parses the supplied packets independently. It evaluates exactly one represented native update on the three parsed inputs using the actual parameter and vector words for any represented load operation. A halted true parameter returns D before it evaluates any mutable instruction. Thus the guard applies equally to valid states with forged work registers.
4. Form the live result packet or put message-empty at every position for a dead result. Use a fine-site flag broadcast if needed before writing the full packet, and send the central fifth to each of the five output fifths with fresh injection, shifts and overwrite as in step 1. Maintain the completed bank while the other repetitions run.

The interpreter may use copy, bitwise Boolean operations, modular addition, comparison, one-bit shifts, endpoint and zero tests, and population count on a word. Each operation on two half words uses a constant number of lower-width operations. For addition pass the low-half carry, for comparison compare the high halves first, and for shifts use the seam bit. The sum of the two half-population counts fits in one lower word for $\ell_j \geq 4$. Address and schedule arithmetic concerns $O(w)$-bit integers even when the simulated word is much larger. A fixed universal one-tape interpreter uses an initialization preamble and a quoted finite template p_{prog} acting on the data (p_{prog}, w). A represented initialization reconstructs the same preamble and supplied template from these data. When the represented state is a running head, evaluation performs its single elementary local instruction. This is a fixed finite tape algorithm. Compact-field scans and bounded arithmetic need a fixed polynomial budget $w^{O(1)} < P_0$ at sufficiently large w. Transfers, invalid broadcasts and fresh tape initialization use $O(P_0)$ Main cycles. A budget of $c_{\mathrm{prog}} P_0 + w^{k_{\mathrm{prog}}}$ per repetition for fixed sufficiently large constants, plus the stated clock paddings, suffices. The table is specified by these finite algorithms and fixed timing constraints and can be enumerated by evaluating the resulting local map on every legal neighborhood.

3.4 Verifier and ballot

After readout reset the verification scratch and set the active bit to the majority of the three completed liveness bits at the same site. During each epoch of $M + 1$ fine ticks, inject each current active bit in its two directed lanes, clearing the counters, then shift the lane bits at unit speed for M ticks. Each shift uses a strict ordinary in-home bond with consecutive addresses and the same clock. A missing input is zero. Count positive arrivals in the two directions separately (saturating at M). A counter exceeding $4d$ at the end of the epoch sets the local bit active. It remains active until death, reset or the ballot. Birth initializes all verifier fields to zero. At the final ballot an inactive site dies. An active site takes the payload vote by the same letter rule used at input, using a fixed live filler if that vote gives message-empty. The surviving site sets its age to zero, normalizes its address and clears every scratch register. The positive decision and the distributed payload are controlled separately

in the induction.

Lemma 3.1 (Clean full-row identity)**.** *On an aligned all-occupied row of encoded inputs with the stipulated read-only operands at the addressed loads throughout the three repetitions, one work period outputs exactly the aligned blocks E_j of the represented results, including $E_j(\mathsf{D})$. Every row at a date $0 \le u < T$ of the period is fully occupied, with the ordinary addresses of the inputs and the common clock u; only the closing ballot at date T can make a result block empty.*

Proof. The only weak bonds are aligned wraps a full home apart. A zero in the Purge propagation cannot reach the right end of a home in q ticks, so all Purges retain the complete row. Shrink, withdrawal and provisional removal find no zero face, and proposals cannot replace incumbents. Hence every row before the ballot is fully occupied. The input, work and readout transactions consequently compute and send the exact represented result in each repetition. In a home with a live result all the bits are initially active, and in a home with a dead result none can activate because the verifier's bonds are in-home. At the final vote the live result is encoded with its entire scratch reset, while the dead result becomes the all-empty block $E_j(\mathsf{D}) = \mathsf{D}^P$, since every site of its home is inactive and dies at the ballot. □

3.5 Read-only parameters

For the fixed-rate family the physical literal is the exact initial configuration of $U(n)$, hardwired into the native rule. The universal computation uses a one-ended tape with fixed bitplanes for symbols, state and head. At each promotion extend the word by a blank half and execute one true step. One step is implemented by a bounded list of masks, Boolean operations and one-bit shifts. The enlarged width accommodates the one newly visited site per step. If the literal for a transition is halted, its result is D on every triple. Every substep reads only the old site and its two immediate neighbors: structural tests use those three old states, head requests and mail move one site, lane copies move one site, and the three addressed LOADs are gathered by two nearest-neighbor transfers. Local word operations read only local registers. The family rule is therefore radius one. The geometry and clock do not depend on the initial literal or word width. For the single rule, the fixed physical literal includes SEARCH, unary counter zero and the strict majority threshold for the number of external channels, with base width at least $m_h + 4$. SEARCH with three fresh words z_0, z_1, z_2 compares

$$\operatorname{popcount}\big((z_0 \operatorname{xor} z_1) \operatorname{and} (z_1 \operatorname{xor} z_2)\big)$$

to that threshold. If a strict majority of channels make both flips, enter LOCKED with the true initial tape for the current counter index. Otherwise add one unary mark and use zero-extended z_0 as the next external word. The external word after locking is zero and subsequent promotions advance the locked tape as in the family. The native step for any represented load takes these designated read-only operands. Mutated latch values remain ordinary mutable data.

3.6 Physical embedding of the helper rule

Use the auxiliary rule of Sections 4.2–4.3. Take a fixed stride D_{lane} large enough to assign disjoint offsets in $[0, D_{\mathrm{lane}})$ to the data site and all private primitive roles, including every probe and individual replica. The data offset is zero. For every physical center x, the declared register of replica offset a is held in the field for a in the physical symbol at $x + a$. The symbol at each physical site has a field for each such offset, the data field, and a raw bit of deterministic proposal zero. Its deterministic local rule computes all its fields by translation

of the respective logical proposals. Recursions at logical center x use logical neighbors $x - D_{\text{lane}}, x, x + D_{\text{lane}}$, with the appropriate finite offsets. The data update reads the exported vector at its own logical center. Auxiliary proposals never read data. Each physical update takes a single uniform whole-symbol replacement after computing the proposed tuple. In one residue lane the sites serving different primitive roles are distinct. Shifted lanes overlap in physical space, so no joint product across lanes is asserted. The physical alphabet is a finite product and the physical radius is finite and fixed. No geometric coordinate, noise rate, input index or externally supplied clock enters this translation-invariant update.

4 Innovation estimates

4.1 Sparse sets and independent flags

Fix integers P, T, w as in the construction. Write $V^j = [0, P^j) \times [0, T^j)$, intersected with the site–time lattice, and use half-open rectangles throughout. Whenever an error set is restricted to a rectangle, it is taken to be empty outside that rectangle.

Definition 4.1 (Sparsity)**.** A set is 0-sparse if it is empty. It is j-sparse if every translate Q of wV^j admits a translate J of $3wV^{j-1}$ such that $(E \cap Q) \setminus J$ is $(j-1)$-sparse.

Sparsity is inherited by subsets, and $(j-1)$-sparsity implies j-sparsity. The latter follows by applying the restriction property before choosing any J. A full V^j is never j-sparse. For $j > 0$ one can choose a wV^j window containing it, and a deletion of $3wV^{j-1}$ cannot intersect two widely separated full V^{j-1} rectangles inside it. This reduces to the nonempty 0-cell. We use only the sparsity definition and the disjoint-pair argument from [10, Section 4, Lemma 4.1]. That lemma applies under bounds on joint errors. Here the iid case will suffice as the base for the correlated estimates.

Lemma 4.2 (Disjoint sparse-set witnesses)**.** *If the restriction to a full box of side vector $2wV^{j+1}$ fails $(j+1)$-sparsity, two disjoint full child boxes of side vector $2wV^j$ inside it fail j-sparsity. The child positions come from a deterministic list of at most $(2P+3)(2T+3)$ boxes. Consequently, for independent flags of rate u sufficiently small, the probability of j-nonsparsity on a prescribed fixed multiple of V^{j+1} is at most $C_0 u^{c_0 2^j}$, with $C_0, c_0 > 0$ independent of j, of the translate and of u.*

Proof. Place the full child rectangles on a grid of spacing wV^j, using positions flush to the boundaries. The intersection of any moving wV^j test with the parent fits in one of these children. If the bad children pairwise intersect, their lower corners span at most one grid step in each coordinate. Their union then fits a $3wV^j$ rectangle. Delete that rectangle. Each remaining moving wV^j test is contained in a good child or has no errors in the parent. This makes the remaining error set j-sparse. The same reasoning applies to restriction to any moving parent window, contradicting $(j+1)$-nonsparsity. Thus there is a disjoint pair.

Let b_j be the supremum, over locations of a full rank-j box, of its nonsparsity probability. At rank zero $b_0 \le 4w^2 u$. Disjoint children carry independent flags and there are at most $D = 8(P+2)^2(T+2)^2$ candidate pairs. Thus $b_{j+1} \le D b_j^2$, and

$$b_j \le D^{-1}(4Dw^2u)^{2^j} \le D^{-1}u^{2^{j-1}} \quad \text{if } u \le (4Dw^2)^{-2}.$$

A fixed multiple of V^{j+1} can be tested with a bounded cover by full rank-j boxes, arranged so each moving wV^j test intersected with the prescribed rectangle fits in a covering box. If all these boxes are good then every such test is good. The number in the cover depends only on the multipliers and P, T, w, which proves the bound. □

4.2 Autonomous estimation of the window

The windows are $I_n = (2^{-A_n}, 2^{1-A_n})$, where $A_n = 8^{n+n_0}$ and the fixed integer n_0 will be chosen last. Set $p = \varepsilon/2$. In one logical lane the raw probes are previous-row values of physical bits with deterministic proposal zero. They are independent Bernoulli p in all updated rows. The data and helper roles in a given lane use disjoint offsets, as specified in Section 3.6.

The receiver uses 13 private probes for a replicated wall bit, five for a replicated mode bit, and three groups of 59 replicas for wall, mode and work. Propose a wall refresh when the first 12 wall probes are one, adopting the thirteenth bit. Otherwise hold the wall bit. Propose a mode command when the first three mode probes are one, choosing CLEAN if the last two are one and RUN otherwise. Otherwise hold its mode. Independent commands of the two kinds have respective probabilities p^5 and $p^3(1-p^2)$ at a given date. Decode by plurality with a fixed first-replica tie rule. A decoded register error requires at least 30 flags in its group, with probability at most $2^{59}(2p)^{30}$. This necessary event refers only to flags and holds whatever the previous contents. Use the proposal of the decoded recurrence for every replica. The raw proposal bit and every replica group occupy private offsets.

Walls are barriers for all work and control storage. A wall's own work state is canonical. Its mode drives the open segment to its right by a speed-one wave. CLEAN blanks every workspace, display, head and launch register reached by the wave. After a continuous CLEAN, a CLEAN-to-RUN command launches one head at the first open site, one step behind the RUN wave. The head moves at most one site in two steps. The mode-edge test uses a one-step delay overwritten from the current mode. When a boundary is absent the same local rules apply, with conflicting head requests blanked and with CLEAN and barriers having priority. RUN while already RUN does not relaunch.

In a clean segment of length L, the head counts its interior sites by shuttling between an advancing marker and a binary counter near the left wall. It accepts n exactly when

$$A_n/3 \le \log_2 L \le 2A_n.$$

The bins are disjoint. Integer comparisons suffice, for example $L^3 \ge 2^{A_n}$ for the lower bound. Generate candidates until the next lower bound exceeds L. The head writes the result on a boundary track and can copy its binary description to the display. The tape has finitely many tracks, two local end tests, and a default abort on an invalid bin or work overflow. The same head will also compute the stripes in Section 4.3. All prepublication algorithms used here cost $O(L^2 \operatorname{poly}(\log L))$ rounds, for any reference input. The compact counters and finite template fit in the segment for all sufficiently large valid L. Shuttling a fixed number of times per bit of a counter costs $O(L \operatorname{poly}(\log L))$ per visited site. The hardwired offset n_0 enters only the bin arithmetic. At accepted lengths its contribution to the job-time constant is at most polynomial in n_0, hence at most $\operatorname{poly}(\log\log L)$; the final increase of n_0, whose windows shrink doubly exponentially, absorbs this factor uniformly and does not create a circular constant choice. The integer comparisons may stop generation as soon as a lower bound is too large for the observed L, so the cost at the accepted indices grows polynomially in $\log L$ with a fixed template. A fixed allowance covers the launch delay. The estimates below use the room between $L = O(p^{-3/2})$ and the time budget of order p^{-4}.

The finite-prefix assertion below is used only at observation dates $t \ge B_0 + H + 4$, with B_0, H as in its proof; burn-in or a two-sided stationary helper history supplies those dates.

Lemma 4.3 (A correct reader)**.** *For each observation in one lane at a rate in I_n, the receiver has a sufficient correct-job test using innovations in a deterministic rectangle. It passes with probability at least $1 - C\sqrt{p}$ for all sufficiently small p, uniformly over incoming states. On passing, the observing site lies at distance at least $\lfloor p^{-1/2} \rfloor$ from the current walls of a cleanly launched segment of length at most $2\lceil p^{-3/2} \rceil + O(1)$. The correct index and every stipulated*

publication fit before observation. For private copies of the receiver the tests are independent within one lane. The test for a reader is independent of the reference track values in the sense that the same passing innovations work for every static reference pattern.

Proof. Use $W = \lceil p^{-3/2} \rceil$, $H = \lceil p^{-6} \rceil$, $B_0 = \lceil p^{-14} \rceil$, $S = \lceil p^{-4} \rceil$ and $\ell_0' = \lfloor p^{-1/2} \rfloor$. Inspect the site interval of radius $W + 3$ in a prefix of B_0 and suffix of H rounds, with fixed extra boundary dates for the queried raw row. Require no excessive wall flags in this rectangle and no excessive mode or work flags in the suffix. The probability of a necessary flag event in these sets is at most

$$O\big((W + 1)(B_0 + H + 4)p^{30}\big) = O(p^{29/2}).$$

Work first with the ideal bits determined by commands and subtract this fault bound by a union bound. Every inspected site refreshes its wall within the prefix except with probability $(2W + 7)e^{-B_0 p^{12}}$. Given refreshes, the last proposals are independent Bernoulli p bits from the thirteenth probes. The probability of a wall refresh in the suffix is $O(WHp^{12}) = O(p^{9/2})$. The probability of a wall near the root within ℓ_0' or a missing wall within distance W on either side is $O(p\ell_0') + 2e^{-cpW}$. The resulting length satisfies the prescribed bin inequalities for all $p \in (2^{-A_n - 1}, 2^{-A_n})$ once n_0 is large. In particular its binary logarithm lies between $A_n/2 - O(1)$ and $3A_n/2 + O(1)$.

The selected left wall depends on wall probes alone. The last CLEAN in the suffix exists and precedes the final S rounds except with probability $e^{-cHp^5} + O(Sp^5)$. Conditional on its date, the later rounds have no CLEAN and have independent RUN commands of probability $p^3(1-p^2)/(1-p^5)$. The first such RUN is more than $2W + 10$ rounds after CLEAN and fewer than $S/2$ rounds after CLEAN except with probability $O(Wp^3) + e^{-cSp^3}$. Hence all stale work is blanked, and the uninterrupted run before observation has at least $S/2$ rounds for the stipulated publication. All inspected work follows the intended computation outside the charged flag event. The sum of the failure probabilities is $O(\sqrt{p})$. The construction selects no wall using work marks, and the sufficient event ignores the values supplied to the head. □

Lemma 4.4 (Forgetting of full receiver storage)**.** *At each sufficiently small fixed $p > 0$, the receiver has exponentially decaying one-slot all-input disagreement under common innovations. This is uniform for every common exogenous reference trajectory. It includes the undecoded physical storage.*

Proof. Inspect a deterministic interval of radius $2W + O(1)$ in the lane for the same B_0-round prefix as in Lemma 4.3, then eight consecutive H-round slabs. Allow a fixed number of boundary rounds and discard the first update for commands that query the incoming raw bit. The probability of any required excessive flags is $O(p^{29/2})$. In ideal bits the prefix produces common wall and mode values for all inputs except with probability $O(We^{-cB_0 p^{12}})$. For each spatial site the probability of a wall at the prefix end and no refresh in the eight slabs is at least

$$p\big[1 - (1 - p^{12})^{B_0}\big](1 - p^{12})^{8H + O(1)} \geq cp.$$

The events are independent at distinct sites. Thus except with probability $O(e^{-cpW})$ there is an enduring wall in each of the intervals $[-2W, -W]$ and $[W, 2W]$.

A helpful slab has no wall refresh inside the inspected interval, and at every potential left wall there is a CLEAN followed by $5W + 10$ rounds without RUN. The wall failure probability is $O(p^{9/2})$. Split the slab into order p^{-1} disjoint subintervals of length of order p^{-5}. In each, the first half has a CLEAN with probability bounded below, and in the following $5W + 10$ rounds RUN arrives with probability $O(Wp^3)$. These rounds fit in the same subinterval. A union over possible walls gives mode failure $O(We^{-c/p})$. The helpfulness tests on the eight slabs use disjoint primitive innovation rows, regardless of which walls endure, so their joint failure probability is $O(p^{36})$.

In a helpful slab the whole work in every segment between the enduring walls is blanked and becomes common. As long as the two walls persist, even changes of interior boundaries cannot bring in an old discrepancy. A final proposal and common physical replacements give agreement in the raw replica slots as well as in the decoded registers. The argument holds for each common reference stream and uses only noise to certify the conclusion. At a depth $d_{\mathrm{rec}} = B_0 + 8H + O(1)$ the all-input nonconstancy coefficient for a fixed slot is $q_{\mathrm{rec}} = O(p^{29/2}) + O(p^{36})$ plus superpolynomially small terms. Taking a finite union covers the finitely many slots queried by a physical symbol.

If a composition of k such slabs is nonconstant at a site, trace a chain of nonconstant output maps, one per slab. Each map has at most $2r_0 d_{\mathrm{rec}} + 1$ possible parent sites for the chain, where r_0 is the fixed physical range. For a fixed chain the failure tests depend on disjoint innovation slabs. The union bound is $((2r_0 d_{\mathrm{rec}} + 1)q_{\mathrm{rec}})^k$, which decays exponentially for small p because $d_{\mathrm{rec}} = O(p^{-14})$. A finite remainder at the initial end of the time interval does not change the conclusion. □

4.3 Spatial stripes

Choose $D_0 > 1$ after fixing the data geometry. A radius $H_n = D_0^{n+2}$ will dominate all prescribed fixed enlargements of the boxes through rank $n+1$, including the lookup coordinates and collars. At a rate in I_n, $H_n = p^{-o(1)}$, uniformly with arbitrarily small fixed positive exponent after increasing n_0.

Use an odd number m_h of independent reference channels Y_i, each with separate probe and replica offsets in the lane. At each site Y_i refreshes when its first b_h private probes are one and its next two probes differ. Adopt the first bit of that differing pair, otherwise hold the previous plurality. Use $2v_h - 1$ replicas. An uncommanded decoded change requires at least v_h flags. True refreshes have rate $2p^{b_h+1}(1-p)$; their orientation is fair. A refresh produces identical proposals in all copies of the trajectory, and common replacements preserve this equality after refresh. Conditional on refresh in a prescribed interval, the value at the end of the interval is fair: after the first refresh simultaneously complement every fair assigned bit and every subsequent fair replacement bit. At each site and in each channel these operations use private primitive pools.

For reference channel i, a stripe anchor is the right endpoint of the pattern consisting of one followed by $k_n = \lfloor A_n/8 \rfloor$ zeros. At a site x, use the nearest such endpoint o_i at or to the left of x within distance $R_n = 2^{\lfloor A_n/4 \rfloor}$. Publish zero if there is none. Otherwise publish $\lfloor (x - o_i)/P^n \rfloor \bmod 2$. The head records the last eligible endpoint during a left-to-right scan and checks its distance before publishing. A head reading a stable pattern uses at most $R_n + k_n + O(1)$ consecutive input sites for a given publication. It computes n, k_n, R_n, P^n using $O(\log L)$-bit counters for an accepted segment of length L. Keeping the last endpoint, detecting runs and writing the outputs adds $O(L^2 \operatorname{poly}(\log L))$ rounds for any fixed number of channels. As in Lemma 4.3, insufficient space causes an abort. On the sufficient event the wall is farther from the observation site than $R_n + k_n + 2$, since $R_n + k_n + 2 < \lfloor p^{-1/2} \rfloor$ for large n_0, so the output depends only on the intended reference window. Use an odd number s_h of private receivers, all reading the same Y_i, and export the coordinatewise majority of their decoded displayed bits (with a fixed default for malformed storage).

Lemma 4.5 (Locally static stripes)**.** *Prescribe $M_s > 0$ first, choose b_h for the short event, and only then prescribe $M_l > 0$. With these choices there are short and long bad events at every center. For suitable choices of the constants in the order described below, their probabilities are at most p^{M_s} and p^{M_l}, and their innovation supports lie within radii p^{-16} and $p^{-(b_h+5)}$, respectively. Outside both events, throughout the radius-H_n box each exported channel is static in time and a strict majority of the channels have values $\lfloor (x - o_i)/P^n \rfloor \bmod 2$,*

with possibly different phases o_i. The conclusion holds for all incoming helper states before the innovation prefix. At each sufficiently small fixed p, under common innovations the probability that the value of a physical helper slot s steps after an arbitrary start is not the same for all incoming helper states is at most $C(p)e^{-c(p)s}$. The complete autonomous helper process has a two-sided stationary realization determined by its past innovations.

Proof. At every read in the box use the sufficient reader tests of Lemma 4.3. They are independent across private copies at that read and work for every static reference pattern. The probability that the majority test fails somewhere is bounded by

$$O\big(H_n^2 C_{s_h}(C\sqrt{p})^{(s_h+1)/2}\big).$$

Also charge any reference refresh or excessive replica flag event during the suffix in which the completed jobs can scan the inputs. It suffices to use a space range of width $O(W + H_n + R_n + k_n)$ and time length $O(p^{-6} + H_n)$, including fixed row lags. The cost is

$$O_{m_h}\big(p^{-8}(p^{b_h+1} + C_{v_h}p^{v_h})\big).$$

Thus, off this short event, all jobs producing a read in the box use the same time-independent reference values. The prefix for the jobs has depth $O(p^{-14})$, so the short event has the claimed support. First prescribe M_s. Choose $b_h > 20$, v_h, and s_h large enough that M_s is strictly below the exponents in the displayed bounds, allowing slack for H_n^2. Fixed multiplicative factors are absorbed by the final choice of n_0.

For the long event, at the start of the common suffix inspect a snapshot of the necessary references in the spatial neighborhood of radius $R_n + k_n + O(H_n)$. Require each such register to have refreshed in the preceding $\lceil p^{-(b_h+3)} \rceil$ steps. The union bound for missing any refresh is superpolynomially small. Conditional on the refresh condition, the snapshot is a product of fair bits, independent across channels and sites, by the private-pool construction and the complement argument. All snapshot values are determined by the finite prefix. For one product-fair row, an anchor ends at a prescribed site with probability 2^{-k_n-1}. Require an endpoint in a strip of length $R_n/3$ just to the left of the whole spatial span of the box, and require no endpoint inside the span. The nearest endpoint to the left of the span then serves all sites. A union bound for endpoints inside and disjoint tests in the strip give a failure bound

$$O(H_n 2^{-k_n}) + \exp\big(-cR_n 2^{-k_n}/(k_n+1)\big) = O(p^{1/10}).$$

The nearest-endpoint conditions apply independently to the m_h snapshot rows. The probability that a strict majority of rows fail, or that a strict majority of stripe channels is absent, is bounded (for odd m_h) by $C_{m_h}(Cp^{1/10})^{(m_h+1)/2}$, plus the missing-refresh bound. After fixing b_h, prescribe M_l and choose m_h large enough. The dotwise majority over good readers sees the correct locally static function in every channel, including the channels with no single good anchor across the span. This gives the asserted simultaneous statement. The complete long test fits within the claimed support radius.

At fixed p, in a cone of an output site and through dates at least $t/2$ after initialization, the chance that any needed reference register still has input dependence is at most a polynomial in t times $e^{-c_0(p)t/2}$, with the refresh rate $c_0(p) = 2p^{b_h+1}(1-p)$, by the refresh tail and a union over the finite cone. Given agreement of these reads, Lemma 4.4 applies to the common, possibly changing, reference values throughout that part of the cone. The sufficient forgetting tests are uniform over those values, so the combined bound uses a union, with no conditioning on the reference history. Absorbing the polynomial into a smaller positive exponential rate and putting $t = s$, the value of every physical helper slot s steps after the start is the same for all incoming helper states except with probability at most $C(p)e^{-c(p)s}$, for suitable

$C(p), c(p) > 0$. This exponential bound is summable. On a two-sided iid innovation field, fix a site and date and run the auxiliary process from increasingly early dates. Borel–Cantelli implies eventual agreement for all starts at that site and date. Intersect over the countable site–time lattice. The limits satisfy the update, commute with time shifts of the innovations and depend only on the past. In particular their time-zero row is independent of fresh future innovations. □

4.4 Painted errors and long paths

In one lane put into E the data replacement flags, of density $2p$, and every site within a fixed multiple of H_n of any short or long bad center. Affected sites added by the latter operation will be called painted sites. Enlarge the paint multiplier once to accommodate the clipped comparisons in Section 6.1. When helpful one may extend the dominating raw flag field to negative times or take a subset of the declared affected sites. In the probability arguments within a lane, products concern disjoint innovation supports. Counts across different lanes use unions.

Lemma 4.6 (Joint sparse tests and finite-rank clusters)**.** *Fix the test multipliers used in the organization and cap arguments. The helper constants can be chosen so that for each rate in every I_n the sum of the failure probabilities for one prescribed $(j-1)$-sparsity test in a fixed multiple of V^j at every rank $j \geq 1$ is less than $1/4$, uniformly in their translations. For a fixed cap K, mark each z for which some centered test, a fixed multiple of V^{i+1}, fails i-sparsity for $0 \leq i \leq K$. The nearest- or diagonal-neighbor cluster of marked sites containing any specified site has an exponentially decaying max-norm radius tail. The exponent can depend on the rate and on K. The same two conclusions hold for raw flags alone at all sufficiently small fixed rates, with a common choice of threshold for a given geometry.*

Proof. Write $b = \min(P,T)$, $B_* = \max(P,T)$. For small p, let j_s, j_l be the first ranks with $b^{j_s} \geq p^{-17}$ and $b^{j_l} \geq p^{-(b_h+6)}$. Fixed scale factors can be included. There are finite exponents d_s, d_l with $B_*^{j_s} \leq p^{-d_s}$, $B_*^{j_l} \leq p^{-d_l}$ after allowing the fixed scale factors. The first exponent depends only on the geometry, the second on the geometry and b_h. At the respective scales the entire primitive support of a center event whose paint intersects a leaf fits in a bounded enlargement of the leaf. The leaf itself has side vector $2wV^j$. Lemma 4.2 with $u = 2p$, after absorbing constants into a small exponent $c > 0$, bounds all raw nonsparsity probabilities by p^{c2^i}, including bounded test enlargements.

Fix a nonsparse full box of rank $i > j$. Iteration of Lemma 4.2 produces a binary tree with $L = 2^{i-j}$ disjoint leaf boxes of rank j. There are at most C^{L-1} possible trees for fixed root. If each leaf event is measurable in a bounded enlargement of its leaf, the intersection graph of the enlarged leaves has degree at most a fixed Δ_{leaf}: after scaling the two axes each neighboring original leaf is a disjoint unit rectangle lying within a bounded box. Select at least $L/(\Delta_{\text{leaf}}+1)$ disjoint supports in each deterministic tree. If the leaf probability is at most α, the product bound on the selected supports followed by a union over trees is at most

$$\left(C\alpha^{1/(\Delta_{\text{leaf}}+1)}\right)^{2^{i-j}}. \tag{4}$$

At $j = j_s$ use as leaf event raw nonsparsity or a short bad center in the paint halo. The additional term is $O((B_*^{j_s} + H_n)^2 p^{M_s})$. Choose $D_s > 12d_s + 10$ and $M_s > \max(8d_s, (\Delta_{\text{leaf}} + 1)(D_s+1) + 2d_s)$, then choose the short-event constants of Lemma 4.5. For small p, (4) bounds raw-plus-short nonsparsity at $i > j_s$ by $p^{D_s 2^{i-j_s}}$. For $i \leq j_s$ we use raw nonsparsity plus the halo union bound directly. Now choose $D_l > 12d_l + 10$ and $M_l > \max(8d_l, (\Delta_{\text{leaf}} + 1)(D_l+1) + 2d_l)$, then choose the channel count of Lemma 4.5. At j_l the raw-plus-short probability is smaller than any prescribed power of p as $j_l - j_s \to \infty$. The long-center term is

$O((B_*^{j_l}+H_n)^2 p^{M_l})$. Applying (4) again gives full nonsparsity at $i > j_l$ bounded by $p^{D_l 2^{i-j_l}}$. A bounded cover transfers each of the estimates to each of the specified test shapes; strict slack in choosing the exponents absorbs the fixed factor. The eventual n_0 absorbs every fixed prefactor. Both admissions use the chosen data values of P, T, including P^n in the stripe period.

For the nested-test sum, bound the raw ranks through j_s by the raw double-exponential sequence, and charge short centers there by the sum bounded by $O((j_s+1)(B_*^{j_s}+H_n)^2 p^{M_s})$. Between j_s and j_l, sum the raw-plus-short sequence. Charge long centers through j_l by $O((j_l+1)(B_*^{j_l}+H_n)^2 p^{M_l})$, and sum the full high-rank bounds above j_l. All sums tend to zero as $p \to 0$ and are uniform in the test translations. This gives the stated threshold for the total.

For the cluster bound use three kinds of centered test marks: raw for $i \le j_s$, raw-plus-short for $j_s < i \le j_l$, and full for $i > j_l$. The support radius of each is at most CB_*^i with a fixed constant, which includes the multiplier from V^{i+1}. Also mark the short-center balls of radius $CB_*^{j_s}$ and the long-center balls of radius $CB_*^{j_l}$, whose probabilities are p^{M_s} and p^{M_l}. With large enough fixed C, every failed centered test belongs to the union of the marked balls, and their supports are contained in them. Indeed, for low ranks, one nearby painted center suffices for its own marked ball to cover the test center. In the middle ranks only long centers are charged that way, and all short events are included in the middle test. At high ranks the full test fits with its supports.

For each type λ let $r_\lambda \ge 1$ be its radius and a_λ its probability bound at any center. Our choices give, for all sufficiently small p,

$$\sum_\lambda C_1 r_\lambda^4 \sqrt{a_\lambda} < \tfrac{1}{2}. \tag{5}$$

For early ranks, the growth B_*^{4i} is dominated by $p^{c2^i/2}$. In the middle range $B_*^{4i} = B_*^{4j_s} B_*^{4(i-j_s)}$, where the first factor is at most p^{-4d_s}; the choice of D_s gives summable decay including this factor. The late range uses d_l, D_l in the same computation. The choices $M_s > 8d_s$ and $M_l > 8d_l$ give the short- and long-ball terms. For raw flags alone this argument uses only the first type at all ranks.

A finite path from the root reaching max-distance t has a finite covering list of marked balls. Retain disjoint balls greedily in decreasing radius. Every discarded ball lies within the triple expansion of one retained ball. A chain of distinct triple-expanded retained balls connects the endpoints, allowing unit adjacency, and the sum of their original radii is at least $t/8$. The event supports of the retained balls are disjoint. For an ordered list of types, the first center is at distance at most $3r_1$ from the root, and successive centers can be at distance at most $3r_i + 3r_{i+1} + 1$. The product of the corresponding two-dimensional lattice counts is bounded by $\prod_i C_1 r_i^4$, since each radius contributes twice to a quadratic count. This holds with an absolute geometric C_1, which can be used in (5). Through fixed K only finitely many types are needed, so $\kappa = \min_\lambda[-\log a_\lambda/(2r_\lambda)] > 0$. Reserve half the product probability for the factor $e^{-\kappa t/8}$. Summation of the other half, including center counts and chain lengths, is bounded by the geometric series from (5). This proves the tail bound and finishes the proof. □

5 Organization of the simulation

We give the deterministic argument for a fixed assignment of read-only operands. In applying the result to the single rule, Section 6.1 specifies the local comparison in which those operands are fixed. All coordinates in a one-rank proof are child coordinates unless explicitly rescaled. Tiles, properness and nested rectangles adapt Gács [10, Sections 10–11, including Lemma 11.2];

weak then strong computation follows the organization of his Lemmas 14.1–14.2. Ordinary relaxation, gap filling and the damage/recovery bounds adapt his Sections 12 (Healing) and 13 (Integrity of blocks). Terminal separation and end relief adapt his Sections 15 and 16, respectively.

Definition 5.1 (Observations on a tile)**.** Use all translates of the grid of half-open j-tiles of width P^j and duration T^j in physical units. At rank zero the nominal and served observations are both the physical data value. At rank $j > 0$ a matching child is a live served child of the expected clock and address residues for this grid, including legal extended addresses of that residue. On the bottom face define x_j to be the unique live state s for which the three central fifths match the encoded packet $\delta_{j-1}(s)$ and the core

$$C_* = [12A + 4q, P - 12A - 4q)$$

is occupied by matching children, allowing a single common interval of at most $2q$ erroneous or missing positions for both tests. Otherwise put $x_j = \mathsf{D}$. Uniqueness follows because two such exceptional intervals together cannot hide all three homologous copies of any one letter, since $2q < P_0$, and because packets are parsed literally (Section 3.3); the proof of Lemma 5.11 gives the details. The three central fifths form $[P_0, 4P_0)$, which lies in C_*. At child date s_c of that tile, that is s_c child ticks after its bottom face, set $g_j = 1$ if more than $P/2$ of the P child positions of its own home carry matching children, and $g_j = 0$ otherwise. The served observation is $y_j = x_j$ when $g_j = 1$ and $y_j = \mathsf{D}$ otherwise. Call a tile computing if $g_j = 1$. The nominal output x'_j is the nominal observation at the following boundary of the same home.

Use the same-origin child grid inside each tile. The bottom observation uses child date 0, and the service test uses child date s_c, with $0 \le s_c < T$ since the service date lies in early computation (Section 3.2). By induction the physical queries of x_j, y_j stay inside the spatial tile and have lookahead bounded by $b_0 = 0$ and

$$b_j \le s_c T^{j-1} + b_{j-1} \le T^j - 1.$$

The observations are defined before choosing the sparse test.

Definition 5.2 (Properness and protection)**.** A tile is proper dead if its matching bottom occupancy fits in a span of width $q/2$. It is incoming-clear if every tile of a different rank-j origin whose time interval overlaps it and whose spatial home meets its $64AP^{j-1}$ horizontal pad is proper dead. A tile is proper live if it is nominally live and incoming-clear and has matching bottom occupancy at all positions of $[1, P-1)$ except possibly one interval of width at most $2q$. A nominal group uses the residues of x_j; a served group uses the residues of y_j. The term *tribe* will refer to either such group when the observation being used is specified. A protected test includes a symmetric collar of $m = \lfloor w/5 \rfloor$ neighboring owner widths and periods. Every internal shifted tile query uses its own requisite work history in that collar.

Proposition 5.3 (Organization induction)**.** *For every $j \ge 1$, every translated rank-j grid and every data history with j-sparse deviations from the prescribed rule in a protected test, assume the lower layers are ordinary and unhalted. Use the true read-only operands. At every protected query with the requisite work and collar:*

(i) *Each tile is either nominally live or proper dead. Nominally live tiles of different geographical origins cannot overlap. The middle of three matching nominal positive tiles is proper live.*

(ii) *A proper-dead or an incoming-clear live input has a proper nominal successor. A proper-live input is served.*

(iii) *The computing tile has* $x'_j = F_j(y_-, y, y_+)$*, where* F_j *is the prescribed rank-j local map, and the noncomputing tile has* $x'_j = \mathsf{D}$*. For either kind of dead issue the actual matching occupancy at the new boundary is confined to the exceptional column and is empty if that column is empty. In early Grow it has no genuine extension apart from a short transient from that column. For a positive issue the three central payloads are exact outside the column,* C_* *is occupied off a single interval of width at most* $2q$*, and surviving incumbents outside the column have a true positive decision. One common interval of width at most* $2q$ *contains both the support and tested-payload exceptions.*

(iv) *In a stencil of nine proper incoming-clear inputs, the central output equals* F_j *evaluated on the served triple, including a requested birth.*

At rank one the child transition on an unflagged update is the actual native rule. The cut assertion (iii) *also holds when* F_j *is identically* D *because of its true halted literal. In a wholly* $(j-1)$*-sparse test the rank-j exceptional column is empty.*

For the proof, restrict the flag set to the buffered domain and extend it by zero. By the count in the proof of Lemma 5.4, every query of the one-rank argument at a tile, including lower protected calls, service lookahead and shifted-grid rounding, lies within 611 owner units of the tile, hence inside the tile's m-collar enlarged by one unit. That region has extent $2m + 3 \le w$ owner units in both axes, so the translate Q of wV^j with the same lower corner contains it. Remove the single translate J of $3wV^{j-1}$ that Definition 4.1 permits in Q. If the footprint of a lower call meets J, its tile and next row lie within $m + 1$ child units of a point of J in space, and between $m + 1$ units before and $m + 2$ units after such a point in time, up to an error below $2(m + 1)/P$. The child-grid cells meeting this enlargement of J form one common column and time interval $K \times S$, with spans at most $3w + 2m + 4$ and $3w + 2m + 5$, both less than d. The common span $3w + 2m + 4 + 2 \cdot 608 < d$ of Appendix A allows padding by $m + 2 + 608$ child units on every side; it contains this $K \times S$, and its term $2 \cdot 608$ is fixed slack. Lower calls whose footprints miss J use the lower-rank induction with the actual restricted history. If the restricted deviations in Q are already $(j-1)$-sparse, take $J = K = S = \varnothing$; this is the convention for an invocation that needs no deletion. No raw substitutions of smaller rank are thereby removed, and all actual deviations, including those in J, stay in the history. At the base one uses ordinary native transitions off flagged sites.

The lemmas below establish the implication from rank $j - 1$ to rank j. The regularity used within a structural step is always that of the lower rank. First start ordinary work from an arbitrary input and obtain a common full seed. The two-flank clearance takes a specified old interval as hypothesis, independently of any new issued cut. Establish weak descriptions at clean centers and at general centers, then the strong issue from preceding weak descriptions. Next classify the issued faces and prove genuine source attachment, density, backward inheritance and overlap exclusion. The end repair applied to a proper current input gives service without requiring a positive successor. For a proper-dead input with a requested birth, advance gives a computing target. The strong issue gives the prescribed live result, its incoming clearance follows by the backward inheritance implication, and the positive-successor part of end repair establishes properness. The clearance and startup parts of advance do not use that properness. For a halted literal stop at its strong cut and do not induct through an ordinary layer above it. The additional two-period conclusion will be used for the initialized descent.

A component refers to one exact extended source during Grow and to one ordinary home after normalization. The words *retired strobe* refer to a sampling date after a complete provisional retirement, positioned relative to the scheduled true instruction as specified in each lemma. Ordinary Heal means a fresh complete cycle of provisional attempts and removals, with the in-home restriction. In the terminal withdrawal calculations r denotes r_{last} from Section 3.2. Write $C = [12A, P - 12A)$ for the ordinary core before the $4q$ reserve at the

issuing cut. The local bounds below may use additional quantities defined in the individual calculations. At searching ranks F_j is owner-specific through the true external words; no step compares maps of different owners. The served-ancestry consequence of Proposition 5.3(iii) is that every protected live output off the exceptional column (at rank zero: on an unflagged update) has an actual served incumbent or donor within radius one at the preceding row, with the same parent clock and address residues; the native transition and Lemma 5.4 give this consequence at each step.

5.1 Protected structural calculations

Lemma 5.4 (Lower structural transitions and the size of the protected test)**.** *On a protected unexceptional structural step a component is measured inside one exact extended source during Grow, or one ordinary home afterwards. On a step with no birth or ballot and no address normalization, a strict interior of such a component survives. On a recruitment step an apparent vacancy can fail to take the proposed birth for one step if its old nominal state is live but unserved.*

Proof. A strict interior is between two positive matching served neighbors with consecutive exact integer addresses. Their three nominal states are live and matching since any served live state equals its nominal state. The lower triple forces the middle proper. It is therefore computing, and its next nominal state is its native transition. Its own left and right bonds have grade two. Purge has no weak right bond for a kill, Shrink sees no zero face, and temporary and true withdrawals likewise have no zero face. Other nonballot maintenance operations hold a live incumbent. Proper successor regularity and proper-live service at the lower rank turn this positive nominal result into an actual served result. At grade one, even an occupied adjacent true home cannot provide strong Purge backing. For example if address $P - 1 - k$ is vacant for $2 \le k < q$, the k remaining sites immediately to its right receive zero in the q propagations and the rightmost has a weak right bond at the wrap. They can all die in the killing pass.

If x is live but y is D, the definition gives $g = 0$. On this step x' is D by the noncomputing clause even if the native map on y requests birth. The actual-support clause at the new cut makes this tile proper dead. Its served value is also D. This accounts for one failed transition at that slot. A subsequent recruitment still needs its full lower advance hypotheses. If x' is instead live from $y = D$, then $g = 1$, so $x = D$. Classification makes this old slot proper dead. It has a proper successor and the live output is served. An undisturbed clear preparation row followed by a complete Purge before the attempted birth is sufficient to remove the first obstruction as shown in the buffered clearance proof.

Every nondead nominal output has a genuine previous served incumbent or donor within radius one, by the noncomputing clause, exact computation, and the three-dead quiescent rule. Its spatial residue and clock identify the same parent phase on tracing this edge backward. A proper next state is served when live. A matching lower member uses its extended address modulo P, including in deciding the occupancy of a displaced geographical home. Undisturbed means incoming-clear in the $64A$ pad.

When invoking classification on a shifted home, take its own preceding work, its three suppliers' preceding transactions and their one-neighbor input lanes. Adjoining inputs or gates are queried inside the outer protected box. Each work interval includes its own service lookahead. The transfer of disturbance requires one preceding predecessor comparison and the two issuing epochs for the overlap argument. In a single terminal comparison only the two adjacent authentic root cores are relevant to a remote marker: a farther source cannot reach it across an intervening core and never has an overhang exceeding $1.1P$. The marker, two source origins, and $300B$ and $0.2v$ margins occupy fewer than eight parent home widths

since $(300B + v)/P$ tends to zero. Paired seed and two-period tests cover six input homes and two further periods. The directed recipient checks nine home widths and at most five source homes for each queried displaced home. These are finite shifts in units of the current parent tile; $608 < m = \lfloor w/5 \rfloor$ at the global choice $w \geq 10^{10}$. Their lower-rank calls take their own m lower tiles of margin, of width m/P and duration m/T in current units, and are not enlarged by m *parent* tiles recursively. Large T/P in the unrestricted cap ancestry cone is accounted for separately in the centered tests of Lemma 6.3.

Use. The no-split argument, all stencil birth calls, the actual cap ascent, and the initialized two-period service assertion use these conventions. The exceptional numerical allowances are specified where the corresponding geometry is proved.

For the localization count, a direct same-rank query spans at most 16 additional owner widths: the largest directed stretch is $[-2P - 8B, 3P + 8B]$ plus two widths on each side for the potential prior source in a normalized home, four lower guard positions and the neighboring home clipped by rounding. Six seed homes, two adjacent message routes, lookup at the three central positions, and all terminal, $300B$ and $0.2v$ margins fit this same bound. The rear-wall radius $3v\ d_0 + 2w^6 + 12d_0 + 20$ and the polynomial early-seed cone are $< P/10$ for sufficiently large w; these are the unrestricted front cones. A direct temporal lookup takes at most six periods: two past weak levels with the neighbors' computations, the prior work for disturbance, and two periods for the initialized descendants, including a rounding period where necessary. Count the following ordered within-rank calculations separately: observed rows, ordinary relaxation and elementary wound response, common full seed and transport bridge, clean weak issue, general weak issue, strong issue, classification, authentic source, thin trace, wide trace, density, disturbance inheritance, overlap/triple, prospective contender exclusions, common withdrawal calendar, current service/proper successor, advance, two-period descent, and the protected query. The elementary response is applied first with a supplied long interval and subsequently with the actual issued source. Allow two extensions per entry for the pairwise comparisons or the two charged searches. Direct displacements therefore fit within $19{\cdot}2{\cdot}16 = 608$ owner widths and $19{\cdot}2{\cdot}6 = 228$ owner periods; no within-rank request starts the same induction all over again. A complete actual lower test needs m lower units, and the finer residual margins sum to at most $2m/P$ and $2m/T$ current units. These are less than one and the entire region fits in the own collar. In child units a common enlarged exception has both spans at most $3w + 2m + 4 + 2{\cdot}608 < 4w + 8$ for the chosen w; this even pays extra fixed stencil offsets if grouped in advance. The cap uses separately centered tests with the larger $2T$-child ancestry zone, which is not charged to one ordinary owner's collar. □

5.2 Ordinary precommit

The ordinary precommit adapts Gács [10, Sections 12–13].

Lemma 5.5 (Ordinary relaxation from arbitrary rows)**.** *Work from an arbitrary normalized input in one home, with the lower transition convention and the common exception. An interior vacant band of width $4B + 4$, disjoint from K and the two true endpoint sites, at a retired strobe persists at subsequent retired strobes during withdrawal and ordinary work until the ballot, with a vacant center throughout. A home capable of genuine computation has a single genuine interval with both end gaps smaller than $16B + 2d + 24$ at a clean sample at most $50Ad_1$ after normalization. At that sample the $3A$ and $6A$ threshold positions and their neighbors at both ends are proper strict interior sites. The other branch is erased except for immediate raw-only material.*

Proof. Let $H = [0, P)$ in child coordinates. Sample after a complete provisional retirement, immediately before each actual Shrink(0) instruction. Thus the step after the sample is a

true Shrink. All following samples include the q intervening Main cycles with full Temp resets, B attempts and B removals. At a Grow-to-ordinary transition normalize first. Allow one whole initial cycle to clear marks and start fresh Cn. All statements concern one clock and address residue group actually supplied by the lower level. Lower optional omissions and failed births are allowed. No second source can put a positively consistent ordinary member in H across a true wrap: the proposed address must be in the donor's own ordinary home.

Empty band. A strip of $4B + 4$ consecutive vacant positions at a retired strobe, disjoint from K and ordinary endpoint positions, retains an empty center on every subsequent date of the final Grow withdrawal and ordinary work until the ballot. It returns to vacancy at each complete retirement. No true Grow births occur on this interval. In each probe all entrants into a currently empty strip are newly marked, except raw assignments which were excluded. At most B sites can be entered from each edge because recruitment moves by one per attempt. The two entering portions therefore never meet. At each removal the furthest surviving entrant from each side, if any, has no positively graded neighbor towards the center. A positively graded neighbor would itself belong to this same clock and residue group. This entrant is marked and its facing site is not its true endpoint (in Grow all potential true endpoints have the same pair of residues modulo P). It dies. Inducting through B removal rounds eliminates the entrants, even if an outside original donor has died. Holes in an entrance only accelerate this bound: the furthest coordinate retreats at least one until its original strip edge. Each true withdrawal, normalization, Shrink and ordinary computation between retired probes can only erase members of the group. A ballot is excluded. From an arbitrary partial phase, an initially empty $8B + 8$ band contains an empty central $4B + 4$ band at the next strobe: from either edge there are at most B remaining attempts before a complete retirement. The central protected site remains empty up to this next strobe, even if old flags were arbitrary outside.

Clean gaps. In an ordinary home a clean pre-ballot step cannot split the descendants of a single component (the no-split convention of Lemma 5.4). In an incumbent's strict interior the two exact consecutive bonds and properness force a hold on the structural instructions. On a birth, an old end with an inner neighbor and an actual outer child stays between matching nominal successors. The nominal triple forces the old end served on the new row. A singleton with two children likewise stays served, and with one child its set of descendants remains connected. On a clean nonbirth step only ends of the same ordinary home/component can disappear. Components can merge but cannot pass one another: a holder cannot accept a proposal at an occupied position, and there is no swap through a nearest-neighbor pair. These arguments concern actual descendants up to the next step and allow an entire component to vanish.

Take two adjacent components still separate at a later inspection. Propagate their component classes back one sampled interval; if either descends from several components, keep the rightmost contributing old component on the left and leftmost contributing one on the right. They have ordered, disjoint ancestors. If there are intermediate old components, any that contribute to neither final class have vanished or have merged into another class. In the latter case a surviving intermediate class would spatially separate the supposed adjacent two, which is impossible. Fix any vacancy between the two ancestral groups at the previous strobe, for instance the gap immediately after the last contributing left component. Both extreme endpoints surrounding an old internal vacancy see a grade-zero face and fail the outward-exemption test on that face. Shrink(0) removes each, enlarging the empty interval. If an entire edge bank dies, the enlarged vacant positions still exist. Before the next sample no old material moves into an empty coordinate; only marked provisional attempts can enter. To persist at a retired strobe, such a frontier directed into the vacancy must either reach a true end, or meet a matching incumbent or a facing probe on the far side. A true end cannot

be reached while the other surviving ancestral class still separates it from the end of the home. A meeting of the two final ancestral classes is impossible because they could never separate again. A meeting with a formerly intervening third component makes that component an ancestor of the joining class if the joined interval has any surviving descendants: once joined, label the entire interval by the union of its component labels. Consequently choosing the gap immediately after the last contributing left component excludes this possibility for persistent advance from the left. On the right, choose the gap immediately before the first contributing right component if there are intervening disappearing components. A discarded middle component may be crossed only after it has disappeared; it cannot leave a lasting supporting front for either side without becoming a contributor. It follows that the final gap contains the chosen old gap enlarged by at least one cell of the bordering contributing component on its populated side. Equivalently its left edge is no farther right than the last old contributing left edge minus one, and its right edge is no farther left than the first old contributing right edge plus one. This latter statement follows separately for each class by looking at its outermost probe at each of the q complete cycles. Without a join to a new contributing class its marked outside extension has a zero unexempt face and the B removals peel it completely. Thus an internal gap at a later sample has an internal ancestor gap at every preceding clean sample; tracing repeatedly gives width at least 1 plus the number of clean Shrink(0) calls. A single genuine interval, once reached, is preserved in clean ordinary work up to extinction or a cut.

Early certificate. After the initial cycle take two consecutive trial stretches of $20A + 1$ Shrink spacings, leaving two complete spacings between them. S has length at most $d < d_1$, so at least one stretch is wholly clean including boundary samples. Its terminal sample is within $50A$ d_1 of normalization since $2(20A + 1) + 2 + 2 < 50A$. If it has more than one component it has an internal gap larger than $20A$. Such a gap contains a sterile band of width $4B + 4$ disjoint from K and from both true ends: remove K and the two true endpoints, splitting the interval into at most four pieces; $(20A - d - 2)/4 > 4B + 4$. The same holds for a true-end vacant interval of width $16B + 2d + 24$: after removing K and the one true-edge pair at the boundary there remain at most two nonempty pieces with total length at least $16B + d + 22$. One contains a sterile band. On every clean stage-zero inspection the strobe point nearest to that band on either side, if any, disappears. Its face toward the spine is broken at a strictly internal depth; the exposed site is eligible for Shrink. No forward provisional gain persists beyond the newly recessed strobe frontier: there is no other same-tribe material between it and the band, the band's center never becomes occupied, and no true endpoint lies between the frontier and the spine on its home side. Each clean inspection reduces the available distance to the far end of H by at least one per such side. Hence $P+1$ consecutive clean inspections suffice to erase both sides, including a last singleton abutting the opposite true end (its spine-facing face is internal). More than $3P + 10$ calls remain because $\lfloor 3^w/q \rfloor - 50A > 3P + 10$. If the pulse is still ahead, removing its touched spacing and the two adjacent partial spacings leaves more than $3P + 6$ calls in at most two stretches. At least one stretch contains $P + 1$ consecutive clean calls. If this stretch follows S every remaining occupant dies. If it precedes S the home is empty at S, and every later occupant traces to K. Before the next complete initialized clean Purge there are at most four recruitment instructions in this phase, so its whole span is at most $d + 8 < q$. With fresh Cn, after q propagations it has no positive right tip, and all its components disappear from the right within q killing ticks. Ordinary recruits from outside H are forbidden. Thus the only late alternative on this erased branch is immediate raw-only material.

A home capable of a genuine productive computation must therefore have had a single genuine interval at the above sample, with *both* true-end gaps smaller than $16B + 2d + 24$. Each $3A$ and $6A$ threshold and its two neighboring positions on both sides of H are strictly interior, since $16B + 2d + 25 < 3A$; at this clean sample each is the middle of a matching

nominal triple, hence proper by the rank-$(j-1)$ clause of Proposition 5.3(i). Their proper incumbents then follow every native live hold by lower proper-successor and service. Any later genuine reoccupation starts with proper birth. This is the input to the one-event ledger, independent of nominal separation at the new rank. □

Lemma 5.6 (Clearance between held flanks)**.** *Suppose at the row just before S there is an unbroken canonical stretch between two intervals of 24 old positions held continuously from that row to t_E. From t_E to the contemplated fresh Heal attempts suppose that each flank remains one canonical interval, has at least 24 positions, and in one tick its left edge can move right by at most one and its right edge left by at most one. Enlargement when it absorbs another compatible interval is allowed. Stop when flanks meet or either disappears. For the questioned nine-slot cross-section assume that its furthest spatial coordinates extended by three on either side, across the current and four preceding dates, are within or between the flanks, allowing use of their innermost 16 positions as a buffer. Then every foreign-or-wrong-phase served obstruction in that middle region descends from $K\times S$; on and after the last affected date plus $4q+10$ there is no such obstruction. The same conclusion holds simultaneously for all the foreign nominal incoming stencils at dates from four later, until the stopped comparison ends. For lower native clean births a previous clean row gives properness of the input slots; the recipient's possibly improper old nominal incumbent may block one attempt. If the walls come from an earlier common full seed of a pair of homes, with an unmarked strip of $q+9$ proper old sites disjoint from K on each side of an interior pulse (in the same home or, when K is within $q+10$ of the true wrap, flush with the seam in the neighboring full home) and the survivor cores held, the paired corridor between those cores is repaired in one pass within $d+4q+14+2d_0 < w^6$ ticks after the start of S, with proper interior triples throughout the restored corridor.*

Proof. Use Cartesian child rectangles of unit width and time. A shifted child grid has offsets a,b in $[0,1)$; its tile with indices i,t has rectangle $[i+a, i+a+1)\times[t+b, t+b+1)$. On one fixed shifted grid, any nominal live output has a served live parent in the preceding row at spatial index distance at most one, and a served live state is nominal live. On the same geographical grid a different native phase cannot occupy a slot already occupied by a canonical state. On distinct geographical child grids the protected lower overlap exclusion forbids intersecting tiles both nominal live. All coordinates below are at that child scale. An incoming pad for a queried lower child is $64A/P < 1$ unit. At rank one there is only the one physical grid and native mandatory births can be tested directly, without a lower properness call.

Wall test. For times up to t_E choose a fixed interior site in each 24-site old run. Afterwards track an interior site in each flank by clamping the previous site into $[\text{left}+8, \text{right}-8]$, inclusive. This interval has length at least 8. Its left boundary can move right by ≤ 1, its right boundary can move left by ≤ 1. The selected site therefore moves by at most one per canonical tick, including when a flank enlarges abruptly. It has at least eight uninterrupted occupied positions on both sides. For the region outward of the selected site but inside its own held flank, every possible competing served rectangle meets a canonical occupied rectangle and is excluded directly. For the region between the two selected tracks, consider any traced foreign tile path through successive rows, following a served parent each time. At the observed tile choose a time inside its intersection with the canonical queried time row. For the served parent of a nominal tile starting at σ with $t-1 < \sigma < t+1$, its own interval $[\sigma-1, \sigma)$ meets $[t-1, t+1)$, so choose a comparison time in this intersection (for $\sigma = t$ use the preceding row). Follow this same fractional offset inside every earlier tile of its lineage, subtracting one per backward step. Compare its spatial midpoint to the selected index on the canonical row containing that time. Successive comparison times differ by one and the

index of the canonical row increases by one. Only the current and previous canonical rows are queried. The spatial midpoint moves by ≤ 1 and the selected index by ≤ 1; their signed difference changes by ≤ 2. A path switching sides therefore includes a midpoint within 2 of the selected position. Every such foreign tile intersects an occupied canonical position in that row (half-site boundary rounding costs at most one). It is forbidden by lower exclusion, or by unique occupancy on the same grid. The same calculation handles a tile crossing the old bottom row just before S: the whole segment between the old selected tracks was occupied. Thus each permissible traced foreign lineage must encounter the exception. This also proves that none can be supplied from beyond a shortened but still enduring flank. After a join the comparison ends.

Uniform killing time. Follow a trapped age/residue class backwards beyond its last raw ancestor. On uncompromised steps a changed spatial position entails a successful birth; lower omissions alone cannot produce a new position. A fixed native clock can birth only on a handle instruction, once per $H_0 = 2q + 2$ ticks or less often. From first S through last $S + 4q + 10$ there are $\leq d + 4q + 11 < 3H_0 + 1$ ticks for $w \geq 100000$; even upon rounding both shifted temporal boundaries, allow four opportunities. Each descendant is therefore within K expanded by six shifted positions on either side, counting two more for the spatial rounding and the old site versus updated site. Its total width is $\leq d + 12 < \mathrm{q}$. For each alien clock after the last affected update, at most $H_0 - 1$ further ticks precede the next clean Purge initialization; the entire initialized Purge uses $2q + 1$ ticks. Two rounding rows suffice to translate these dates back to the canonical time grid, giving a common upper bound last $S + 4q + 6$. No outside immigrant can supply strong backing during this operation. Wrong phases, residues or geographical grids at the two canonical walls cannot supply a strong link to the trapped alien. In a fixed alien age class with distinct ordinary sources, only exact consecutive addresses within one root can give grade two; the whole trapped age class in any case fits $d + 12$. If a normalization occurred earlier, bound the union of the original footprints, all in this same column. Each surviving $\mathsf{Cn}{=}1$ after propagation would require q strong left predecessors, which do not exist in this short span. The right break eats every remaining zero component in at most q killing ticks. No birth occurs during the Purge. The common date bound holds separately for every shift, with no sum over the number of origins. After this time a new alien interior site would need another raw pulse or a path across a held wall, neither of which is available.

From served to nominal. A foreign nominal non-proper-dead blocker is nominal live by lower classification. Its preceding served parent is at most one spatial shift index away and one time index earlier. A lower stencil querying four neighboring slots and a < 1 pad can intersect a shifted tile that extends by one further unit; use its parent, then one additional prior row to certify lower proper successor. The total extra spatial reach from the outermost stencil site is < 1 pad $+1$ tile width $+1$ parent step, hence at most three slots; the corresponding temporal reach is at most four rows, including the previous-row proper-successor test and both rounding endpoints. All these areas are within the buffered wall hypothesis on rows up to and including the queried row; no prospective future wall beyond a first actual join or disappearance is required. Canonical undisturbed live inputs at that previous date acquire proper successors. Canonical nominal D inputs are proper dead by lower classification, even if a foreign native residue lives in a different slot. A geometric competing origin is gone by the previous calculation. Thus the required nine inputs are simultaneously proper and incoming-clear at the actual recruitment step. For an old vacancy that could conceal a canonical unserved nominal-live state, any preceding nonbirth step with no canonical incumbent there has nominal output D regardless of its service indicator. The initialized Purge just before every attempt provides such a preceding date or a true prior matching incumbent. If the latter is a proper incumbent and outputs nominal live, its positive

successor remains proper and served. An initially improper incumbent can be omitted for one attempt. Consequently the exact native birth into a qualified clean vacancy is served.

Previously full paired seam. When separated walls originally form a full seeded pair, static inner backings suffice. On each side of an interior pulse choose an unmarked contiguous strip of $q + 9$ proper old sites outside K in the same home. Its $q + 1$ interior survives the at most two Shrink calls before the next requested complete pass, by the fixed-backing proof in Lemma 5.8. Take a one-tick properness and edge reserve. If K is within $q + 10$ of the true wrap, an inadequate strip on the near side of the wrap is replaced by one in the neighboring full home, flush with the seam and disjoint from K. If K straddles the seam take one strip in each home outside it. If a fully backed hole lies deep inside one home take both walls there. At the resulting in-home faces one reaches each damaged position or its legitimate true end from a $q + 1$-backed strip in at most $q + d + 32$ attempts: the short old near-wrap portion is $< q + 10$, K has $\leq$d positions, the two-possible-Shrink and stencil reserves use $\leq$20 more, and at most two attempts are blocked by initially improper nominal incumbents. For a deep cut its width including these reserves is $\leq d + 20$. Outside the two stationary $q + 1$ backings the source core used for transport lies in the already held part of each computing home. This is applied only to seams in the corridor between the two inner cores of the full pair; its two exterior true ends are irrelevant.

Wait for the common clearance above, and use the first subsequent reset / growth pass (at most another d_0 until its reset and d_0 until completion). On the B growth attempts all exposed in-home front sites adjacent to the hole have legal proposals, unoccupied recipients see either matching age or D one step farther, and all nine inputs are clear as just proved. A legal competing bid from the opposite bank in the same home has exactly the same canonical address at the recipient; a bid across the true seam is illegal for that bidder's ordinary home. Existing correctly occupied positions along an advancing front need no replacement. During this pass no real Shrink occurs, every new birth is proper, and each connected extension containing the old backing survives each intervening Purge: its rightmost site has q strong predecessors, and all earlier sites have a strong right neighbor. Allow at most two attempts blocked by initially improper nominal incumbents. Every other attempt fills the next vacant site of each backed front or finds that its target has already been reached. Since $q+d+32 < B$, the two legitimate ends or the joined same-home fronts are reached in one pass. Between the restored full seam and the two cores all interior lower triples are proper; a wrap is positively bonded but each home supplies its own strong Purge backing. Subsequent clean operations preserve that corridor. All response dates fit $d + 4q + 14 + 2d_0 < w^6$. A later good input transfer and all its shifts require that their whole transaction interval is disjoint from this charged slab, not merely that its first read date is disjoint.

Dependencies. The unconditional application to a general old single interval requires the two enduring, sufficiently buffered flanks and a bound on their cavity width until the chosen pass. Lemma 5.8 and Lemma 5.12 discuss the marginal cases. The paired-seam result additionally needs an earlier actual common full seed as supplied by Lemma 5.9, and the survivor cores of the computing homes. Thus this is a proven barrier/clearance mechanism conditional on the explicit spatial and temporal wall premises. □

Lemma 5.7 (Local response from a supplied macroscopic interval)**.** *Assume the lower-rank hypotheses, the structural transitions of Lemma 5.4 and the held-wall conclusions of Lemma 5.6 (not its paired-seam clause), with one exception $K\times S$ of spans at most d. Supply an interval I of length at least $2q + d + 44$, consisting of one exact extended source during Grow or of one ordinary home after normalization, at its last fully retired strobe before S. Assume that its strict interiors are proper and that its evolution is clean from its latest scratch reset (a* Temp *reset, or the fresh reset of a positive vote whose live decisions keep its proper incumbents outside K) until S. Times are measured from the start of S. Put $t_c = 2d + 9q + 40$;*

the selected pass is the first full B-growth pass after t_c, or the remainder of a pass then within its first ten growth instructions. Stop at normalization. A ballot during the response does not stop it, provided that (i) every ballot decision differing from the native one lies in K, (ii) after cutting the pre-ballot old occupancy at K, each remaining old panel of length at least $q-1$ loses at most one site per side and no interior site across the ballot, (iii) scratch is freshly reset at it and the handle cadence continues through it, and (iv) at most four birth opportunities occur between the start of S and the first complete clean Purge that follows both S and every non-native ballot decision, and the last of these exceptional dates plus $4q+14$ is at most t_c. Otherwise stop at the ballot. A reopening of the same column K at a later ballot after a completed restoration is a second wound, to which the lemma applies with the pre-ballot interval as I and the ballot's non-native decisions in K as the exception. To the first complete clean Purge the bounds are $25d+100$ new or reassigned positions and $3d+12$ components, with a continuously held 24-site old run in every surviving component of length at least $q-6$. A persistent internal cavity has width at most $B+30q+40$ and stays inside K enlarged by $2B+40q+40$ on each side. Within $2d_0+10H_0<w^6$ ticks after the start of S it is filled or a distal flank disappears, provided that time remains before the stopping boundary. On a truncated response only localization is asserted. During Grow, every run of $q+1$ old proper positions with ordinary addresses, disjoint from K, carried continuously through the latest reset and trimmed by one site at each end, loses no site for the rest of that Grow: no Purge with arbitrary exterior, temporary retirement, native attack or true Grow withdrawal removes any of its sites. A run of extended addresses loses at most one outer layer per true Grow withdrawal. The same conclusion holds in ordinary precommit with no ballot during the response, using in-home bids and at most three real Shrinks in place of the Grow losses.

Proof. The structural facts used below are those of Lemma 5.4; the only clearance call is the held-wall part of Lemma 5.6.

Count to the first clean Purge Fix one exact source and clock, before its ballot or a change of address convention; in ordinary work this is one ordinary home. Immediately before S its actual matching occupancy is one interval, possibly empty. Let t_E be the end of the first complete initialized Purge starting after S. With $H_0 = 2q+2$ an instruction occurs only once per H_0 ticks. Starting after the last exceptional transition one waits at most H_0-1 transitions for the first clean initialization and then takes $2q+1$ transitions through its last killing tick. Including S gives at most $d+4q+2$ transitions. Since $d<q$, a sequence of four distinct instruction dates (first and last at distance $3H_0$) cannot fit entirely inside these transitions. Keeping four opportunities is conservative even if both bounding rows are counted.

During this time let k be the number of canonical recruitment opportunities elapsed; for purposes of the proof count a full instruction even when it does not recruit. Apart from K enlarged by k positions on both sides, each of the two remaining half-lines meets at most one interval of the actual phase. At $k=0$ initially this follows by intersecting the single initial interval with the two rays. Consider one future update. On a nonbirth tick an old interval within a ray can lose only endpoints. Its strict interior has two matching neighbors, hence is proper by the lower nominal triple property, computes its true update, and persists when the native rule holds it. Purge kills only at a right break, Shrink only at a zero face, Temp retreat only at a zero face, and a true Grow withdrawal only at a zero face (Lemma 5.4). On a birth tick clip each ray by one additional inner position. A child of a site outside the old ray cannot reach the new ray. Descendants of the one old ray interval still constitute an interval, possibly empty. An old interior has a proper positive hold. If a previously exposed tip with an old next-inner neighbor issues an actual outer child, then the old tip, its inner neighbor and the child all have matching nominal outputs (the first two have actual old

served states and compute positive native holds). The middle of this next nominal triple is therefore proper and served. A singleton with children on both sides has the same argument; with only one child its descendant set is connected whether or not the old tip persists. A right- or left-end loss with no outer child adds no component. At the old ray boundary, a previously excluded parent's child can occur at most in the one newly clipped position. This proves the induction.

Inside the enlarged band there are at most $d + 2k$ occupied coordinates, giving at most this many components there and at most $d + 2k + 2$ components overall. For $k \leq 4$ this is $\leq d + 10 \leq 3d + 12$. On each of the at most four opportunities a clean birth can only adjoin one of the at most two exposed ends per canonical component at the previous row. At most $8(d + 10)$ assignments are therefore clean new births, counted with multiplicity. All remaining forged canonical sites, or lost-and-recovered canonical sites whose recovery is forged, occupy one of d coordinates. At most $9d + 80 \leq 25d + 100$ distinct coordinates in total can be newly attached or forged. If a surviving component at t_E has length at least $q - 6$, delete these coordinates. Deleting k'$\leq 25d + 100$ coordinates leaves at most k'+1 runs. If every run had length at most 23, the total initial size would be at most k'+$23(k'+1) = 24k'+23$. The inequality $q - 6 > 24(25d + 100) + 23$ holds for $w \geq 100000$ by division by w (right-hand linear coefficient 2400). Hence some 24 consecutive original coordinates in the component have survived continuously since before S. Each lost original site can only have returned by one of the counted assignments.

A canonical fragment with no outside flank at the start of that clean Purge is erased if its whole strong-link component has $\leq$q sites. Along any rightward string of strong links a surviving Cn=1 after q propagation ticks requires q strong predecessors on that interval: unfold the copies of positive left Cn one tick at a time. The finite string lacks such predecessors for every site if its length is at most q. No new site is born in the Purge. Starting from the resulting zero flags, the rightmost remaining site of each component has a weak or broken right link and is removed at each killing tick. Optional lower omissions only shorten the strings. This proves extinction within q killing ticks. A genuinely raw-only independent phase confined throughout to K expanded by four has width $\leq d + 8 <$q. A newly joined flank instead belongs to the resulting whole component.

Run protection. A run of $q + 1$ old proper ordinary positions, carried continuously through the latest reset except for the particular phase boundary being analyzed, cannot lose a site to a Purge with arbitrary exterior. In q propagations the last site's q predecessors all lie in this run and were positive at reset. During the killing ticks induction from the old row preserves the last site by its positive Cn and every earlier one by its strong right neighbor. This also applies starting within a partial sweep when the run was present for its clean prefix. A raw alteration only outside this run cannot place a zero signal inside except by the same leftward-to-rightward propagation of at most q links counted from the reset. A positive cut is followed by fresh initialization. Native attack never replaces incumbents; true Grow withdrawal cannot remove an ordinary address. Originally unmarked proper positions cannot die in a temporary retirement. These assertions are simultaneous invariants of the full proper run, via lower proper-successor service. They prove protection throughout Grow of each original $q + 1$ run of old proper ordinary positions away from K, after trimming one site for initial endpoint properness. If J hits an extended arm, one true Grow may remove one outer layer of a run of extended addresses; trimming three layers per exposed face in the short interval considered below pays up to three true Grow withdrawals. Each still existing strict interior has a nominal matching triple and is therefore proper. On the native step before a Temp reset, old strobe positions can still be marked from the preceding completed retirement. No further temporary retirement happens before the next unconditional reset, so these marks give no additional loss at old positions outside K. If J hits immediately after a positive vote,

the vote resets all scratch and keeps the proper ordinary run. If it straddles the vote, the run had a common continuous ordinary history through the vote and the fresh first Grow reset.

Bookkeeping at a single wound. Write $H_0 = 2q + 2$ for the handle length. From the last exceptional transition to the end of one full clean canonical Purge takes at most $4q + 2$ ticks. The same bound applies independently to the first full clean Purge on each displaced grid; allow d more ticks of exceptional rounding and four previous rows for removing nominal ghosts. For definiteness put the clearance time after the start of J at $t_c = 2d + 9q + 40$, an upper bound on these waits even if a terminal reset restarts a purge at a phase boundary, since Purges do not straddle the boundary on the fixed clock. If a reset restarts, the nominally preceding handle still ends at its proper instruction, and initialization starts at the very next handle. Take the first full B-growth pass after t_c, or the remaining part of a pass already in its first 10 growth instructions at t_c. There are at most ten skipped instructions in this latter case. To see that no full extra growth pass occurs between the original interrupted cycle and this selected pass, note that $t_c < 5H_0$ for $w \geq 100000$ and begins at the start of J: even when J begins at the very last instruction before a cycle boundary, it cannot consume eleven growth instructions in the next cycle. If it begins later than ten instructions into the interrupted pass, the next pass is selected. If it begins in a cleanup or idle part, one simply proceeds to the next pass, paying at most the same ten initial opportunities. The selected pass ends, and the following complete Purge ends, in less than $2d_0 + 10H_0$ after J. This is $< w^6$ for $w \geq 100000$. At most three Main instructions intervene even when counting both ends. Close to normalization use only the valid prefix of this count and stop at the label change.

Take $[u, v'] = I$ at the last fully retired strobe preceding J in this phase. If the wound coincides with an incoming vote, use the preceding ordinary retired interval, followed by the ballot on its actual old occupants. If the raw was materially earlier and K reopens at the ballot after restoration, use the then actual pre-ballot interval for the second wound, with no new column. Before the selected growth opportunities begin, genuine descendants outside I may have advanced at most $B + 14$ on either side: at most the one interrupted B-probe pass, up to ten initial opportunities of the selected pass, and at most three Main births. Only the one interrupted pass is charged as a full pass in this count. Any part first forged in K and not advancing from a genuine donor either dies at its first full Purge or attaches by then to a genuine component; its amplitude before that Purge is $\leq d + 8$. Accordingly no enduring member extends beyond $[u - B - d - 22, v'+B + d + 22]$ at the selected start. A remote raw-only group cannot certify a more distant old tip here. If J is spent early in the previous ordinary work, ordering starts anew at the genuine restored interval.

Write $K = [a, b]$. Where K cuts two original pieces each of length $\geq q + 22$, there is a $q + 1$ fixed old proper subrun on each side reaching within 10 of K. For example use $[a - q - 10, a - 10]$ and $[b + 10, b + q + 10]$; when K is inside I these are wholly inside the corresponding piece after discarding one old extreme. The ten-site margin covers one initial omission and three possible native true Grow losses at each facing edge. Apply the run argument at each purge between these operations. Then the between-flank gap is at most $d + 20$. If the left piece is shorter (including length zero when K strikes the left tail outside I), use the first right protected run starting at $\max(u + 10, b + 10)$, whenever the right old piece is present. In this case $b \leq u + q + 22 + d$ unless K is disjoint far to the right, which is handled by the opposite orientation. The distance from any lasting distal left-hand member to the fixed run is at most $B + q + 2d + 55$, by the preceding support bound. Exchanging u with v' and left with right gives the same estimate for a short right piece. The argument uses $q + 1$ strong links ordered *left to right* for both protected runs; it does not reflect the orientation of Cn. Both old sides cannot be short when K intersects the macroscopic old interval. When K misses I and also misses its $B + 14 + d + 8$ neighborhood, it cannot hit

a genuine descendant during this time and produces no genuine split. Thus any persistent between-flank cavity has length $\leq B + 30q + 40$, since $q + 2d + 55 \leq 30q + 40$. Its entire interior stays inside K enlarged by $2B + 40q + 40$ even on partial earlier trials, since each relevant original fringe is within $B + q + 2d + 55$ of K, and one incomplete trial moves a front by $\leq$B. For a deep K use the two protected neighboring runs directly. This argument allows a genuine, macroscopically long far part outside K; only the missing interval is localized.

Durability and actual refill. Stop at first reunion, extinction of a substantiated side, or the extra full Purge after the chosen pass. The cone count to the first clean canonical purge is $\leq 25d + 100$ new or reassigned positions and $\leq 3d + 12$ components, by the count just proved, which uses only intra-source bonds and genuine birth/nominal-next rules. After this purge a wholly new separated component of sustained width at most $d + 8$ is impossible, even if pinned at a grade-one or true-end contact. Otherwise a lasting flank of size $\leq$q-2 on any intermediate row will have at most two new sites before its next complete Purge (there is at most one intervening instruction, giving at most one birth at each end). It will disappear unless a genuine join occurs. On the branch where neither stopping event happens, both flanks consequently have size $\geq$q-1 on the needed rows, including those whose next such Purge is the extra final check. At the first clean purge each contains 24 continuous old positions, by the pigeonhole proof above. The positions between these two initially stationary panels contained only the authentic source before J. Thereafter each bounding interval moves its endpoints by at most one per structural tick and has width at least q-1. Successive slices thus share $\geq$q-3 sites. A displaced lower lineage moving at most one cell per one of its own ticks cannot enter the cavity across the panel without occupying a simultaneously protected tile of the panel; shifted endpoints use two units of space and one extra row. Lemma 5.6's buffered clearance argument applies to both enduring inner faces (eight spatial guards and the four prior rows are within the panels or cleared interior). It supplies mandatory births on the actual nine-cell stencils. A genuinely competing exact source planted between these panels has only K enlarged by the four birth opportunities of its own phase for backing and dies in its own full clean purge. If its normalization or ballot intervenes, its aggregate width is still $\leq d + 8$, and the next initialized purity sweep treats that width. If one side disappears earlier the cavity no longer requires a directed fill.

The two inward-facing donors propose legal intermediate extended integers because both face the same exact source label and the legal integers comprise an interval. Where one vacant slot separates them the proposals agree on its address and age; priority is harmless. During the selected pass each lasting face makes at least $B - 12$ genuine inward births (allowing two attempts blocked by an initially improper nominal incumbent), or joins. Its inward *net* movement need not equal its birth count. The following larger allowance replaces B-3 in uses of this Grow estimate. At the start of the clean pass all but its two initially exposed endpoints are proper, and each genuine birth thereafter is proper. Thus at most two optional losses can ever occur in that flank during the pass. On a fresh Purge with no such loss, a component of size $\geq q + 1$ stays intact by the strong padding argument, while one of size $\leq$q disappears. If any of the at-most-two optional losses occurs within a Purge starting with $s \geq q + 3$ sites, all but the two possible losses persist: the rightmost original survivor has original index at least $s - 3 \geq q$, and a new left break caused by a single loss of the old left endpoint cannot propagate farther in q ticks than original index q (if it disappears on reset before propagation starts). In the latter extremal case $s = q + 3$ and loss of *both* original endpoints still leaves rightmost original index $q + 1$, outside that signal. In detail if the left original endpoint is removed at initialization, propagation on the shortened string starts with reset ones; its rightmost after loss of the other endpoint has at least $s - 3 \geq q$ old predecessors in the remaining run. If the left loss occurs later, the zero front can only move less far. If no left loss occurs the zero front reaches at most original index q-1. On a

Purge with $s \le q+2$ and an optional loss, the enduring branch leaves at least q-1 on its final row, hence the total sites removed in that Purge is $\le$3. At most two Purges can be exceptional in this accounting. Allow two additional optional losses on instruction dates for a conservative bound of eight sites lost. Hence each still separate enduring inner face gains at least B-20 inward sites during the chosen pass. The two gains exceed the initial hole bound, since $2(B-20)-(B+30q+40) = 98q-80 > 0$. There must be a join or earlier extinction. After joining, the true internal bonds have grade two, so the temporary cleanup cannot reopen this genuine hole. After distal extinction, a new detached genuine piece would require another pulse or an authentic donor already beyond the vanished piece; neither exists for this isolated source. A side whose only long connection was to *another* exact owner is handled by its own strong backing at that owner boundary and the same alternative. With no time before normalization for the full pass, retain only the proven localized-wound bound and restart ordinary geometry on normalization.

A ballot inside the response. Under conditions (i)–(iv) of the statement a ballot is a no-birth step whose non-native decisions lie in K. By (iv) the count to the first complete clean Purge uses at most four birth opportunities, and every non-native coordinate lies in K, so the bounds $25d+100$ and $3d+12$ hold as above; the last exceptional date is cleared by t_c, so the same pass is selected. By (ii) a panel loses at most one further site per side across the ballot, which the ten-site margin pays together with one initial omission and three withdrawals or Shrinks; by (iii) no old temporary mark crosses it. The remaining steps apply unchanged. The issued-source application verifies (i)–(iv) at a positive vote.

For the ordinary instantiation, there are no true Grow births or withdrawals. Each of at most three real Shrinks removes at most one layer on either exposed face, and an initially improper tip may additionally be omitted; the same ten-site reserve pays for these losses. All other Main instructions here are nonrecruiting. In a trapped internal cavity both inward ordinary bids are legal because its addresses form an interval in $[0, P)$. If one flank disappears, the gap is exterior and no two-sided closure is asserted. No same-rank issuing conclusion is used. □

Lemma 5.8 (The ordinary damage and recovery bounds)**.** *In the ordinary precommit calculation from the lower-rank hypotheses, a home satisfies the erased branch or the surviving backed-interval branch of Lemma 5.5. For a survivor the single interruption costs one localized response of length at most w^6. At clean retired strobes outside it the stage-zero and stage-one end budgets are $3A+3\Delta$ and $6A+3\Delta$, respectively. By $s_c - w^6$, either authentic material is extinct or $C = [12A, P-12A)$ is covered outside the one local wound through pre-ballot work. The wound is confined to K enlarged on each side by $2B+40q+40$. To the first complete clean Purge there are at most $25d+100$ new or reassigned positions and $3d+12$ components; every surviving component of length at least $q-6$ contains 24 consecutive old positions held continuously since before S. An unmarked contiguous run of $q+1$ proper old incumbents inside one ordinary home and disjoint from K survives a full initialized Purge with arbitrary exterior states, and also a Purge already in progress at the first exceptional transition if the run existed through its clean prefix. Such an unmarked run of $q+5$ sites retains its middle $q+1$ through any interval containing at most two real Shrinks. The erased branch can have only its bounded column remnant until an initialized Purge.*

Proof. All times below count child updates, with S a set of consecutive exceptional transitions and K its interval of child coordinates. The common enlargement for shifted-grid queries is applied before using $d = 4w+8$. Assume the earlier-rank interpretation of the native rule, including proper survivors, proper births, and the nominal triple rule at both the old and new row.

1. Counts to the first clean Purge Fix one ordinary home and clock, before its ballot or a change of address convention. Immediately before S its actual matching occupancy is one interval, possibly empty. Let t_E be the end of the first complete initialized Purge starting after S. With $H_0 = 2q+2$ an instruction occurs only once per H_0 ticks. Starting after the last exceptional transition one waits at most $H_0 - 1$ transitions for the first clean initialization and then takes $2q + 1$ transitions through its last killing tick. Including S gives at most $d + 4q + 2$ transitions. Since $d < q$, a sequence of four distinct instruction dates (first and last at distance $3H_0$) cannot fit entirely inside these transitions. Keeping four opportunities is conservative even if both bounding rows are counted.

During this time let k be the number of canonical recruitment opportunities elapsed; for purposes of the proof count a full instruction even when it does not recruit. Apart from K enlarged by k positions on both sides, each of the two remaining half-lines meets at most one interval of the actual phase. At $k = 0$ initially this follows by intersecting the single initial interval with the two rays. Consider one future update. On a nonbirth tick an old interval within a ray can lose only endpoints. Its strict interior has two matching neighbors, hence is proper by the lower nominal triple property, computes its true update, and persists when the native rule holds it. Purge kills only at a right break, Shrink only at a zero face, and Temp retreat only at a zero face. On a birth tick clip each ray by one additional inner position. A child of a site outside the old ray cannot reach the new ray. Descendants of the one old ray interval still constitute an interval, possibly empty. An old interior has a proper positive hold. If a previously exposed tip with an old next-inner neighbor issues an actual outer child, then the old tip, its inner neighbor and the child all have matching nominal outputs (the first two have actual old served states and compute positive native holds). The middle of this next nominal triple is therefore proper and served. A singleton with children on both sides has the same argument; with only one child its descendant set is connected whether or not the old tip persists. A right- or left-end loss with no outer child adds no component. At the old ray boundary, a previously excluded parent's child can occur at most in the one newly clipped position. This proves the induction.

Inside the enlarged band there are at most $d + 2k$ occupied coordinates, giving at most this many components there and at most $d + 2k + 2$ components overall. For $k \le 4$ this is $\le d + 10 \le 3d + 12$. On each of the at most four opportunities a clean birth can only adjoin one of the at most two exposed ends per canonical component at the previous row. At most $8(d+10)$ assignments are therefore clean new births, counted with multiplicity. All remaining forged canonical sites, or lost-and-recovered canonical sites whose recovery is forged, occupy one of d coordinates. At most $9d + 80 \le 25d + 100$ distinct coordinates in total can be newly attached or forged. If a surviving component at t_E has length at least $q - 6$, delete these coordinates. Deleting k'$\le 25d + 100$ coordinates leaves at most k'+1 runs. If every run had length at most 23, the total initial size would be at most k'+23(k'+1) = 24k'+23. The inequality $q - 6 > 24(25d + 100) + 23$ holds for $w \ge 100000$ by division by w (right-hand linear coefficient 2400). Hence some 24 consecutive original coordinates in the component have survived continuously since before S. Each lost original site can only have returned by one of the counted assignments.

A canonical fragment with no outside flank at the start of that clean Purge is erased if its whole strong-link component has $\le$q sites. Along any rightward string of strong links a surviving Cn=1 after q propagation ticks requires q strong predecessors on that interval: unfold the copies of positive left Cn one tick at a time. The finite string lacks such predecessors for every site if its length is at most q. No new site is born in the Purge. Starting from the resulting zero flags, the rightmost remaining site of each component has a weak or broken right link and is removed at each killing tick. Optional lower omissions only shorten the strings. This proves extinction within q killing ticks. A genuinely raw-only independent

phase confined throughout to K expanded by four has width $\le d + 8$<q. A newly joined flank instead belongs to the resulting whole component.

2. Fixed backing in the old genuine interval An unmarked contiguous interval of $q + 1$ proper old incumbents inside one ordinary home, outside the damage, survives a full initialized Purge with arbitrary exterior states. Induct through its propagation: its last site has a length-q chain of strong left predecessors within that interval. At the end of propagation its Cn is one. It cannot die during the killing ticks, while each earlier site has a strong right link to the next and remains. A Purge already partly executed at the first exceptional transition has the same conclusion when the old interval existed through its preceding clean prefix. Preceding clean birth instructions, old proper exact holds and a reset on an already occupied interval do not mark these original sites. On a Temp retirement their own original mark remains zero. On an ordinary Shrink either exposed edge of this fixed interval may be lost, costing at most one from either side on one such update. Hence for a time interval containing at most two real Shrinks, any original piece of $q + 5$ such sites retains at least its middle $q + 1$, including when a Shrink occurs before the first clean Purge. Original strictly inner sites at an old single interval are proper by lower triples; discarding one old endpoint on each side gives the needed originally proper section. This proves survival of deep old backing without using recovery.

3. A cavity opening onto a true end A vanished distal flank leaves an exterior gap of a single surviving interval. There is no requirement that a single donor close distance $B + 30q + 40$ in one B-opportunity pass. Ordinary proposals are in-home, so the adjacent home across a true end is never a legal donor for the remaining in-home vacancies. If the gap is not relatching and no other canonical material exists ahead, the authentic strobe end moves inward by at least one on each eligible Shrink: a provisional trial ending strictly short of the true end has a zero face there and is fully removed by B retreats. Failure to advance on an attempted birth only reduces trial length. This is the open-end episode counted in Lemma 5.9. A positive true seam already fully occupied in both homes has the stronger fixed-backing repair discussed in Lemma 5.6.

4. Relative size of the threshold constants With $B = 128q$ and $\Delta = 1024q$, the separate small-run test $16B + 2d + 24$ is strictly greater than Δ and smaller than 3Δ. A guarded strip beyond a threshold uses the 3Δ comparison in the guarded erasure calculation. The tighter Δ in the survivor calculation bounds only the extra displacement of an authentic restored frontier during the short local response.

5. Guarded erasure with the pulse at an arbitrary point of the timetable Use fully retired ordinary strobes, also between two successive real Shrink calls. In particular one can select a retired strobe at distance at most d_0 from any ordinary fine date. Unless explicitly called an inspection, such a strobe need not be followed by a real Shrink. On a fully retired interval, a missing end at positive depth cannot be advanced permanently toward its true edge unless the edge relatches or another matching component ahead of that end has been encountered. The foremost site of the whole population on either side of a permanent sterile strip never has a component ahead of it on that same side. A trial of $\le$B marked sites into that open gap, stopping short of a true-end exemption and any other matching component, is erased by the B no-birth retreats: at each step its exposed forward site has grade zero and is nonexempt. This proves the assertion even for failed attempts or if an honest lower end is independently omitted.

Let $W = 4B + 4$. A W-site vacant strip at a fully retired strobe, entirely inside the ordinary home, disjoint from K and from the two true endpoint sites, is a permanent barrier. At most B marked sites enter per side on a pass. Its central $W - 2B > 0$ sites remain vacant; every still entering tip is nonexempt and withdraws by at least one per removal until its entrance is empty again. Other phases yield zero grade to its advancing tip unless they count as this matching population. Lower omissions can delete more sites but cannot create an unmarked new site inside. Thus its middle stays empty at every fine date and the entire band is empty again at the next strobe, including during the pulse elsewhere. Across a transition from one ordinary stage to the next the same true homes, reset and trial restrictions apply.

Here is a quantitative interruption charge that does not require a macroscopic old interval. At the last retired pre-S strobe let f denote the nearest position of matching material on one side of a permanent sterile band, ordered by distance toward that band. If no old material exists on that side, only a raw-only family of width $\le d + 8$ is possible up to the first complete clean Purge, and vanishes there. Otherwise at the end of that Purge genuine descendants can be farther toward the band than f by at most $B + d + 16$. From the old strobe up to S there is at most one partial ordinary pass, giving $\le$B steps beyond the former largest coordinate; if the chosen strobe is just before the previous ordinary handle and a real Shrink intervenes it can only subtract sites. From S to Purge end there are at most four additional birth opportunities on this clock. Descendants with no raw assignment in their ancestry therefore lie at most $B + 4$ beyond f. Those with a raw ancestor and now outside that previous half-line lie in K enlarged by four. If this last band has no contact with any genuine old-descendant site by the time of the complete Purge, its entire width $\le d + 8$ dies. For a contact its far side can be at most $d + 8$ farther than $f + B + 4$, plus one for adjacency. A pre-existing old group at a different location behind f does not increase this maximum. This proves the stated $B + d + 16$ bound, whether or not the pulse creates numerous components in K. Through the first following retired strobe pay at most another B for the remainder of a pass. Thereafter the foremost frontier can only withdraw again at the clean strobes. Between the old and the new strobes at most three real Shrinks have been missed: the elapsed time is $\le 2d_0 + d + 4q + 4 < 3d_0$, and true Shrinks are $d_1 = q\ d_0$ apart. Use the symmetric argument for a leftward order; this is an order calculation on front positions, while the actual small-component Purge proof still kills from the right.

Suppose the left endgap at an inspection (or any retired strobe in the eligible stage) exceeds $\theta + 3\Delta$, where $\theta = 3A$ or $6A$. In the vacant part strictly between θ and the actual old population choose a W-site strip not meeting K or a true endpoint. There is room for two disjoint W strips separated by $d + 2$ and both avoiding the true endpoints since $3\Delta > 2W + d + 4$. K cannot meet both. No old material lies to its left. There can only be an independent, first-Purge-extinct colony there after S. Every true Shrink with threshold θ removes the leftmost remaining site to its right, as its zero left face is at depth $> \theta$. The sterile band prevents any probe from permanently latching at the left true end or encountering another matching site ahead. On clean strobes the ordinal position of the nearest surviving authentic front increases by ≥ 1 per eligible real call; disappearance of a component exposes a farther front and only speeds this up. Counting from the start of the θ stage, at most $P + 2B + d + 19$ eligible inspections suffice even if S falls in its middle, by the preceding interruption charge $2B + d + 16$ and three potentially missed calls. This is $< 1.1P$; there are $2P$ such calls in a whole stage. When the large gap was already seen in the preceding lower-threshold stage, the same band remains empty and all authentic sites on its populated side still have depth $> \theta$ when the θ stage begins. The independent late raw-only case does not produce a genuine late comparison. The reflected end uses the depth $(P - 1) - x$ in the eligibility test and the rightmost remaining authentic front facing the strip. If the sole survivor is exactly at the opposite true endpoint, the face toward the strip is still internal

and eligible.

For an initially arbitrary normalized home at retired strobes, each interior gap surviving n clean stage-0 inspections contains at least n old vacant coordinates in the following precise size sense. At an inspection a still-internal gap has unjoined, retired banks: no clean birth creates a new internal gap, so trace one preexisting internal gap inside its predecessor whenever two gaps merged. At each true Shrink its exposed bank loses its inside tip (or the entire bank). A recurrence on the ancestral gap's length gives the next internal gap at least one longer than that chosen gap. If the gap becomes exterior, it can never become an internal gap again without an independent newly born component, unavailable on the clean ordinary evolution. Therefore after $20A$ clean real inspections the home is one interval or has a gap of length $> 20A$. Choose a W-site sterile band within that gap avoiding true ends and K as above. On each side the matching front facing this strip erodes by ≥ 1 per remaining Shrink(0) call, using the interruption charge above on both sides concurrently. More than $3P$ stage-0 calls remain, so both sides are gone in at most $P + 2B + d + 19$ of those calls, except short immediate raw-only debris. If S intersects the initial $20A$ inspections, restart once after S and a clean ordinary reset; $50A$ d_1 covers both attempts and the boundary spacings for $w \geq 100000$. If the sole surviving interval on that early clean row misses a true end by more than $16B + 2d + 24$, there is again room for two disjoint W candidates separated by $d + 2$ between that end and its front; a sterile strip forces the same stage-0 extinction. Thus on a branch capable of authentic late occupancy the simultaneous early row is a genuinely long interval and both θ thresholds lie strictly in its proper interior. A site outside K which was a proper incumbent stays proper on every subsequent live native hold; an honest re-entry after absence is a proper birth by the previous-rank proper successor clause.

In the later stage-0 interval, $P + 2B + d + 19$ calls in the following whole stage 1 dispose of any authentic population whose endgap ever exceeded $3A + 3\Delta$ at a genuine retired strobe. Likewise a stage-1 depth above $6A + 3\Delta$ is excluded by the whole following stage 2. This establishes the guarded-threshold bounds on every surviving clean retired strobe without presuming a finished local refill. In particular prior to an unspent S in stage 0 or 1, either one of these erasure tests applies or the two true ends enclose an old interval of length at least $P - 12A - 6\Delta - 2$. A virgin late pulse in stage 2 or computation follows a wholly clean stage 1; an initially active left stage-1 gap $p > 3A$ cannot latch because $p > B$ and grows each Shrink until extinction, and an inactive $p \leq 3A$ cannot first lose a proper occupied $3A$ threshold spontaneously. In detail, while the interval is long, Purge misses the strongly backed last site, the endpoint at depth θ has no eligible outward Shrink, and a genuine strobe end at positive depth is unmarked (a nonjoining B-long outward trial has completed all its removals). An uncharged omission there would require improperness, excluded at an initially proper or honestly reborn threshold site. If either end was already active, once the old interval becomes short its continuing exposed active face only accelerates its exhaustion; no clean outer trial from its interior can cover an end gap already exceeding $3A > B$. Therefore in this clean-before-S case either the home is already extinct or its stage-1 exit has both end depths at most $3A$. The same clean comparison applies during stage 2 up to the first later pulse.

6. Local response on a survivor and final-stage timing Apply Lemma 5.7 to the old macroscopic ordinary I just obtained; its length is at least $P - 12A - 6\Delta - 2 > 2q + d + 44$, its strict interiors are proper by lower triples, and its evolution before S is clean. The issued-source application is given later in Lemma 5.12; that is a pointer, not an input to the ordinary calculation. There are no true Grow births or withdrawals in these ordinary stages. At most three actual Shrink instructions occur in the response. Each can remove one layer on either exposed face, and an initially improper tip can additionally be omitted. The ten-site reserve on each face pays for these losses; all other Main instructions are free of structural

attack or withdrawal. Both in-home inward proposals in a trapped internal cavity are legal, since the intervening ordinary addresses form an interval in $[0, P)$. No ballot occurs on this ordinary response if it is used to establish geometry before computation. For a pulse near the computation ballot, the ordinary assertions here stop at its pre-ballot row with the local wound still charged. The two event horizons and partial pass in that calculation yield at most $B + 30q + 40$ sites of persistent internal cavity, or disappearance of the distal piece. The successful-pass estimate uses $B-20$ instead of $B-3$; $2(B-20)-(B+30q+40) = 98q-80 > 0$. If one side is lost and the hole is exterior, use the earlier true-end cavity calculation. Deep surviving old strips have no interior loss, and the discrepancy before restoration is confined to K enlarged by $2B + 40q + 40 < .2A$. The full response uses $< 2d_0 + 10(2q + 2) < w^6$ ticks and at most three true Shrinks. To be precise, the portion of a fresh pass used after its uniform clearance may skip up to ten initial growth instructions, already included in $B - 20$. This local argument does not import any same-rank strong vote or any Grow-specific conclusion into an ordinary precommit calculation. The guarded-erasure and survivor steps use the held-wall implication of Lemma 5.6 and the supplied-interval lemma only when their wall and old-interval premises have been established.

At either original strobe end of I, a surviving old proper strip begins within $q + d + 40$ sites of that end throughout the response on the macroscopic branch. Indeed choose two candidate strips of length $q + 12$, separated by $d + 1$, strictly inside the old interval beginning one site from that end. K can touch at most one. Trim four old positions at each end of the surviving choice for potential Shrinks and an unprotected extremity, leaving more than $q + 1$ old sites. The nearest preserved point has depth at most $q + d + 18$ beyond the previous old endpoint. Conversely, by the first-Purge contact calculation in step 5, a *persistent* edge cannot shift outward beyond the old retired edge by $> B + d + 16$ at first cleanup. Up to the completion of the chosen pass there are at most two full or partial additional B-growth passes and up to three Main steps; all ordinary Main steps here are nonrecruiting. Including one more B safety allowance, the absolute displacement of a restored strobe end beyond the extrema of the original I is $< 4B + d + 40 < \Delta$, with the inward loss paid by the preserved strip. If a true-edge latch occurs during this interval the newly latched endpoint must also lie within this paid distance, since each recruited layer is counted. These assertions are for the authentic strobe frontier, not for a distant unjoined raw-only satellite.

For the final threshold test choose one true end, put $h = 3A$, h'$= 6A$ in its directed depth coordinate, and separately test whether K lies within 2Δ of h'. Since h'$-h = 3A > 5\Delta$, this cannot also put K within 2Δ of h. Work on the potentially surviving macroscopic branch above; if an active facing population instead disappears only the immediate small raw alternative remains.

1. When S starts before the beginning of stage 2 and K is farther than 2Δ from h', the site at depth h' was proper on the early long row. Any later authentic re-entry there was honest and proper, because the site is outside K. If both ends of the restored interval have depth $\le$h' at stage-2 entry they cannot first cross that bound: on a long interval at an inactive threshold its genuine outward Shrink is exempt, interior bonds prevent other removal, no improper omission is possible there and backing prevents purge. The two-end first-crossing argument uses length $\ge P - 2h$' until the first such crossing. If an end started at $p > h$', it faces a gap of width $> B$; unsuccessful clean trials fully retire, each stage-2 Shrink removes that foremost member, and by the first $P + 2$ genuine calls the old population is gone. The other end's behavior cannot slow this one. If S itself straddles the stage-1/stage-2 transition or began fewer than $2d_0 + 10H_0$ ticks before stage-2 entry and restoration continues into it, delete the $\le$3 initial interrupted calls and start this same clean test at restoration. Original properness at h' is preserved through the job or supplied anew by honest re-entry outside K. This leaves far more

than $P + 2$ calls.

2. When K lies within 2Δ of h', it is farther than 2Δ from h. On a clean inactive stage-1 trajectory $p \le h$ the proper h site cannot fail uncharged. If S comes while this boundary is still inactive, h lies in an old protected strip directed inward of length $q + 12$ beyond h disjoint from K. Through the local response the left face at h remains exempt in stage 1, its inward $q + 1$ backing persists across the $\le$3 Shrinks by the reserves, and if h was a strict interior it either remains so or becomes that exempt exposed endpoint. The properness of its old incumbent follows at every live step since h is outside K. When the interval is restored, $p \le h$ still. If $p > h$ already at stage-1 entry, clean calls increase p until S. The interruption cannot relatch it across h. For $p \ge h + \Delta$ at the last pre-S strobe, the permanent outward gain proved above is strictly $< \Delta$. For $h < p < h + \Delta$, K is $> \Delta$ from the frontier and the $\le$B partial probe from it. A raw-only patch here cannot contact that frontier: a birth history of $\le$four instructions before a complete Purge stays in K expanded by four, and no persistent independent short patch is possible. Nor can a change deeper than Δ force a new inwardly supported outward latch from this frontier: there are true old proper sites of length $q + 12$ immediately inward of it, and the exposed depth $> h > B$ forbids a completed true-end latch; a genuine vacancy ahead has no other old matching component. It therefore remains $> h$. The same argument applies when S was already spent in stage 0: if h is occupied at stage-1 entry it is proper, and if the end is beyond h it dies through the clean stage. Accordingly an active stage-1 end facing a wide gap throughout pays at most Δ extra retreats plus three missed calls and is exhausted within $P + \Delta + 4$ of the $2P$ calls. On the surviving alternative this side has depth $\le$h at stage-2 entry. If damage is still being repaired over the boundary, the surviving old strip beginning within $q + d + 40$ of its former end gives depth $< h + \Delta$, far short of h'. Thereafter at most the one uncharged loss of a previously improper exposed tip can occur before any fresh authentic endpoint birth, and that loss while backed cannot open a depth-h' breach. Thus a possibly forged h' site is never asked to retain an already marginal h' edge arbitrarily late in stage 2.
3. If S has not started before stage 2, the wholly clean stage-1 exit has $p \le h$ on both ends, or has gone extinct. Until S any exposed proper h' site cannot be endangered from an edge of depth at most h plus one optional endpoint loss. The macroscopic old strips bounding inward setback throughout S put the new strobe depths below $h + \Delta + 2 < h'$. Their original interior extent is again much larger than q, and the at most $d + 8$ subsequent improper-tip omissions cannot move an edge from this depth to h'. Erasure of a remote deep piece has already been localized and repaired by the internal-cavity lemma. An S in computation instead encounters the same long old interval under no true Shrink and the same one local wound.

Both ends undergo the same numerical test, sharing a single K rather than independent exceptional columns. In ruling out a spontaneous first crossing while applying the two tests, stop at the first crossing at either end, at which all protected opposite-side buffers are still macroscopic. If an end begins a stage active (>threshold), its exhaustion argument depends only on its own positions and remains valid when very little material remains. By the time of $s_c - w^6$ either authentic material is extinct (temporary raw occupancy only), or $C = [12A, P - 12A)$ is covered outside the one local wound through pre-ballot work. The long stage-2 erasure branches terminate by $P + \Delta + 10$ calls, leaving an exponential margin before the end of its $2P$ calls. The two stage-0 and stage-1 bounds of step 5 and the clean interval persistence used for the nonfull-date count now apply to this latter alternative. The true-end one-sided case and the distinction between a grade-one bridge and an independent Purge backing are unchanged. □

A home is *nonfull* at a fine date if at least one of its P positions lacks an actual matching

served child.

Lemma 5.9 (Number of nonfull dates and a common full seed)**.** *Work after the early stage-0 relaxation. Assume the survivor assertion of Lemma 5.8: apart from one response interval of at most w^6 fine dates, this is one genuinely contiguous ordinary interval at retired strobes, with each endgap $p \le L = 3A + 3\Delta$ at each clean stage-0 Shrink strobe. Assume also that on a clean partial pass lasting through a strobe, a trial which fails to latch its true end returns to the previous exposed frontier, and that through any fully covered true end a $q + 1$ backing of previously proper incumbents is maintained on this branch. Then after the relaxation and throughout stage 0 the number of nonfull fine dates is at most $100A\ d_1$. In particular $N = 200A\ d_1$ is a uniform search allowance. With the same hypotheses for up to six neighboring homes and the common forbidden dates below, a full common seed of $5d_0$ consecutive ticks occurs.*

Proof. Cut the calendar at the response interval, leaving at most two clean time intervals. Consider one true end in one such interval. After a genuine birth authenticating the true-end site the newborn is proper; it remains served while its native update holds. Its outward zero face is exempt from Temp and Shrink, and its in-home link is positive. The $q + 1$ same-home true-end backing prevents an erasure from the tail of a Purge, even though the seam's wrap link itself is only grade one, by the last-site backing argument of Lemma 5.8. Consequently if the true end is initially present and improper, that initial site may be lost at most once, and after any subsequent relatching birth there is no further clean loss in this fixed clean interval. An initially absent endpoint counts as one initial missing episode. Allow two missing episodes per end per clean time interval, including a partially open first or last episode.

Until relatching, any incomplete attack into the remaining true-end gap has only B marked sites along its nonjoining trial; each removal tick retires its newly exposed eligible tip, hence at a clean fully retired strobe this trial cannot decrease the old depth. At each Shrink(0) with positive depth the end with its outward zero face is removed, increasing the depth by at least one before the following strobe, unless a true-end latch has since occurred. Thus successive fully internal strobe depths of a missing episode are distinct increasing positive integers $\le$L, so there are at most L such strobes. With d_1 between strobes an episode including the two arbitrary partial spacings takes $\le (L + 2)d_1$ dates. Fine rows temporarily covering the true end inside such an episode only reduce its nonfull count. Counting both ends, both clean intervals and two episodes gives $\le 8(L + 2)d_1 + w^6 + 4d_1$, padding both exterior sides and both response boundaries by a strobe spacing. At $w \ge 100000$, $3\Delta/A = 3072/w < 0.031$; $d_1 = q(2A + 4)(2q + 2) > 4w^7$ implies $w^6 < d_1$. Hence $8(L + 2) + 5 < 100A$. A hidden interior gap would invalidate this argument; the genuine interval premise and the response charge are essential.

For up to six neighboring positive branches, let n_i denote their nonfull dates. Each $|n_i| \le N$. For each common raw temporal interval requiring exclusion, enlarge it into the full forbidden date set, of total cardinality at most $4w^6 + 20$ since the single immediate response has length at most w^6, any associated delayed ballot response is confined to the same wound and has length at most w^6, and adding a w^6 preparation halo on each temporal side, with 20 rounding dates in total, covers all queries. Within stage zero itself there is no new ballot. Define $b = 6N + 4w^6 + 20$. After the early cutoff take $b + 1$ adjacent disjoint trials, each of length $5d_0 + 1$ integer dates. Every failed trial contains at least one distinct forbidden date (use its first such date). The union of six nonfull sets and the slabs has size $\le$b. Therefore at least one whole trial is full in every needed home. In particular, this costs $(b + 1)(5d_0 + 1)$=poly(w) dates, while the remaining Shrink(0) timetable has exponentially many calls and $d_1 \ge 1$. A full trial of this length includes a Temp reset and a complete Purge whose initialization follows that reset, because resets recur every d_0 ticks, and a Purge takes

$2q + 1 < d_0$ ticks. For adjacent full interiors, the lower triples which straddle a positive seam put both seam endpoints in proper status during the seed interval. Coverage after a later raw interruption requires the paired repair of Lemma 5.6.

The common full trial supplies the paired input corridors and the six-home service calculation. The survivor and backing premises were established by the ordinary relaxation and damage bounds.

Discharge on the surviving stage-0 branch By Lemma 5.5 and steps 5–6 of Lemma 5.8, away from the one response interval there is a genuinely single ordinary interval at each retired strobe and each of its two directed end depths is at most $3A + 3\Delta$. Its length is therefore at least $P - 6A - 6\Delta > 2q + 30$. When it reaches a true end by growth the new end has an actual proper birth. Its near in-home $q + 1$ sites form the beginning or end of this interval; the true outward Temp/Shrink exemption, positive inward link, old proper holds and rightmost q-Purge backing imply that an already achieved latch cannot be lost through a later clean retreat. For an initially full but not-yet-proper endpoint, allow its first free loss; strict old interior sites were already proper by triples and later legitimate births are proper. A previously nonjoining trial on an exterior gap that never reaches the true edge has no matching population ahead and regresses its still-marked facing tip at all B removals. The current single interval could not meet another matching component across that exterior gap. Thus precisely the clean-episode premises of the counting lemma are discharged by the preceding ordinary geometry. If the end was attained for the first time on the Bth attempt, the facing true-site exemption already applies at the very first retreat, and its positive inward link provides the other face. With a full seeded pair any ensuing seam interruption is treated by Lemma 5.6. □

5.3 Transactions and classification

The weak-description then strong-issue order follows Gács [10, Lemmas 14.1–14.2].

Lemma 5.10 (Weak descriptions and the strong issuing cut)**.** *Use the ordinary precommit alternatives and the lower-rank transition and regularity clauses for the owner and its three suppliers. There is a weak positive description by one common live-letter string on the central fifths, or a weak dead cut with actual support in the exceptional column. Three prior weak descriptions give independent parsing of the served suppliers and the strong cut of Proposition 5.3*(iii)*. The latter includes the noncomputing branch and a possibly halted true literal.*

Proof. In each unaffected evaluation the full central three working fifths are occupied with scratch reset. The three LOAD sites and both neighbors on their actual shifts are lower-rank live matching triples. Their exact native step loads the true assigned operand at that position even from arbitrary pre-existing latches and program tape. In every such repetition the same three fixed geometric positions and parameter position are used. Conditional on the specified frozen environment the target F is therefore a single fixed owner-specific map, including on forged legal mutable states. Section 6.1 identifies the actual operands and proves localization.

Geometric inputs. A weak positive description means one common live-letter string on the central fifths off the common physical column, with the narrow support test specified below; a weak dead description means actual matching support confined to that column. The inputs are the late core of Lemma 5.8, the common full seed of Lemma 5.9, and the paired corridor of Lemma 5.6. Use the lower exact-or-dead transition, true lower birth/hold and proper-successor clauses. For the current centered transaction and its three suppliers assume the ordinary dichotomy, the late-date core C, and the continuous paired corridors of

Lemma 5.8, Lemma 5.9 and Lemma 5.6, wherever the full time interval of an operation misses the raw slab enlarged by w^6. The empty ordinary branch has no matching material beyond the immediate $5q$ allowance. A disturbance during ordinary work changes an old survivor's interior only in K enlarged by $2B + 40q + 40$ (or the wider $.2A$ used here); outside it the survivor endpoint ledger keeps C protected after the common seed. Work in the one rectangle K×S in this entire buffered domain; the possible delayed wound at a subsequent vote is in the very same K. A complete good repetition spans its reset, input gather, immutable loads, computation, result writing and final output writes. The separation P idle Main instructions gives pairwise disjoint $2w^6$-padded transaction intervals and permits the stipulated margins before the first and after the last prior cut. No claim here assumes the current nominal output.

Old carrier invariant. Number fifths 0,...,4; the three inputs of each supplier are in fifths 1,2,3. In a previous weak positive description there is one common live-letter string z, possibly malformed, on all its central fifths off K, and C_* is served and matched off one common interval of $\leq 2q$ containing its possible missing positions and any wrong tested letters. The narrow actual issuing cut localizes any such loss at K and its immediately adjacent short piece, or in the $4q$ reserve outside C_*. In the absence of a raw column there is no such defect. In particular the error sources for X in the central fifths have the one common physical column; this is an explicit part of the narrow weak cut proved below.

Any $q + 1$ consecutive old positions deep in C_* outside those exceptions and K, with one adjoining old match at both ends, are initially proper by lower triples. Their original scratch, including Temp, is reset at the preceding healthy ballot and their original addresses are ordinary. Through an entire Grow no true Grow withdrawal reaches those addresses, recruits ignore incumbents, and no Shrink is active. A complete Purge preserves the whole old strip: the rightmost Cn remains one after its q internal strong predecessors, and all earlier sites retain strong right bonds in every killing tick. Temp stays zero on old incumbents, even when nearby births are marked. The fixed-strip proof and lower proper-successor clause apply inductively through all handles to every old strip and are independent of the exterior. This preserves actual held X through the long Grow. Normalization keeps these ordinary addresses unchanged.

For an explicit later spatial margin set $N = 200A\ d_1$, $D_{\mathrm{tr}} = 50A\ d_1 + (6N + 4w^6 + 21)(5d_0 + 1) + w^6 + 2d_0 + 20$. This exceeds the upper bound from Grow-to-ordinary normalization to the selected common seed. For fixed w all these quantities are polynomial expressions in w; in particular choose w so that $2(D_{\mathrm{tr}} + 10A + 10q) + d < P_0/2$ and $12A + 4q + D_{\mathrm{tr}} + 10A + 10q < P_0/2$. In ordinary times up to this seed, any middle-fifth cell whose D_{tr}-neighborhood is entirely old correct and excludes K retains X: in each next step the strict interior of the previously held segment has two strong neighbors and is proper and held by the exact native transition. No clean interior death occurs in this interval, including in a partially completed Purge, because its old right bond is strong. Trimming one site on each side per step proves the claim. In particular all three central fifths except one physical interval K expanded by $D_{\mathrm{tr}} + 10A + 10q$ carry the original z. Its width is $< P_0$. The outer ends of C_* have more than D_{tr} margin to the middle fifths. After the seed and through good reads every such old unhit position remains present by the earlier surviving-home ledger and the localized single-wound repair; a proper living incumbent does not rewrite X on a nonballot step. At any offset u, positions $sP_0 + u$ for $s = 1, 2, 3$ are P_0 apart, so at least two *of this supplier's coordinates at that offset* continuously carry z. A different u may choose another pair. When the preceding weak decision was D, matching occupancy at the cut outside K was empty. A subsequent matching X at such a site comes from a native birth, writing the playable neutral carrier at its own residue; thereafter it is again held or replaced by a true birth with the same neutral letter. A foreign surviving incumbent cannot switch into canonical matching status across the off-Grow

normalization without the correct original clock and residue; at the preceding cut all off-K survivors had normal addresses. This proves the neutral or physically absent alternative.

Paths and validity. For receiver home h and supplier offset j in {-1,0,1}, let $h_0 = hP$ and source coordinate $a = (h + j)P + sP_0 + u$ and destination $b_{\mathrm{dst}} = h_0 + (2 + j)P_0 + u$. During supplier j/copy s transport each hop at step k has prescribed coordinate $a + k\,\mathrm{sign}(b_{\mathrm{dst}} - a)$, $0 \le k \le |b_{\mathrm{dst}} - a|$, on its assigned lane. The displacement is an integer multiple of P_0 independent of u. At injection only a site with the programmed source residue places its X there; every uninjected lane writes message-empty. At zero displacement the destination takes the newly assigned local value. At a positive displacement count (in either orientation), the old upstream signal from the unique predecessor enters the next slot if the receiver survives that hop and their grade is positive. A missing link gives message-empty, and a newborn's mail fields start as message-empty. No stale destination value is read because receive unconditionally overwrites the indicated bank at that offset. Induction on k therefore gives X(a) on an unbroken path or message-empty on an interrupted one. Any injected garbage on a physically damaged hop belongs to the one affected repetition and is not included in either good read. Across computing homes every hop is within their joined corridor (including a grade-one wrap, which transmits); all paths between homologous data slots remain inside the core-to-core corridor. Every required write and hold is the native update of the actual supplied slot. For $g = 0$ the entire source fifth at a good injection is absent, so all these receives get message-empty or their birth/reset default. The centered evaluation itself is full on the required destination fifths, validator and gathering tape.

For $g = 1$ and a prior weak positive, the at-least-two surviving copies per offset yield z exactly at every offset after the local vote. If z were malformed, any candidate live decode outside its own exceptional interval would be contradicted at some offset: of three homologues one avoids both that interval and K. So malformed z gives $x = D$, exactly the fresh parser output; a legal live z instead yields its serialized x. For a weak dead input, the proper parse of a neutral packet or absence is D. For $g = 0$ the supplier value y is D irrespective of preceding formatting. This proves that each supplier independently parses to its actual y in a good repetition, assuming its preceding weak description. Two good repetitions therefore evaluate exactly the same total native $F(y-, y, y+)$ on the fixed physical loads. The computation work tape and mailbox overwrites are freshly initialized in each repetition. Within output readout apply the same path induction with source $hP + 2P_0 + u$ and each destination $hP + sP_0 + u$ for $s = 0,...,4$. A true dead result cannot insert a live letter on a broken path. Delivery in the middle three fifths is exact at all good writes.

No backward assumption of a strong cut. First obtain a weak description when the center has no direct exceptional update in its late ordinary computation, verification or ballot region (include a $2w^6$ temporal margin before its first read and the entire actual home through its ballot). Incoming data from neighbors may change between all three evaluations. Each fresh full center still outputs one legal live packet with every letter live or all message-empty on its central fifth, so the central three fifths see the same generated letter at each homologue in each repetition. The weak liveness decision is their common majority of the three liveness bits; if live, the vote gives the same string z at every homologue. No validity is claimed for z. A repetition with a dead result can put only dead banks in any unhit shifted path; a direct event in another home cannot inject into the all-in-home output shifts. A direct event confined to an adjacent home has broad damage radius $.2A$ and cannot touch C of this center from outside. For a center whose direct late ordinary work does meet J, the preceding periods of all three suppliers ended a full Grow/Shrink duration before that late work. Thus their late ordinary weak descriptions are all in the preceding clean-center case. Apply the two-good-repetition calculation just proved to this affected center, yielding in particular a weak description without assuming one for any arbitrary earlier predecessor. If S straddles

the bottom of the current period or occurs in its early Grow, the present late-ordinary center is clean; if S straddles its final ballot, it is in the other case and the three earlier supplier work computations are clean. The one separated delayed footprint at a vote uses the same physical column, and its earliest read margins exclude it. Having proved weak descriptions for all internally queried prior periods by this one backward step, apply the good-read argument to the actual current period for its strong conclusion.

Retaining the central votes. In ordinary computation no true Shrink is active. On a late raw pulse take the original interval $[P_0 - q - 10, 4P_0 + q + 10)$ of proper held center slots just before the pulse; on an early pulse instead take the good complete writes after its recovery. Split the original interval on K. Any component of its complement containing a middle data coordinate has $\geq q + 3$ positions, or its short boundary piece lies entirely in the $q + 10$ margin. Such an old component has proper middle triples and had persisted through the initialization of a partial Purge. It survives all later Purges by a contained $q + 1$ strip through the tested position (using the margins for initially exposed sites); unmarked incumbents hold their banks. Thus outside K an earlier good central write persists or a later good write unconditionally restores that bank. In readout step 4 each repetition writes the result to all five fifths, not just the three middle ones. Its transaction interval is padded by $2w^6$, while the single local response lasts less than w^6; hence at most one repetition can meet the enlarged raw slab. In the positive case each of the other two repetitions writes a live liveness bit at every site of C outside K and the refill band. That band is K enlarged by $2B + 40q + 40 < .2A$ on each side, and has width at most $.4A + d + 2$ with rounding. Each such write travels an unbroken in-home path inside C, by the path induction above: its transaction misses the raw slab enlarged by w^6, and at such times the geometric inputs give the late core C fully covered. At every such site an unmarked incumbent retains its bank by a Purge-backed $q + 1$ strip; recruitment cannot replace it. Thus two of its three liveness bits remain intact and live at verifier initialization, modulo the already consumed edge reserves. For the clean-center weak vote all three central writes give the specified weak vote, with the same bound on untouched sites.

Liveness. For a dead majority all initially active sites are in K, because both reliable dead repetitions can only deliver a dead or interrupted dead output, even to an outer site. Raw bank forgery/reset is charged in K. Up to the first purported activation outside it, an $M = 128d$ step sampling epoch has at most d original positive injection coordinates. Any forged directed transmission follows a line $x - t$=constant in the positive direction, $x + t$=constant in the negative direction; each constant gives at most one arrival at a given receiver in that epoch. K×S supplies at most $2d - 1$ values per direction. The resulting $\leq 3d - 1$ count is below the required $> 4d$ threshold. A birth resets lanes, counters and Active. This rules out a spurious positive outside K at any later time.

For a live majority choose a clean no-birth verification interval. The genuine surviving home there decreases as an interval containing C. Any eventual survivor lies within $20A + 2M$ of an intact run of M initially certified positions, excluding the one $.4A + d + 2$ historic band and the $12A$ margin. Along the interval joining this run to the target, all surviving positions and their connecting lanes remain present throughout. One complete epoch certifies the next $M - 4d - 1$ positions: at the furthest new site the incoming distances to at least $4d + 1$ old positive sources are between 1 and M, so those distinct shifted impulses all arrive. Repeating epochs costs at most $(ceil((20A + 2M)/(M - 4d - 1)) + 2)(M + 1)$ ticks, bounded by $100A + 4(M + 1)$ since $(M + 1)/(M - 4d - 1) < 2$. Choose the second plateau if J has expired more than w^6 beforehand. Otherwise the first is clean by the $> 3w^6$ separation, and covers every survivor of that earlier plateau.

After this first certification and before a later J, at each complete Purge all eventual survivors are again certified. In each interval between Purges at most one new clean

recruitment instruction adds two previously uncertified sites to the single clean component. A component beginning the next initialized Purge of length ≤q dies. For every eventual survivor x of a longer component, take its subinterval on the row $2(M+1)+2$ steps into the Purge. During these initial propagation steps there is no birth and only endpoint omissions, at most two losses per step. Length is at least $q+1-4(M+1)-4 > 2M+1$. It therefore contains a solid M-site one-sided interval adjacent to x, held since the start. This time span contains a full sampling epoch, with arbitrary epoch phase. At least $M-2 > 4d$ donors were previously positive. The receiver is certified by the end of the epoch unless already certified. The same proof applies on the first clean Purge beginning just before a later J with its conclusion restricted to the prior clean times as necessary; just before J at most two sites from the last intervening birth need remain uncertified.

At the first complete Purge wholly following J add the at-most-$25d+100$ raw/new coordinates proved in Lemma 5.8. At most $25d+102$ of the donors can then be uncertified; $M-(25d+102) > 4d$. The identical surviving-M-run argument certifies every eventual survivor of this Purge. In each later clean Purge there are at most $2(d+10)$ newly attached positions over the whole home, because the first-clean-Purge component count was $\le d+10$ (step 1 of Lemma 5.8; its statement records the weaker $3d+12$) and clean updates can only merge components. This bound covers several small components merging into one, and again leaves more than $4d$ donors in the one-sided M-run. Finally, if there is no complete clean Purge after J before the ballot, J began at most $d+4q+6$ ticks before it. The final plateau reaches the ballot and has duration $\ge 100A+4(M+1)$, so it already included a complete clean Purge under the preceding certification invariant before this J; throughout the remaining time there are no births. Only K can be reset or forged. At the ballot every surviving site outside K with a live target thus acts on true Active; a missing payload bank at the fringe gives the legal live filler. Central sites off K retain their voted letters. A late raw damage with earlier, widely separated verification interval cannot quietly reuse an uncertified site after the last Purge.

Support. In a positive result split the old C on K; if a side has $\le q+2$ sites its union with K has length $\le q+2+\mathrm{d}<2q$ and lies at the C boundary. Longer old sides have long proper held strips. The $4q$ additional trim to C_* pays the two unprotected outer boundary sites; the same interval around K covers the support test and all potentially incorrect central letters. Absent a recent pulse there is no such actual loss. For the no-computation branch a surviving same-ordinary-home raw-only colony at the ballot has its origin and its residual spatial footprint to check separately: Lemma 5.8 gives the preballot width and dates, and the last no-birth period rules out an off-K newly recruited survivor. In the strong computed-D branch all matching lower occupancy outside K is absent by the negative verifier proof. At an empty K the entire cut is physically empty for this home. A remnant in K can add at most four layers per side before its first complete clean postcut Purge, to width $\le d+8<\mathrm{q}$; different source labels cannot lend a strong link at their contact, so the zero/right-break calculation eliminates it. All ranks use the actual served data, not an assertion that a supplied symbol is immutable after its own $g=0$.

The late-event verification uses the final no-birth interval reaching the ballot.

Positive ballot after a completed restoration The liveness proof restores **all actual survivors**, including a previously hit site in K, at the end of a complete clean post-event Purge once a clean earlier verifier plateau has certified the home. If the event instead precedes that bootstrap, the later clean plateau ending at the ballot certifies all its eventual survivors directly. Between two complete Purges the subsequent restoration argument again applies to every survivor; the final plateau has no births. Thus, when there is at least one clean restoring Purge wholly after S before the vote (or the clean final plateau itself supplies the

bootstrap), every incumbent present in the preballot row takes the native positive decision. At a clean ballot a strict interior incumbent was proper by the previous-rank matching triple. Its positive successor is proper and served by the rank-$(j-1)$ proper-successor and service clauses of Proposition 5.3(ii). Only already exposed improper endpoints can therefore be omitted. In particular such a clean positive ballot cannot split an existing canonical interval anew inside K. This statement is about liveness, without asserting correct restored payloads at a previously forged site.

If the positive ballot can make a new internal hole in K and the first clean plateau did not reach this vote, no complete clean restoring Purge had time to fit after the actual event. The last event thus ended fewer than $(2q+2)+(2q+1)+2 = 4q+5$ ticks before the vote (two endpoint-rounding ticks allowed); its start is within $d+4q+6$. This whole interval is inside the final no-birth plateau. It permits the spatially common original K hit and a ballot loss in K to be handled as one short two-date disturbance without a new canonical recruitment between them. An intervening gap already opened by S can of course persist; the statement excludes only a further internal cut after old liveness had been restored. On a dead decision a forged Active inside K can remain sticky until the ballot and the stronger conclusion is unavailable.

Mail between nonconsecutive transport instructions Successive logical hop instructions can be d_0 fine ticks apart. Following a nondead received value backward across the intervening handles is valid only at a coordinate where the same canonical incumbent survived continuously and its designated lane was in hold mode on all those handles. A birth, including a reappearance in a formerly absent slot, initializes its mail and stored output banks to message-empty; the receiver's current native injection/overwrite clears any further stale value. A physically distinct predecessor on another geographical grid cannot supply this ordinary positive link. Within the same grid, an unhit continuation advances the clock by one modulo T and preserves its address residue even at normalization. Reading the grade at each preceding hop consequently forces the same programmed injection schedule all the way back. During actual input transfer and readout all unhit sites have ordinary normalized addresses; for a fixed clock, residue and location the owner is unique. A wrong clock cannot become synchronized merely by wraparound because every genuine continuation advances once. Therefore a prior message on a different lane, repetition, source phase or former extended label cannot furnish a nondead word after the required path has broken. When a good path is intact, induction through these fine holds followed by its next unique unit hop yields exactly the designated source coordinate from the path equation. This is also the reason an empty supplier cannot inject a live value through a different center's concurrent transaction. □

Lemma 5.11 (Classification at an issued face)**.** *On a face for which its immediately preceding strong issuing cut is valid, exactly one of nominal live and proper dead holds. This implication also holds at a cap cut with the ordinary transition replaced by the constant D. It applies to arbitrary payloads entering the weak-cut calculation.*

Proof. **Precise issuing-cut hypotheses.** On the actual aligned child grid the predecessor has either (i) no computation, (ii) computation with a dead result, or (iii) computation with a live result s in the native state alphabet. In (i) and (ii) all observed matching occupancy on the new face is confined to its exceptional column K, of width at most d, including arbitrary forged live children inside it. An empty exception gives empty occupancy. In (iii) C_* is matched off one interval I_{dec} of width at most $2q$, and every central-fifth coordinate outside K is present with the correct issued letter of the encoding of s. The exceptions in the two tests admit one common interval of width at most $2q$. In particular this follows if the possible C_*

loss is a terminal piece at K whose union with K has width at most $2q$. These are conclusions required of the preceding strong cut, not premises about this tile's classification.

Put $P_0 = P/5$. For $w \geq 100000$, $d = 4w + 8 < q/2$, $2q < P_0$ and $|C_*| > 3q$. At a fixed packet offset u in $\{0,..., P_0 - 1\}$ the three tested positions are $P_0 + u$, $2P_0 + u$ and $3P_0 + u$. An interval of width $\leq 2q$ can hide at most one of them. If two live-state decoding witnesses use different exceptional intervals, at least one of the three positions avoids both. Its carrier therefore equals the two alleged encoded letters at u. Apply this to each u. The packets agree at every offset. The live structural tags and fixed field positions of the serializer determine every scalar and every pair of payload halves, so the two candidate live states coincide.

In cases (i) and (ii) all observed matched children fit a single interval of $\leq$d$<$q/2 sites. This is proper death by definition. Even deleting an arbitrary $\leq 2q$ positions from the C_* test cannot hide its missing sites: at most $d + 2q < 3q$ positions are accounted for. Hence nominal decoding yields D. In (iii) choose the common interval of the issuing cut as decoding exception. The issued live state s passes both central tests, so uniqueness gives $x = s$. At least $|C_*| - 2q > q > q/2$ matched positions exist, excluding proper death.

The pointwise weak vote has no independent formatting consequence. For example if it yields a common malformed live packet z in the three middle fifths off K, it cannot encode a live nominal state: for any candidate decoding exception I'_{dec} and each offset, some replica avoids both K and I'_{dec}, forcing z to agree with the candidate's canonical packet at that offset. Malformation rules this out. This weak description alone gives no proper-death conclusion. When used as a supplier to two unaffected strong repetitions, each correct read receives z at each offset from at least two of the three homologues (one interval shorter than P_0 covers the compromised coordinates). Both fresh parsers on that supplier return D. A weak empty supplier represented by three neutral D packets or absent messages also returns D. Two identical independent evaluations on the resulting three supplied states and fixed operands issue the same canonical packet of a legal live state, or both issue all message-empty. Thus a malformed weak string is never promoted merely because its letters voted live. The geometric/transaction premises that actually provide those two reads, maintain their banks and impose the strong cut are separate obligations.

Shifted faces. For an internally queried protected face run the preceding issuing calculation in its buffered history, with earlier weak descriptions at the three supplier predecessors of that calculation. Apply this lemma after the strong cut at this rank and before invoking classification for proper-death, disjointness, properness, or service of the current face. In constructing that issuing cut the no-split and protected-geometry premises refer to the rank below. Its reading of a weak malformed supplier uses the elementary replica and parser argument above, not classification of that same-rank weak face. □

5.4 Grow sources, density and separation

The terminal geometry adapts Gács [10, Section 15].

Lemma 5.12 (Attachment and local repair of an issued Grow source)**.** *Away from the actual or delayed wound in this epoch, an issued source is an interval attached to its original ordinary core. Its core C_* is immune outside its localized wound, and is wholly present after repair. A genuine distal piece cut off from that interval is included in its ancestry. During repair each of its internal holes lies in K enlarged by $2B + 40q + 40$. If normalization occurs first the assertion ends there with this last bound. An unrelated short satellite may be arbitrarily far from the authentic tip, but cannot supply a long continued arm by itself. Two sources can change the exact division of an adjacent repaired bridge.*

Proof. Use child units. A source is the exact integer base $hP = x - \pi$, with its genuine clock and address residue. Until the final normalization, an authentic child inherits this same

base. A temporary raw-only satellite is counted separately until its first clean Purge. The hypotheses are the preceding-rank nominal and served clauses on all shifted neighborhoods, the established previous ordinary single-interval dichotomy, the prior issuing cut with live decision at all its surviving proper incumbents outside K, and the buffered two-flank clearance implication in Lemma 5.6 (including the previous-row and four-neighbor tests). Lemma 5.7 supplies the actual birth and assignment counts to the first clean Purge for an exact source (the same count as step 1 of Lemma 5.8 for an ordinary home). This statement is about a previously *issued* positive, not a merely formatted array put on an arbitrary initial row.

Source at the vote and real backing. Fix the true old ordinary home of a positive cut. When that cut and its preceding ordinary work are clean the genuine pre-vote interval is nonempty by positivity; no birth occurs in the ballot, scratch is reset and the live decisions retain its proper sites. Only an initially improper extremity may be lost on either boundary of a genuine interval. The new source is its resulting interval. With a late J retain the last pre-J retired ordinary strobe I. If I were empty, the only independent raw support before the first clean Purge would have span $\leq d + 8$; this could not account for a macroscopic positive cut. More precisely choose a positive cut marker in $[P/5, P/3]$ and one in $[2P/3, 4P/5]$, avoiding the formatting exception and K enlarged by $10B$. Each band has width $\geq 2P/15$-$2 > 20B + d + 2q + 10$. Each chosen marker is within $8B$ of I: otherwise a centered vacant interval of $4B + 4$ positions within its host home avoids I,K and both true-edge residues, and the empty-band induction of Lemma 5.5 (also rederived in Lemma 5.15) uses only provisional birth/retirement in this ordinary part. Written independently here, each complete attempt enters $\leq$B sites from each direction and every such still separated entrant has a marked, unexempt, zero-grade face directed to the central empty strip; B removals clear its entrances. Its center cannot be occupied at the ballot. Hence $|I| \geq P/3 - 16B - 4$. This argument only uses the native empty strip and does not assume wide ancestry. When J is in a subsequent Grow round with a clean preceding ballot, the retained ordinary proper pieces of the preceding productive core give the same macroscopic backing.

Run protection. Apply Lemma 5.7 to the actual macroscopic interval just obtained. Its length exceeds $2q + d + 44$ (it is at least $P/3 - 16B - 4$ for a late J, and contains the retained ordinary core pieces in a later Grow round), its strict interiors are proper by lower triples, and its evolution is clean from its latest scratch reset until S: a Temp reset, or the fresh reset of the positive vote, whose live decisions keep the proper incumbents outside K by the prior issuing cut. By the run-protection clause of that lemma each original $q + 1$ ordinary run of C away from K, trimmed by one site at each end, persists throughout Grow.

Clean source. Inside an exact source every adjacent pair is a strong link. At a structural step its strict interior is proper and computes a live result. On a birth tick it also holds nominally, and the next nominal triple of an inner neighbor, old endpoint, and a new outside child makes the old endpoint proper and served. For a singleton with two children use the same triple; for one child, dropping the parent still leaves one interval. Without births only an end of the exact source interval can be lost. The independent clean-evolution paragraph in Lemma 5.13 expands these arguments, including length two; it is a forward pointer and uses no attachment conclusion. Thus a clean genuine component cannot split. Zero-grade true-end exemptions apply to Temp only at the outward true face. A distinct source, including one of the same clock with a positive wrap, has a nonstrong bond to this source at a contact: equality of successive extended addresses would assert equality of $x - \pi$. Therefore a short distal source piece whose only support is the other source's grade-one junction still has a right break and a left break for Purge. When its length is $\leq$q at an initialized reset, its entire strong component clears in that Purge. Pinning by a true-end exemption does not affect the propagation or right-break killing instructions either.

Bookkeeping and refill. Lemma 5.7 supplies the $B + 30q + 40$ cavity bound, the

K enlargement, and reunion or distal extinction within $2d_0 + 10H_0 < w^6$, stopping at normalization when necessary.

The original C had an extra $4q$ reserve beyond C_*. At the cut the issuing lemma retains the long ordinary pieces of C outside K, losing at most the one short edge piece of length $\leq q+2$ there. Every C_* position outside the wound therefore has one of the fixed long ordinary runs of Lemma 5.7, and a missing C_* position at an internal wound has long ordinary pieces on both sides in C, so it is filled by the genuine two-sided case. The only other possible loss of a short edge piece capped at $q+2$ lies within the outer $4q$ reserve. Thus C_* persists except during its localized repair. A raw-only label outside the authenticated interval is precisely the short cleared satellite case.

Uses. This provides the earlier attached interval, surviving core, localization of real gaps and short-pinned-flank alternatives for density, wide comparison, root identity and true-end relief. It does not require the contested geometric seam between two positive homes to retain its old division into exact source labels.

Wound on both sides of a positive vote There is a case in addition to an actually straddling S and two fully separated repairs. Suppose S ends before the positive vote but a completed geometric refill has not yet occurred at that vote. In particular a terminal ballot could appear to reopen K during the previous wound's response time. The following calculation uses the preceding positive-verifier proof in Lemma 5.10; this is available before the current Grow source lemma.

If the ballot's immediately preceding complete Purge started after the last affected transition and the new S is this close to the vote, the first of the two separated verifier plateaus had finished before S (their separation exceeds $3w^6$ while this case lasts less than $2d_0 + 10H_0$). Every surviving position at the final Purge, including positions in K, is certified. Indeed the first complete post-S Purge before the ballot certifies every eventual survivor: in its initial $2(M+1)+2$ propagation rows each such survivor has an intact M-site one-sided interval, with at most $25d+102$ positions not certified by the earlier invariant, and $M-25d-102 > 4d$. A short strong component dies. Each intervening complete Purge recertifies newly recruited neighbors by the same inequality, and the instruction immediately before the final ballot is the complete Purge in question. All these certifications are sticky in the absence of another S. The actual positive ballot thus makes no new inactive hole, even if the payload at K is filler or wrong. Lower masking can remove only exposed ends of an already existing ordinary exact component: strict interiors are proper and have a positive certified native ballot, propagating by the rank-$(j-1)$ proper-successor and service clauses of Proposition 5.3(ii). The ordinary true addresses become the new epoch's normalized addresses at the same coordinates. A panel of length at least q-1 thus loses at most one position on either side across this ballot. Its interior wall and exclusion of incoming paths persist. A short raw-created component already joined to a genuine side does not detach from that side on the clean certified vote. In this case use a single ordinary/Grow response beginning at the original S and continue its birth-handle accounting to the first chosen growth pass; the ballot itself is a no-birth, no-interior-split instruction. The regular handle cadence is continuous at that instant and the fresh reset precludes an old temporary mark crossing the vote. The total time from that S to the end of the selected pass and extra Purge remains $\leq 2d_0 + 10H_0$. Thus conditions (i)–(iv) of Lemma 5.7 hold at this ballot: no ballot decision is non-native, a panel loses at most one site per side, the reset and cadence are as just stated, and (iv) is the standard count because the first complete clean Purge after S precedes the ballot.

If instead the last preballot Purge intersects S or S includes the ballot, then the first raw transition is within $d + 2q + 3$ ticks of the vote. This interval is strictly inside the final no-birth plateau since its length exceeds $100A + 4(M+1)$. Use the clean genuine ordinary

interval just before S. The raw updates and the exceptional decisions at the following ballot are all at positions in K. The latter assertion uses the first clean verifier plateau already completed before S, no subsequent births near this final Purge, and sticky certification off K: the whole immediately preceding plateau contains an earlier full clean Purge that certified its entire surviving group. From the start of S to the first complete clean Purge of the new Grow takes at most $2d + 6q + 10$ ticks, with no native birth before the vote and at most two after it. The authentic intervals outside K are therefore shortened only at exposed ends. A clean proper original ordinary run of $q + 1$ positions off K survives both cuts by the fixed-run proof, and its reset sites remain ordinary incumbents. At most d coordinates inside K have been reassigned at either exceptional operation, regardless of how often a coordinate is replaced, and counting four honest growth opportunities still bounds the assignments by $25d + 100$. A shifted foreign family in the earlier cavity either comes from actual S or an old shifted ballot descendant of an S ancestor; no foreign genuine external lineage crossed a surviving wall. A potentially spurious *canonical* decision at the vote is confined to the same K and the ensuing proper Purge clears its separate short component. Choose clearance after the later of the actual S end and the vote, plus $4q + 14$ rows. This incurs $\leq 2d + 9q + 40$ ticks after the original S, within the t_c of Lemma 5.7. There is no additional full B-probe pass before the selected one. Hence conditions (i)–(iv) of Lemma 5.7 hold at this ballot: (i) and (iv) were just shown, the authentic intervals outside K are shortened only at exposed ends, at most one site per side as in the previous case, and the reset and cadence are as there. Thus the same bound $B + 30q + 40$ for an enduring internal cavity and $2d_0 + 10H_0$ for total elapsed time apply. No passage combines two complete B-growth passes into a budget of four birth instructions.

For completeness, the first raw-only enumeration can leave additional **small canonical inner fragments**: the two rays outside K enlarged by four contain at most one true same-source interval each. A separated component not reaching either ray has at most $d + 8 < q - 2$ members. Before its next whole initialized Purge it can add at most two new members. It therefore dies in that Purge or joins an adjacent compatible flank, including two fragments coalescing inside the narrow middle. The union of all fragments that have not contacted a flank remains in the same original column with one further opportunity on each side, of width at most $d + 10 < q$, so coalescing among themselves does not rescue them. Any contact with an enduring genuine side only enlarges that side. Allow this one additional Purge in the $9q + 40$ buffer of t_c in Lemma 5.7: an initialization occurs every H_0 ticks, so two clean complete purges have ended by last $S + 6q + 4$ (overlapping the elapsed time with the displaced-grid guards). Genuine two-sided pressure may come from different source labels at a repaired seam: the source lemma requires its two internal walls to be the same exact label. When it is applied to a pair across a true wrap it gives only persistence of the two individual backings.

The spatial wall test for the final attempted birth requires only past and present rows. A shifted nominal tile overlapping canonical row t has a preceding served tile whose time interval intersects canonical row t-1 or t for positive duration. Choose a time in that intersection and trace its ancestors backward one unit at a time. A shifted actual blocker is tested directly at its overlap with t. No row later than t is involved; the lower proper-successor test begins one canonical row earlier. At the last attempted birth before a join or extinction the wall need endure only through the old row of that update. □

Lemma 5.13 (Deficit from a component count)**.** *Let $H = [0, P - 1]$ be the home of a positive issue. In its comparison period from $r + \lceil v/20 \rceil d_0$ through early* Shrink(0)*, suppose (i) every authentic source before the single exceptional rectangle was one interval, possibly empty, (ii) the only additional exceptional cut is one affected ballot at which only positions in K can make erroneous decisions among the ordinary pre-vote survivors, and (iii) outside those exceptions the preceding-rank regularity, nominal triple and service clauses hold. A previously earned cut*

may initialize a source with at most one interval of erroneous positions in K, or two genuine flanks around K left by a straddling pulse. All intervals here refer to a single exact extended label until normalization. If every interval of missing positions in H has length less than $L = 60B + 3d + 40$*, the total deficit in H is at most* $(5(4d + 12) + 1)L$*. In particular it is at most 2000000 A for* $w \geq 1000$*. The same bound holds before, within and after the event, and on entering ordinary work. An exact count with* $d + 2$ *components per source is unnecessary.*

Proof. Time and space are integer child-grid indices. The initial source interval comes from the already established issuing cut and the source geometry of Lemma 5.12; the clean-evolution paragraph below is independent of that input.

Proof of clean evolution. Follow a connected interval of one source whose neighbors at both ends do not belong to the interval. During a birth step each old site has a live nominal successor with the correct clock and address and retains its source label. If there are three or more old sites the strict interior is proper by the nominal triple, so each interior site supplies its successor. If the newly born child immediately beyond an end is served, the former end has a live nominal successor between the old inner neighbor's live nominal successor and the live newborn. It is therefore proper and served at the new date. In length two this same argument applies at each side that gives birth. Any subset of a two-site old interval is an interval; two attached children cannot detach. For length one with two born children, the middle has a nominal matching triple and is supplied. With just one child the union of that child and the surviving old singleton is an interval even if the singleton drops. On a no-birth step at fixed source labels all strict interior cells were proper. A strict interior faces grade two on both sides; neither Purge killing, either directional temporary removal, Shrink nor native withdrawal acts there. Thus only interval endpoints can be removed, and every surviving interval is connected. These facts also apply to an interval artificially truncated by a moving boundary if a one-cell margin is lost rather than gained. Two old intervals may merge, but this never increases the count. Matching incumbents belonging to other labels never switch label except at normalization and the ballot.

Effect of a pulse. Enclose the raw spatial interval K (length $\leq$d) at the start in $[a, b]$. At its ith subsequent fine date enclose $[a - i, b + i]$, for $0 \leq i \leq d + 1$. The last interval has length $\leq 3d + 2$. No locus in the remaining left half-line at time $i + 1$ has an old neighbor in the excised region at i; the same holds on the right. On either exterior half-line the surviving descendants of each of its old intervals thus form at most one interval by the preceding argument and restriction to a half-line. This remains valid whichever admissible end omissions occur, without comparing to any counterfactual history. Inside the cut-out region any subset has at most as many connected components as it has sites. Starting with a single old interval, the total is at most $3d + 4$ per source at every intermediate date. Starting with two intervals costs two additional exterior components. A delayed affected ballot can delete any subset of K or forge a matching member there. For arbitrary membership replaced on an interval of length d the number of components increases by at most $d + 1$: restriction outside it can split at most one existing component into two, and the new inside has at most d components. If the ballot is inside the excision it has no exterior cost. An honest positive ballot preserves the occupied interval apart from possible improper ends (all incumbent live decisions off K are valid); an honest negative ballot leaves it empty. Counting both exceptional operations, including a vote outside the pulse, gives at most $4d + 7$ per source. We allow $4d + 12$. If the changed cut is first, cutting K out leaves at most two exterior pieces of the preceding genuine source and interior K at that cut. The expanding cone argument then applies directly since it encloses precisely the same column K. A distinct actually empty source can start only in K. After this pulse clean evolution cannot add a component. At normalization the matching material of H is the union of its contributing sources intersected with H. Taking a union does not increase the sum of component counts.

Counting homes and gaps. Translate so a home starts at zero. The base of a possible exact source is hP for an integer h. An occupant x in H has legal address x-hP in $[-1.1P, 2.1P)$. If $h \le -3$ this is $\ge 3P$, and if $h \ge 3$ this is $\le$-2P-1. Hence h belongs to {-2,-1,0,1,2}. Intersecting an interval with H leaves an interval. There are at most $5(4d+12)$ occupied components, hence at most $5(4d+12)+1$ missing intervals counting both boundary gaps. Multiplying by L proves the estimate. If normalization occurs before the pulse, apply the expanding cone to the two remainders of the whole matching home. Splitting old components at the two walls adds at most two, then the cone adds at most $3d+2$; before a first pulse at most five authentic intervals have entered and at most five remain. A delayed cut in the same column satisfies the same replacement bound. All these counts are smaller than $5(4d+12)$. For $w \ge 1000$ one has $d \le 4.008w$ and $B = 128w^2$, so the first factor is at most $81w$, $L \le 7681w^2$ and their product $\le 622161A$.

Uses and dependencies. Combined with the early productive-home sterile-strip implication (no interval of L missing positions), this estimates a single deficit on a common date for all markers. During later ordinary dates the computing core bound supplies density. After the ballot use the held core of each issued source. These conclusions enter nominal separation and the comparison against terminal contenders during end relief. This lemma does not derive its starting premise of genuine source intervals, the true classification of decisions at the vote, or the productive-home sterile-strip implication. A source-preserving component bound remains valid even if the source has not yet healed its one wound. □

Lemma 5.14 (Thin ancestry and productive-home density)**.** *Fix one matching tribe in its final v-round withdrawal and ensuing ordinary period before its next ballot, under the protected lower structural clauses outside K×S. Every genuine terminal observation at n has a member of the same tribe within* $100B$ *of n at each earlier date of this phase interval; an isolated immediate raw-only fragment is instead allowed in the short column transient. In a home H whose ordinary work issues positive at T, no interval of* $60B + 3d + 40$ *sites in H is wholly vacant at any fine date of the final withdrawal or of ordinary work before the early seed cutoff (at all such dates at least* $3P + 10$ Shrink(0) *inspections remain). With Lemma 5.13, the end budgets and late core of Lemma 5.8 and the held core of Lemma 5.12, at every date of* $[r + \lceil v/20 \rceil d_0, T + s_0)$ *the deficit in H is at most* 2,000,000A*; no dates in that interval are excepted from this deficit bound.*

Proof. Use one matching child tribe in its final v-round withdrawal and following ordinary period before its next ballot. Work on a buffered test for which the previous-rank exact-or-dead and structural conclusions hold outside K×S. Write J empty if no top exception is needed. Only the B provisional attempts per Main cycle can recruit in this interval. Normalization preserves residues and changes no coordinate. An unflagged newly matching state must arise from a prior matching incumbent or donor. Distinct exact extended sources of this tribe have different $x - \pi$, so cannot lend strong backing across a source boundary.

Thin trace. For a genuine terminal observation n at t, each earlier date u of this phase interval has a member of this same tribe within $100B$ of n. An observation which has only isolated recent raw ancestry in a span of $d + 8$ at K, before its first clean complete Purge, is allowed as an alternative.

Suppose the interval of radius $100B$ is empty at u. Work separately in the integer bands of distances from $3B$ through $90B$ on either side of n. Remove the coordinates in K and those congruent to either true endpoint residue. Each candidate band is shorter than $P - 2$ and contains at most one adjacent endpoint pair or single endpoint. The deletion leaves at most three intervals of total length at least $87B - d - 2$, so one has length at least $8B + 8$. Select such a band on each side. Starting from u the remaining first partial probe can move at most B places towards the central portion of a chosen band from either edge, and the two directions

together cannot reach its central $4B + 4$ places. No true Grow attack is scheduled here. Only provisional recruitment enters these particular vacant strips, and its marked entrants cannot join across the vacant center or reach a true endpoint. The empty-band clause of Lemma 5.5 gives a fixed vacant central site at every subsequent fine date and restores a $4B + 4$ strip to vacancy at every strobe. The intermediate region between the two permanent barriers was empty at u. No actual matching descendant can arrive from the exterior on any tick: it moves by at most one, its boundary crossing requires the fixed vacant site, and a true nonbirth cannot spontaneously create a matching state. Thus any later occupant near n needs a raw assignment in this interior. From the first exceptional step through the first initialized clean Purge after its end there are at most four genuine recruitment instructions ($d + 4q + 6 < 3(2q + 2)$ for sufficiently large w, allowing rounding at both end dates). Its spatial support lies in K enlarged by four on either side throughout. At complete purification every same-label strong component has length $\leq d + 8 <$q; by q propagation and q killing ticks it is empty. If normalization intervenes before the full Purge, count the whole matching union, still in the same $d + 8$ span. No scheduled Purge straddles the normalization. The first complete clean Purge still finishes within $d + 4q + 6$ fine ticks of the first exceptional step. This gives precisely the stated short alternative or a contradiction. The proof fixes u arbitrarily, so it does not require a single material particle surviving at every earlier date. If the observation immediately follows a no-birth cut, use its actual preballot incumbent or a recent raw substitution before applying this statement.

Empty interval incompatible with a positive issue. In the same home H whose ordinary work issues positive at T, no interval of at least $L = 60B + 3d + 40$ sites in H is wholly vacant at a fine date in the final withdrawal or before the early ordinary seed cutoff, provided at least $3P + 10$ Shrink(0) inspections remain. Removing K and the two true boundary sites leaves at most two nonempty interior pieces in H (the two endpoint deletions in H do not cut an interior piece). Their total length is $\geq$L$-d - 2$, and one contains an empty band of length $8B + 8$. The first partial phase and sterile-band proof carry its central vacancy through normalization and all ensuing stage-0 inspections. On either side of that strip the nearest occupied strobe frontier recedes by at least one per eligible clean Shrink(0), by the erasure proof in Lemma 5.5. After discarding at most three inspection spacings intersecting the one pulse and its boundary strobes, one of the two remaining stretches has more than $P + 1$ inspections. If it follows the pulse all matching home material is extinguished. If it precedes the pulse only the later raw-only family can repopulate H, and it disappears in its first complete clean Purge unless the terminal no-birth interval and the pulse overlap. In that last case only cells in K can remain at the ballot. Such a group cannot meet the positive nominal support test. This contradicts the positive issue. The six-home seed search uses only a polynomial number of fine ticks past normalization, leaving more than $3P + 10$ calls. Once the ordinary interval and threshold ledger apply at clean strobes, occupied C and its possible spatial band of width less than $.4A + d + 2$ give deficit $\leq 24A + .4A + d + 4$. The band argument applies at every date with more than $3P + 10$ Shrink(0) calls left. This includes the entire initial $50A$ d_1 relaxation and the polynomial common-seed search, so it overlaps the strobe ledger. In the later stage-zero and stage-one work the respective end budgets are $3A + 3\Delta$ and $6A + 3\Delta$. Provisional attempts between strobes do not decrease old occupancy, and the next true Shrink and optional old-end loss cost at most the stated edge allowances. The single internal wound lies within K enlarged by $.2A$ on either side. Stage two and computation use the $12A$ core bound. After T the positive source's held C_* across Grow gives deficit at most $24A + 8q + 2$ plus its single localized wound of width $\leq 4B + 80q + d + 82$. Thus throughout $[r + \lceil v/20 \rceil d_0, T + s_0)$ a constant $2,000,000A$ bounds the common deficit, using the explicit component count of Lemma 5.13 for the earlier dates. All those counting hypotheses concern the already proved previous ordinary cut and the genuine-source geometry, not the conclusion

of disturbance inheritance.

Uses. The thin conclusion supplies an actual earlier witness for the remote-marker estimate and for end relief. The positive-issue deficit is counted once per date against an entire competing interval. □

Lemma 5.15 (Remote sources during terminal withdrawal)**.** *Fix one child-grid tribe in an epoch whose final v-round withdrawal starts at r. Let n be the coordinate of a genuinely supported matching member observed at withdrawal age at least* $7v/10$*, or later through the final ordinary preballot date. Assume authentic-source attachment to enduring cores reaching within* $12A + 4q + 2$ *of both true edges, K-only extinction at the first complete clean Purge, repair length at most* $w^6 + 2d_0$*, and the thin trace of Lemma 5.14, under the protected lower structural clauses. If the spatial home H containing n has dead nominal input, then a live home is within* $300B$ *of n, or matching material covers* $[n - v/5, n + v/5]$ *at every fine date in* $[r + \lceil v/20 \rceil d_0, r + \lfloor 3v/10 \rfloor d_0]$*, apart from at most* $(220A + 2w + 20)d_0$ *fine dates. An immediate cut remnant is first traced to its actual preceding preballot observation.*

Proof. All spatial coordinates and fine times refer to one child grid. A strobe is the row immediately after its true Grow withdrawal. Write $k = 12A + 4q + 2$, $W = 4B + 4$. Enlarge w to at least 10^9. Hypotheses supplied by Lemma 5.12 and the earlier structural lemmas: throughout the final withdrawal, the authentic sources have enduring ordinary cores reaching within k of their two true edges at each clean strobe, each source is connected to its core outside the single repair interval, and an independent K-only source perishes at its first complete clean Purge. The repair interval has length at most $w^6 + 2d_0$. Lemma 5.14, including its raw-only alternative, gives an actual member within $100B$ of a genuine late marker n at each earlier date in this withdrawal. We use that lemma on a date away from repair, not an inference that the same individual lived throughout.

Sterile band. Start from W consecutive empty sites at a strobe in a band disjoint from K and positions congruent to 0 or P-1. In the final withdrawal no Main instruction recruits. At each subsequent cycle the only matching new material is a marked Heal birth, which requires a previously occupied adjacent donor. In a growth pass at most B sites can be entered from either side. The central $W - 2B$ sites stay empty. Each entering one-sided component is marked at every interior point of the band. Its inward extremity has grade zero facing the empty center and is not exempt there. At each of the B temporary-removal instructions the foremost remaining site is removed, unless an intervening Purge or a prior omission already removed it. Its distance from the corresponding band boundary decreases by at least one until the entrance is empty. Raw sites and true-end exemptions never occur in the band. This induction applies to the whole matching tribe, so a different same-phase label cannot make the two trials meet. The central subband stays empty at all fine dates. The full band returns to empty at every next strobe. Across normalization, if needed for the thin lemma, use its separate strobe boundary with a new Cn initialization.

Geometry at a clean strobe. Take the nearest authentic home to the left of H and to its right, omitting a side which has no contributing home. Their order follows from the enduring core and attachment hypotheses: a further connected source would have to occupy the core of the nearer one. At a joined row their seam is a positive bond, since the left address is at least $P - 1 - k > P/2$ and the right address at most $k < P/2$. Each side has its own long strong backing, and all seam sites are in a nominal triple. Neither a true Grow, a temporary removal, a Purge nor a clean service omission can open this seam. At a first contact the two provisional chains already have positive facing bonds, so the B temporary removals do not undo their join. The next Temp reset makes the retained sites unmarked. The whole interval between the two deep cores is thus occupied until the next event. It contains the claimed band at every such early joined date, because each core extends into its own home a distance

greater than $P/2$ and H is between the homes.

Otherwise write p for the left source's rightmost actual site, z for the other's leftmost site, e,f for their facing true edges. Define $U = \max(e, p)$ and $V = \min(f, z)$. The actual open gap has length z-p-1. Its excess over V-U-1 is $(U - p) + (z - V)$, an integer in $[0, 2k]$, since the cores are occupied. Here $2k = 24A + 8q + 4 \leq 25A$ for $w \geq 10$. There is at least one whole home between e and f, hence $f - e - 1 \geq P$. On a retained clean unjoined cycle an exposed terminal tip can have a positive outward bond only to material of the same tribe, with the clock and address-residue alignment required by the bond rule. A neighbor with consecutive exact address belongs to the same root, contrary to extremality; a grade-one neighbor from the opposite authentic root is the joined case. The only other possible supporting label is a J-made remnant of width at most q. Its repair cycles are discarded, and at its next initialized Purge it has no q strong predecessors within its own exact label and dies; a grade-one junction supplies no strong Purge backing. Thus no foreign tribe or short supporting label can pin the tip on a retained unjoined cycle. Each extended exposed site then has a free outward face and withdraws or has already disappeared. No new unjoined temporary site can remain beyond the true edge, since each faces an empty site, carries a mark and has the full B withdrawals. If a true endpoint is reached from within its own home the projection already pays for that advance. In particular $U_{next} \leq \max(e, U - 1)$, and $V_{next} \geq \min(f, V + 1)$. If the projected gap is below $P/2$, both variables cannot simultaneously equal e,f, so the projected gap grows by at least one. On the first open strobe $V - U - 1 \geq -2k$ because $z - p - 1 \geq 0$. Consequently after $110A$ clean unjoined rounds, either $V - U - 1 \geq P/2$ or the actual gap is at least $110A - 2k > 40B + 2d + 40$. An interval of length $40B + 2d + 40$ is shorter than P-2; deleting K and the pair of true-end residues removes at most $d + 2$ places and leaves at most three intervals, unless the residues lie at opposite endpoints of the tested interval. In the latter case deleting endpoint residues creates no additional intervals. Some remaining interval has at least $(40B + d + 38)/3 > W$ places. It contains a sterile band. Once such a band is certified the two sides cannot subsequently meet. Discard the first $110A + 1$ comparison cycles and the first $110A + 1$ after the single disruption, plus all cycles meeting the repair interval. This works for every good early strobe. With only one side, no separating band is needed for its clean-tip recurrence.

Permanent upper bound across the disruption. This argument uses a second band and does not bound the number of provisional passes by a constant during the broad w^6 repair. Fix a retained unjoined early strobe. If the disruption lies between it and the late comparison, take U in the gain estimate at the last clean strobe before the disruption: the band certified at the early strobe persists to that strobe. If no disruption intervenes, use the clean recurrence alone. Consider the left root at the selected strobe. The two-root case at such a date has already waited out the $110A$ grace period and has a certified band; joined dates already supply the wide coverage. If a previously certified separating band exists, call its leftmost site h. It is empty at this strobe, and the genuine left root lies to its left. If $h \leq U + 30B$, the farthest connected site of that root after restoration is strictly less than $h + W$, since it cannot cross the center of the sterile band. Thus $U_{after} \leq U + 35B$. If $h > U + 30B$, or no opposite root exists, the first $30B$ places to the right of U contain no matching member at this strobe. In the two-root case the right source lies to the right of the original band. In the one-root case there is no other connected root on this side of its own core, and the observation window need only extend $30B$ beyond the relevant tip. Remove K and the pair of true-end residues from these $30B$ places. The remaining pieces number at most three and total at least $30B - d - 2 > 3W$. Select a new sterile band there. Its central gap stays empty throughout the disturbance, and the restored genuine source is connected to its original core, which lies behind that band. Its rightmost postrepair site is at most $U + 30B$. A temporary same-label source forged on the far side must either die, or merge

with a similarly labeled authentic source. Once it is restored and connected to this root it cannot be on the far side. In both cases $U_{after} \le U + 35B$. Exchanging left and right proves the opposite bound. The test uses only old certified vacant space and postrepair attachment; it needs no claim that a particular exposed tip remained the tip during the pulse.

Iteration and dates. Choose the late comparison strobe between rounds $\lceil .64v \rceil$ and $\lfloor .66v \rfloor$. Fewer than $w + 6$ strobes meet the repair interval, by $d_0 = (2A + 4)(2q + 2) \ge 4w^5$, even including two rounded endpoints and the liberal extra $2d_0$. Such a clean strobe exists since $.02v - 2 > w + 6$. Thin tracing supplies a live site within $100B$ of n there. Its source is authentic by clearance of the short transient. A clean joined early strobe was already dealt with, irrespective of a later break. At an unjoined retained early strobe choose its permanent separating band as above. The late supported root cannot have passed through this band, and it is connected back to its original core. Thus the root on the late side really is the same rooted label on that side at the early date. This also treats an earlier joined configuration opened by J: its early joined dates already have full coverage, its first $110A$ postrepair rounds are discarded, and every subsequent separated date bears its own newly certified sterile band.

For the late left root put $N = w + 6$, $D = 35B$, and M equal to the difference of the two strobe indices. At clean rounds the projected recurrence holds; during the single interruption bounded by N skipped rounds use the permanent gain D once. Composing a clean block of a steps maps u to at most max(e,u-a). Composing two such blocks with the gain between them gives $U_{late} \le \max(e + D, U_{early} - M + N + D)$, also valid if the gain is absent. If all live homes are more than $300B$ from n, left-side supply within $100B$ forces $U_{late} \ge n - 100B > e + D$. Therefore $U_{early} \ge n + (.64 - .30)v - (N + 135B + 2) > n + .33v$. This is a truly extended tip, so its connected interval extending back to the old core covers the required band, even allowing $B + 2$ places of loss when passing from a strobe to a fine date. More precisely, on a full clean round a macroscopic connected proper side cannot suffer a native interior loss. Its post-withdrawal tip and its preceding tip differ inward by at most two (one improper endpoint and one actual withdrawal); on partial rounds no interior point common to those strobe intervals can disappear and reappear behind an unfilled break. Alternatively repeat the calculation at both endpoint strobes and use persistence of the large strict interior. The right-root derivation follows upon replacing every position x by -x and defining its projected outward extent as -V; the estimate concerns the true Grow and temporary directions on exposed sites and relies on separate q-backing for Purge on each actual large source. The count of discarded cycles is at most $220A + 2w + 20$, including neighboring strobe cycles for all sampled fine dates. The sets deleted depend on the common disruption and the epoch start, so they apply simultaneously to every n.

Uses. Nominal inheritance at a deep issued marker, and terminal exclusion for end relief, retain the $300B$ near-source option and the $.2v$-wide band. Both tolerate the explicit $35B$ permanent gain, which is already paid by the $300B$ cutoff. The linear recurrence and the two sterile bands require the preceding authentic-source attachment and core lemmas. Neither hypothesis follows from formatting an arbitrary input cut. □

Lemma 5.16 (Backward inheritance across an incoming pad)**.** *Suppose two nominally live tiles U',V' of different geographical origins have overlapping temporal rectangles and one spatial home meets the other's* $64A$ *pad. Then their same-home predecessors U,V one epoch earlier were both nominal live, and the translated pair retains the same overlap/pad relation.*

Proof. This proof uses the earlier positive issuing cut of Lemma 5.10, productive-home density of Lemma 5.14, protected lower exclusion and Lemma 5.15. All examined work rectangles and the preceding supplier computations are inside the common sparse test. Positions and dates below are child coordinates; fractional offsets of different child grids are retained when intersecting rectangles. At most two positions at the ends of any interval need be discarded

to select whole child tiles of one grid. Select all comparison dates in the interiors of fine tiles. Choose $w \geq 10^{10}$, large enough for the explicit component-deficit constant 2000000 and the at most $(220A + 2w + 20)d_0$ date deletions for each of the fixed comparisons.

Identity. There is no assertion that an individual physical site carries one exact extended address from the terminal ordinary date through Grow normalization. At a genuine late marker thin tracing supplies a supported matching member within $100B$ at a clean strobe in $[.64v, .66v]$ of the same final withdrawal. Before normalization its extended integer identifies one exact source. At each good early strobe, there are the following exhaustive cases.

1. There are two roots and their intervals join. Their union at that very date is the required band; no terminal-label conclusion is needed then, even if J later opens a break at their seam.
2. There are two separated roots. The $110A$ clean-round waiting bound yields a fixed previously empty $4B + 4$ strip between their tips disjoint from K and all true ends. It is empty in its center throughout the remaining withdrawal. The two cores persist on opposite sides of this strip. A same-label forged distal site cannot change this order after restoration, since the restored interval connects to its core and cannot cross the empty center. The assignment at the late clean strobe therefore identifies the same authentic side and same exact label at the early comparison. A root farther out is separated by the nearer core and cannot be the source at the late strobe. Apply the projected-tip recurrence to this particular side, charging at most the single $35B$ gain across J by the second sterile strip. This covers a raw-created fragment meeting a genuine source, since a postrepair connected member cannot be beyond the certified band. A gap first opened by the disruption is either sampled in its $110A$ discarded rounds or already admits its own newly certified band.
3. Only one side has an authentic source. With no opposing root the single free-tip recurrence applies, using the second sterile strip to bound the gain over J.
4. Neither side has a source. At a clean strobe with H dead this contradicts the existence of a genuinely supported marker.

These are the alternatives at the two sampled unexceptional times; no side attribution at an exceptional strobe is presumed. In particular a live-state label at a sampled damaged physical position is never treated as a certificate of descent.

Translate previous issuing periods of U' and V' to $[0, T)$ and $[b, b + T)$ with $0 \leq b < T$, interchanging names if needed. Let $l = \lceil v/20 \rceil$, $h = \lfloor 3v/10 \rfloor$; the reference comparison window is $[r + l\ d_0, r + h\ d_0]$. For each positive issue the earlier density lemma covers the whole time interval $[r + l\ d_0, T + s_0)$ in its issuing coordinates with deficit at most C A at every date in its spatial home. Here $C = 2{,}000{,}000$; if necessary enlarge w so $w > 100(C + 1000)$. This choice can be made after the constants are established. We have $r + h\ d_0 < s_0$, hence the entire shifted comparison window starting $b + r + l\ d_0$ is in both issue-density spans (a fixed one-tick rounding margin may be deleted). Pick a common retained date outside S, deleting this additional interval of at most d dates. If the lower grids are different, lower nominal exclusion forbids physically overlapping occupied lower rectangles from the two tribes. If the lower grids are identical, the required lower age or address residues differ, because simultaneous equality of both residues at a common lower index would identify the two upper origins modulo their time and space periods. The same lower state cannot realize both. At most C A lower cells are missing in each home, so spatial overlap has length at most $2C$ $A + 4$; otherwise a common occupied physical position would remain after removal of both sets of missing intervals and the two boundary roundings. If the homes are disjoint, their gap is already $\leq 64A$. Their facing edges thus differ by at most $(2C + 64)A + 4$.

Suppose first that V was nominally dead. Choose a terminal marker inside V between $14A$ and $18A$ from its facing edge, avoiding K enlarged by $100B$ and the at-most-$2q$ exception.

The available length $4A$ exceeds $200B + d + 2q + 10$. Use the last preballot row: the positive cut supplies occupancy off the common $2q$ defect throughout this band, its vote cannot recruit into a vacant slot and any affected coordinate is in K. Consequently this is a genuinely supported terminal member. Since V is dead its own home is excluded as a positive input; any other true home is at least $14A - 2$ coordinates from this marker, exceeding $300B$. The remote alternative yields actual matching coverage across its $.2v$ band at all retained dates of the shifted comparison window. Towards U its center is at most $(2C + 82)A + 6$ from U's facing edge; it extends at least $.2v - (2C + 82)A - 8$ into U. This exceeds C $A + 2$ by the choice of w, and the remaining width of U exceeds v. The band and U's density at a common retained date contradict lower exclusion. Thus V was live.

If U was dead, choose the matching terminal marker in U identically. Its own remote band now lies in the unshifted comparison window. The density spans of the positive issues V and V' begin at $b - T + r + l\ d_0$ and $b + r + l\ d_0$, respectively, and their ends are $b + s_0$ and $b + T + s_0$. They overlap since $r + l\ d_0 < s_0$. Their union contains the whole unshifted comparison window (use $b < T$ and $r + h\ d_0 < s_0$). Removing the remote-comparison date set for U (the density spans of V and V' have no excepted dates) and the at most d dates of S still leaves a common date, since $h - l - 2 > .24v > 220A + 2w + 21$ for the stated large w. At that date either V or V' supplies density over their common home. The identical spatial calculation contradicts lower exclusion. U is therefore live. Translating both tiles back by T preserves their physical overlap and the $64A$ pad, proving inheritance.

Uses. Apply inheritance to genuinely overlapping U',V'. Positive predecessors force the preceding computing periods, which allows the separate full-home common-date exclusion using the long Shrink(0) stages. The contrapositive supplies incoming clearance of a positive successor when its predecessor was proper dead or undisturbed live, once competing nonproper-dead tiles are classified. A positive current input alone does not assert service of a successor. □

In the next lemma use the earlier rank-j inheritance, seed, issuing, classification and held-flank implications as well as the rank-$(j - 1)$ hypotheses.

Lemma 5.17 (Overlap exclusion and properness of a matching middle)**.** *Under the rank-$(j - 1)$ hypotheses, distinct nominally live rank-(j) tiles in protected work rectangles do not overlap in space–time. The middle of three adjacent matching nominal positive tiles is proper live.*

Proof. Suppose two genuinely overlapping positive homes U',V' of different origins existed at this rank. Their input faces and their predecessors, shifted back one whole epoch, have the same relative overlap. Lemma 5.16 supplies live U,V and, applied to them, their two live predecessors. In the ordering of the inheritance proof one foreign issuing period $[b - T, b)$ lies in the union of the two canonical issuing periods $[-T, 0)$ and $[0, T)$. All three works compute, since each has a positive nominal issue. Foreign Shrink(0) lasts more than $.85T$. At least one of the two canonical periods contains at least $.425T$ of it and less than $.15T$ of this intersection lies outside that period's Shrink(0). Exclude each home's first $50A\ d_1$ dates, its at most $N = 200A\ d_1$ nonfull fine dates afterwards, and at most $4w^6 + 20$ for the common slabs and four fine boundary rows. The excluded length is at most $2(50A\ d_1 + N + 4w^6 + 24)$, a polynomial in w smaller than $.275T$ for sufficiently large w. A common remaining physical date has both whole homes occupied. If the grids are displaced the overlapping physical spatial intervals contain a point on both occupied lower rectangles, violating lower exclusion. On a shared child grid two distinct higher origins require different child age or address residues at a common index, impossible for a single served state. Half-open rectangles with mere touching at a single endpoint are excluded by the overlap hypothesis.

Now take three consecutive matching nominal positives. All three prior works were computing. Their common full ordinary seed from Lemma 5.9 and two inside bridges keep the entire middle home and the two neighboring rims joined through the ordinary work except for the one localized response. Its true end sites start proper by the adjoining lower triples at the seed. A late split at K leaves its long proper old sides intact in computation; either both reach length $q + 3$ or the only short side lies next to K, of length at most $q + 2$. The paired repair of Lemma 5.6 otherwise restores both in-home sides well before the ballot. If K straddles an edge use the independent backing of the adjacent home; grade one at the wrap does not carry Cn. The positive ballot preserves all unhit properly covered cells (Lemma 5.10). Hence $[1, P - 1)$ is present off one interval consisting of K and its possible adjacent short piece, of length $\leq$d$+q + 2<2q$. For an early deep Shrink hit the at-most-two additional losses are healed by the buffered repair before the long computation. A different origin with a nonproper-dead tile touching the middle home's $64A$ pad is classified as nominal live by its preceding issuing work (Lemmas 5.10 and 5.11). Since $64A < P$, it would overlap the middle home or one of the two neighboring live homes in space and time. This contradicts the overlap exclusion just proved. The middle nominal positive is therefore proper. Neither separation nor this triple argument uses current-rank service. □

5.5 Current service and proper successors

The end-relief strategy adapts Gács [10, Section 16].

Lemma 5.18 (Service of a proper endpoint under live holds)**.** *Let a lower site at space i, time a have a proper live nominal state. On any integer interval $[a, b]$, suppose at each $t < b$ that the actual native instruction F evaluated on the three served lower operands at (i,t) prescribes a live output at i. Then at every t in $[a, b]$ the site is proper live and served, and for $a \leq t < b$ its next state is the specified output. Neither spatial neighbor is required to be live after the transition. The hypothesis is pointwise on F and thus includes a positive ballot.*

Full-home specialization. A full home of proper ordinary sites with a genuine common clock and a fresh Cn reset remains full, proper and served up to its ballot; the same is true for a mid-period start. If every native ballot result is live, it stays so across that ballot and its scratch reset, independently of exterior neighbors.

Proof. Fix a step of the rank-j induction. Work on the common aligned rank-$(j - 1)$ grid, inside the domain where the preceding-rank organization conclusions apply at both times. When a rank-j exception is allowed this lemma applies only to steps so protected. At rank zero take actual unflagged transitions and identity service. The following uses lower proper-successor clause in exactly the form "proper-dead or undisturbed live input has a proper nominal successor."

Write x,g,y for nominal state, service indicator and served state. Suppose x(i,t) is proper live. Its definition includes incoming clearance, so it is undisturbed. Lower proper-live service gives $g(i, t) = 1$; consequently $y(i, t) = x(i, t)$. Lower exact computation then gives $x(i, t + 1) = F(y(i - 1, t), y(i, t), y(i + 1, t))$. This value is live by hypothesis. Apply lower proper-successor clause to the undisturbed input at (i,t). Its actual nominal successor is proper; the classification and its live value place it in the proper-live case. Lower proper-live service at $(i, t + 1)$ serves it. The first application uses the given start; iteration up to b proves every assertion, and for $b = a$ lower proper-live service alone suffices. No step applied the triple rule at time $t + 1$. At the physical base $x = y$ identically and the unflagged rule performs F directly; no lower properness conclusion is needed there. All invocations below positive rank use the already established preceding-rank conclusions, never a same-rank assertion.

Native full-home specialization. On the old row suppose all P cells of one home have proper ordinary addresses 0 through $P-1$ and the same genuine clock, and are freshly reset at the start of a work period. Induct up to but excluding its ballot. The same induction applies from a mid-period start, which needs only ordinary addresses, a common genuine clock and a Cn reset in addition to fullness and properness: the Temp-removal clause exempts the outward true face even for an initially marked endpoint, and every other face is positively bonded. At a birth instruction an occupied destination ignores proposals. A true Grow withdrawal cannot target any of these ordinary addresses. A temporary retirement at an interior face sees a positive bond. An outer face at an extreme address is exempt in its outward direction, and the opposite face has a positive bond. For Shrink(i) each interior face has positive grade and a true outward face has depth zero, which fails the strict comparison with $3iA$ for $i = 0, 1, 2$. At the start of every Purge the sites have Cn=1. In its q propagation ticks a site at address $k \geq q$ has Cn=1 afterwards: after h ticks this holds for all $k \geq h$ by induction through the exact integer increment from $k-1$. During killing, every $k < P-1$ has a strong right neighbor on the old row and thus cannot die. At $P-1$ the bit is one since $P-1 \geq q$. These facts hold successively on all q killing rows since each row retained all sites and killing does not change Cn on a survivor. Normalization fixes these addresses; other pre-ballot instructions leave incumbents live. The established pointwise lemma transfers each native hold to a proper served next row. Synchronous induction over the entire home therefore proves all P sites served until the pre-ballot row for any outside neighbors. If each native vote at the ballot returns live, the same pointwise lemma transfers all P votes across that ballot, even when the two neighboring homes vote dead. The ballot's scratch reset, ordinary addresses and zero age give the initial condition for another work period.

For a newborn site, proper deadness of its predecessor and the lower successor clause yield properness once the positive birth is known. The endpoint lemma applies to protected updates with an established positive native result. It does not assert that an eligible birth occurred, that a physical exception leaves a site proper, or that an arbitrary ballot votes live. □

Lemma 5.19 (Terminal competitors at an ordinary end)**.** *Use the issued all-date density, lower overlap exclusion, and the terminal near-source or remote-band alternatives, with the proper-current-input or known-positive-successor anchor described in the proof. For the four-site and four-row stencil about a true end, at every query date of a recovery search whose request lies in $[s_0, s_1 + 20vd_0]$, through the following $3vd_0 + 2w^6 + 12d_0$ search (all such dates are below $T-2$), terminal and early-age contenders are excluded off one common physical raw/repair date deletion, of length at most $w^6 + 2d + 50$ including its four-row halos. Remaining genuine contenders are attached active Grow sources whose positive homes are disjoint from the canonical $64A$ pad and, given the retained middle wall, lie on the far side of the pad from the home (strictly to the right of the pad at a right end), with intermediate Grow age and epoch start strictly between 0 and T. The end's retained near and middle walls are hypotheses up to the first excessive-loss test.*

Proof. This is the terminal and early-age elimination used in the uniform end-relief argument, prior to selecting a single active competitor. It uses all-date issuing density (Lemma 5.14), held cores, the whole-tribe thin trace and the near/remote/short alternative in Lemma 5.15, the protected lower exclusion, and backward inheritance in Lemma 5.16. None assumes current-rank service or directed advance. The right end is P and its prospective query strip $Z = [P - 6A, P + 15A]$. For a left end use the physical interval $[-15A, 6A]$ and distances from zero; spatial overlap, held cores, thin tracing and birth counts below do not depend on pressure propagation direction. In proofs using Purge on that side, it is still the rightmost position with q old predecessors that prevents the purge of a long run.

Claims and surroundings. The canonical work period of U is $[0, T)$. The query has time t between s_0 and $s_1 + 20v\ d_0$, allowing the following $3v\ d_0 + 2w^6 + 12d_0$ search. These dates are all below $T - 2$: after Shrink(0) the later two stages alone have 4Pq d_0 ticks, exceeding this extension. The following anchors are distinguished.

1. A proper live U has its own current incoming clearance. Up to the first excessive loss its near-end ordinary wall is sustained. Its issued-source preceding C_* through Grow and its preceding productive epoch's density give actual canonical material in a band of width C $A + O(A)$ near this edge at the earlier comparison dates.
2. A proper-dead or incoming-clear live U has a known positive successor U'. Disturbance inheritance gives U' incoming-clear. Only if U is live does one additionally have direct current-time clearance of U. U' supplies density of U's actual productive period from age $r + \lceil v/20 \rceil d_0$ onwards; the simultaneous full stage-0 seed supplies the near wall at the start of this search.

Take w large enough that $v/5 > (10C + 500)A$, where C is the fixed density constant. A nominal blocker on a lower displaced tile is converted to a served native predecessor by its nonzero nominal issue and the native no-spontaneous-birth rule. If a native increment at the boundary crosses a ballot, use the incumbent at the pre-ballot row. Radius one in the foreign grid, the pad $64A/P < 1$ and two displaced-grid roundings cost at most four canonical rows and four sites each way. Use the following argument for the served observations with that extra buffer. Any physical wrong-age occupant in the opposite neighbor of a proposed vacancy must be considered a served contender as well.

Early, raw, and shifted-cut observations. Delete the one common physical date interval running from four ticks before the common raw rectangle S to $w^6 + d + 12$ past its last date, with an additional four-row halo. This covers all actual raw-only new families at a late age and their bounded source repair. At a foreign vote at b an older perturbation may give a delayed small source in K. If its new age is $\leq 5q + 2d{+}12$ then trace the particular served occupant backward to the genuine final incumbent of $[b - T, b)$, unless this meets the deleted raw dates. Each step is a prevailing member or a birth from an adjacent served donor; across the vote a vacancy cannot give a live ballot. In these $\leq 5q + 2d{+}12$ ticks there are ≤ 4 native recruitment dates for the fixed foreign phase since $5q + 2d + 12 < 4(2q + 2)$. Consequently the ancestor lies within four sites of the query. The old epoch meets $[0, T)$ because $0 < b < T$ ($t \geq s_0 - 4$ and $s_0 \gg 5q$). Its terminal-age incumbent is considered below. After its full initialized Purge the independent small source has disappeared. If instead a genuine current source is responsible and age $a \leq w^6 + 6q + 20$, its authentic frontier could advance beyond its ordinary home by no more than $(B + 1)(ceil((w^6 + 6q + 20)/d_0) + 2)$ plus $d + 8$ for a short forged group joining it. This counts every provisional birth as permanent and every true attempt once per cycle, and includes partial cycles. As $d_0 \geq 4w^5$ the bound is $\leq 32A + 400q$ for $w \geq 100000$. In Z with the four-site buffer this forces the candidate's positive home to intersect the canonical $64A$ pad. Its start $b = t - a$ satisfies $0 < b < T$ and its epoch meets both $[0, T)$ and $[T, 2T)$. This contradicts the pertinent incoming-clear anchor in cases 1 and 2. This treats each early phase at its actual time without deleting its individual delayed ballot time from the prospective common window.

Terminal age means $a \geq r + .7v\ d_0$, before the candidate's ballot; permit the previous epoch obtained in the shifted-cut case. Its period begins at $b \leq t$ and the query date, including its four-site ancestry, exceeds $s_0 - 5q - 2d - 20$. For a live spatial home use its core; for a dead spatial home apply the remote-marker result, yielding a live home within $300B$, actual wide coverage at every retained comparison date in $b + r + [.05v, .30v]d_0$, or the raw-only transient already removed. The bounds stated here use the nearest edge of the true home and permit a fixed site of its held core within $300B + 12A + 4q + 4$ of that edge, outside the bounded wound on retained dates.

When U is live and incoming-clear, a positive foreign home within $300B$ of Z (including the observed spatial home if live) directly violates its pad since $15A + 300B + 10 < 64A$. Its period still intersects $[0, T)$. Otherwise the foreign wide band contains canonical positions near the right end. The earliest possible comparison is at least $s_0 - 5q - 2d - 20 - T + r + .05v\ d_0$. In the preceding epoch this is later than $-T + r + .05v\ d_0$, so the density of U's positive preceding issue is available, or, if in $[0, s_0]$, the held core of U is available. Between s_0 and the starting ordinary strip use the known U' density if U' is assumed. For the current-input service branch the near-wall bootstrap instead supplies an actual member in $[P - 4A, P - A]$ up to the first-loss test and the wall retained over the pending search supplies it on later dates. With a deficit of at most C A one obtains an occupied site in the portion of the foreign $.2v$ band inside U by counting $> C\ A + 2$ whole canonical sites. At a held-core or near-wall date there is such an occupied site directly. Remove the at most two earlier canonical source-repair date sets (issuing density itself has no excepted dates) and the candidate's wide date set from the order-$.25v\ d_0$ comparison. Their total length is at most $3(220A + 2w + 20)d_0 + 3w^6 + 20$, smaller than $.24v\ d_0$ for the chosen w. Disjoint lower grids cannot occupy the same physical point; two same lower grids requiring different native time/address residues cannot supply both values either.

When U is proper dead and U' positive, select the canonical deep productive late marker in U at distance between $14A$ and $18A$ from P, avoiding K enlarged by $100B$. Its previous-ballot existence is proved in Lemma 5.16's inheritance argument by the positive output's cut coverage. The distance to every positive current canonical home is $> 300B$ since its own home is dead. Hence there is a canonical wide band in $r + [.05v, .30v]d_0$. Consider a foreign terminal contender.

1. If its alternative is wide and $b \geq 0$, the foreign comparison dates fall after the beginning of the canonical issuing-density interval for U'. They also fall before the observation by at least $(.7 - .3)v\ d_0$, so before T. Its band reaches more than C A inside U. The same common-date count contradicts exclusion.
2. If its alternative is wide and $b < 0$, the canonical wide comparison is later than $b + r + .05v\ d_0$ and earlier than the observed date, which exceeds $s_0 - 5q - 2d - 20$. At each such canonical comparison date the whole-tribe foreign terminal thin lemma gives an occupied point within $100B$ of the observed coordinate. This thin implication covers the final withdrawal as well as the following ordinary part and remains valid across normalization by the sterile-band proof for the whole tribe. The marker's distance to the canonical deep center is $\leq 33A + 100B + 12 < v/5$, contradicting the canonical band.
3. A nearby foreign live home with $b > 0$ overlaps U' and the $64A$ pad, directly contradicting U' clearance.
4. For a nearby foreign live home with $b \leq 0$, the canonical comparison dates are after b and before the foreign observation. Until its own normalization $b + s_0$ use a point of its old held core at distance $\leq 300B + 12A + 4q + 4$ from the observed coordinate. If the observation is later than that normalization, on its ordinary prefix from $b + s_0$ onward use the whole-tribe thin trace of the observation itself, within $100B$. If the observation is during Grow, use just the held core. Thus on each retained canonical comparison date there is a foreign point at most $300B + 45A + 4q + 20 < v/5$ from the deep center, again contradicting exclusion. At $b = 0$ no overlap of successive half-open parent epochs was assumed.

In every case the earlier witness date may depend on this specific competing origin. The conclusion holds for every terminal or early contender at every prospective candidate time outside the *single physical* deletion, and requires no union of historical foreign bad dates over all clocks. A genuine remaining contender must have age strictly between the early cutoff and $r + .7v\ d_0$ and epoch start strictly between zero and T because $t \geq s_0 - O(q) = r + v$

$d_0 - O(q)$. It is an attached positive Grow source away from its early/delayed and actual repair. Its live home cannot touch the canonical $64A$ pad. At the right end it must approach from the right: a left home outside the pad would have a connected arm passing through the old canonical middle wall (whose independent strip selection is given in Lemma 5.20), forbidden by lower exclusion or incompatible physical native residues. This is the input to Lemma 5.20. □

Lemma 5.20 (A common clear recovery interval)**.** *At a true end, let the request be the first date of a one-ended ordinary recovery search: an existing end deficit may grow at* Shrink(0)*, an interior edge position is missing, or the proper-input bootstrap has its initial shallow shortage. Require its date in* $[s_0, s_1 + 20vd_0]$*. Assume source attachment, protected lower exclusion, ordinary precommit, the preceding strong issuing cut, the terminal exclusions of Lemma 5.19, and the old near wall and interior support for its proper-current-input or known-positive-successor anchor. Then either the first full ordinary Heal after the request whose complete pass and four prior rows avoid the common raw/repair deletion of Lemma 5.19 is clear and is completed within* $w^6 + 2d + 4d_0 + 50$ *ticks of the request, or an interval of length at least* $0.05vd_0$ *starting within* $3vd_0 + 2w^6 + 12d_0$ *of the request is simultaneously clear throughout the* $100q$ *end strip and its stencil halos.*

Proof. The hypotheses give a fixed near wall and an interior support, using the two alternatives for the current input and the known positive successor, respectively. The distance to an actual donor when the later pass starts is a separate premise for filling: Lemma 5.21 supplies an enduring $q + 1$ guard within depth $97q$. An initially shallower missing strip does not by itself give that later bound. All rank-level clearances, including the old-row halo, refer to original protected physical rectangles.

Walls independent of successful births. For the right end work on $Z = [P - 6A, P + 15A]$ with the four-site/row guard of that lemma; the narrow tested strip for actual births and ancestry lies inside $[P - 100q - 20, P + 100q + 20]$. At the ordinary full seed, or in the proper-input bootstrap before excessive debt, the region $[P - 4A, P - A]$ has old proper ordinary pieces on both sides of a single possible short missing region. An authentic initial local wound and a later K use that same spatial column; any initial edge shortage $< 60q$ is outside this region. Remove K, the possible single $2q$ initial support exception with its neighbors, and, if using the temporary expanded wound before its clearance, its band of length $\leq 4B + 80q$+80+d. The three-A region has at most four remaining pieces and retains more than $2A$ sites, hence an intact piece of length $> .1A$. Trim the exterior two sites to ensure proper old incumbents. If one piece boundary is an active probe rather than an old site, take its old interior before this event by the core/old-run bounds from Lemma 5.12 and the ordinary seed. Old proper, unmarked ordinary positions survive from one true Shrink to the next by the $q + 1$ strong-run calculation: their rightmost Cn retains q predecessor support and each earlier site has a right neighbor. One Shrink can remove at most one layer at each end. During $3v$ $d_0 + 2w^6 + 12d_0$ there are at most $ceil((3v + 2w^6/d_0 + 12)/q) + 1 < 4q + 4$ actual Shrinks, since $v/q = q$ and $d_0 \geq 4w^5$. Taking the original subrun trimmed by $4q + 6$ each way leaves a nonempty fixed wall in $[P - 4A, P - A]$. Source continuation and lower-rank properness apply to that wall throughout the prospective search. A canonical middle barrier in $[P/4, 3P/4]$ is obtained by the identical selection and trimming on that much wider interval. Both barriers therefore exist without already having a successful one-sided repair.

Skip the common physical raw/repair interval of the preceding lemma if necessary and test the next full ordinary Heal together with its four prior rows. At most $w^6 + 2d + 4d_0 + 50$ ticks are needed for this first trial even if the one deleted interval cuts the first candidate pass. If it is clear this is the immediate calendar alternative. Given a surviving in-home guard within depth $97q$ on this pass, each step to the true edge is mandatory and needs at most $97q < B$

opportunities, as proved at the end below. Otherwise take an actual incompatible served witness to failure at a particular query time $t0$, including the predecessor of a conflicting nominal test. All four-halo enlargements are already in Z. By Lemma 5.19 this witness belongs to a genuine active Grow source Q, whose home lies strictly to the right of the $64A$ pad. Its age exceeds the early cutoff and precedes the terminal threshold. It is attached to an enduring core on this clean date. In particular its interval from the tip to the core fills $P + [.25P, .65P]$, up to a four-child rounding. This follows whether its home starts inside this strip (its core's far edge is beyond $P + 64A + P - 12A - 4q$) or farther to the right (its authentic arm must traverse the whole strip).

Let $L = ceil(3v\ d_0 + 2w^6 + 12d_0) + 20$. Since $P = 5{\cdot}2^w$ and L is polynomial in w, choose w so $10L + d + 100 < .01P$. Remove K enlarged by $L + 4$ to both sides from $P + [.25P, .65P]$, and trim the ends by $L + 4$. The remaining length exceeds $.35P$ in at most two intervals, one of which contains an interval of length $\geq .1P$. In its interior retain an interval D of width $.05P$ so every point of D during the next L ticks has its backward cone to the witness date in the original occupied region avoiding K. The date of the recovery below is before Q's own next normalization and ballot. At each Grow update in this cone, a strict interior of three contiguous members of Q is proper, has exact grades, and prescribes a positive hold. This is deterministic also while the source is interrupted elsewhere. Induct on ticks, shrinking the protected occupied witness by one site at each end. Since the cone intervals were chosen with margin L, the whole of D is continually occupied by the same genuine exact label. This is the rear wall of Q.

Universal exclusion over a fixed horizon. Fix an arbitrary later date u in the selected portion of Q's current withdrawal (specified next), excluding the common physical slab and a four-row guard. Inspect any actual incompatible member in Z. At u apply Lemma 5.19 anew. It rules out the tentative member if it is early, terminal, or a pure postcut residue. A residue at someone else's fresh vote is judged by its actual past incumbent; the historical comparison dates used to exclude that incumbent may vary, but this does not change u. In the remaining case it belongs to a connected active source with a positive home disjoint from the canonical pad, as at $t0$. A left home would have a connected interval crossing the canonical middle wall. A right home must reach past D's right end from Z for the same spatial reason that Q supplied the original D. Thus its interval contains D at this very date u. Different lower grids violate lower exclusion there. On a shared lower grid, at a given index (x,u) the actual lower state has a single age and residue, so a second different upper phase cannot match both. If both lie on the same phase and grid, its exact extended address at x in D fixes the source base $x - \pi$; it is Q's base, so the second source is Q. A contender already inside the wall when certified, or one born on its near side later, is subject to the same argument at u: the new age/old-vote cases have already been eliminated, a still-unattached raw lineage requires the removed common slab, and every remaining member must be connected to its own far core. No transport of a distinct source THROUGH the wall between $t0$ and u is assumed. This proves the universal statement: in the single chosen clean portion, for each u and every conflicting member in Z (and its specified halo), the only possible source is Q.

Retraction and uninterrupted window. If Q is in attack, its next withdrawal begins within $2v\ d_0 + d_0$. If in withdrawal with at least $.3v\ d_0$ available, begin now, allowing one partial strobe. If in a nonfinal withdrawal with less remaining, the next full one begins within $2.3v\ d_0 + 2d_0$. This last case cannot occur in the final withdrawal because the source would already be terminal. In all cases restrict to the first $.3v\ d_0$ of this available portion and remain inside Q's same Grow. Delete the physical raw/recovery interval(s), with integer four-row margins and round to complete strobes. Even allowing two deleted intervals of total length $2w^6 + 4d + 100$, at least one of at most three pieces has length greater than $.09v\ d_0 + 2d_0$, since $(.3v\ d_0 - 2w^6 - 4d - 100)/3 > .09v\ d_0 + 2d_0$. Choose its first complete strobe and retain $.09v$

d_0 of clean Q withdrawal (drop endpoint fractions with an additional d_0 guard if needed).

At the selected strobe Q's left tip is to the right of the retained near wall, since otherwise Q's connected interval would occupy a position of that wall. For a tip between $P-4A-4$ and $P+12A+4$, its address is negative because its home starts beyond $P+64A$. If a neighbor on its left has positive grade, it either has the immediately preceding integer address, in which case it is another member of Q contrary to extremality, or its address exceeds $P/2$ by the grade-one definition. In the latter case its own base is less than $P/2+12A+4$. If its positive home meets the canonical pad it violates the selected anchor (same active clock $b>0$). If the home is wholly left of that pad its authentic arm must pass through the canonical middle wall. In particular an honest left neighbor from a second phase cannot pin Q; positive grade requires the very same clock and residue alignment. A raw-only supporting label could last only to its next full Purge and is absent in this chosen clean portion. Therefore at each true Grow there is a zero face at its leftmost extended site or the tip is already clear of this entire strip. The B attempts of an unjoined provisional advance were marked and fully withdraw before this true Grow, as long as the entered front is in this guard, where no true-end latch exists. A site at most B beyond the guard cannot evade the argument by carrying an untested tip through the guard: reaching the narrow tested strip would first require traversing the additional gap of order A outside it. The first $\lceil .02v \rceil$ clean complete withdrawals exceed $16A+B+15$. They carry the tip past $P+12A$, or find it already beyond. Choose a $4B+4$ virgin band near $P+8A$ inside this now empty margin, avoiding all true endpoint residues of Q's tribe. Such a band exists, for instance choose within $[P+7A, P+9A]$, since its length $2A$ is $< P-2$ and the possible paired true ends remove at most two places and split it into at most two nontrivial intervals. This band contains no competing matching source: the universal exclusion just proved applies to any Q-tribe member in Z that could positively pin/repopulate the narrow strip, since that tribe differs from the canonical one. The whole-tribe sterile-band argument prevents later return through its center at a strobe or fine tick during this withdrawal.

There remain at least $.06v$ d_0 clean ticks, which contain a contiguous interval of length $.05v$ d_0 after deleting four-unit rounding and four preceding halo rows. At all its dates every competing served member is excluded by the universal assertion and Q's retraction. A conflicting nonproper-dead nominal tile at any needed prior stencil would have an actual conflicting served predecessor within the already cleared halo, impossible. On the preceding row each of the nine ordinary lower cells was proper-dead or live incoming-clear; lower proper-successor clause makes its successor proper. The initial search + $2.3v$ wait + $.3v$ chosen portion totals $< 3v$ $d_0 + 2w^6 + 12d_0$ even after the stated rounding. **Conditional filling remark.** The current nine cells are incoming-clear and their true central births are consequently mandatory. A same canonical clock/residue neighbor has the unique actual ordinary address at its position by off-Grow normalization. At the actual outer donor-age test it has the correct age if occupied. Across the true wrap its out-of-home ordinary proposal is illegal, and any legitimate in-home proposal installs the intended matching address. A $q+1$ old backing persists through every intermediate Purge, whether advancing to the left or the right. Each of the B attempts recruits the next vacancy or joins an already properly aged in-home incumbent, so the distance from the enduring guard to the true endpoint decreases at least one per successful attempt. With the separate $97q$ guard supplied by Lemma 5.21 this repair takes at most $97q < B$ attempts, irrespective of the end depth at the earlier first-loss request. At its own true endpoint the new trial sees a positive inward face and the outward directional exemption; it persists at retirement.

For a left-end recovery all temporal and home comparisons use $Z = [-15A, 6A]$. A right-side hostile root would have to cross the canonical middle. The actual attacker lies left of $-64A$ and has positive extended rightmost address $\geq$P. A protecting right neighbor of a

different exact base must have address $< P/2$ at the high-to-low wrap, so its home meets the pad or its arm crosses the canonical middle from the right. Select the permanent rear wall as the reflected subinterval of width $.05P$ in $[-.65P, -.25P]$, with the same collar avoiding K; the near wall lies in $[A, 4A]$. Replace the empty band location by $[-9A, -7A]$ and require the tip pass $-12A$. The same distances and withdrawal counts apply. In the Purge and one-sided fill arguments continue to use the rightmost incumbent's q strong predecessors, exactly as in the right-end proof. □

Lemma 5.21 (Current service and a proper positive successor)**.** *In a protected rank-(j) work period with the lower-rank hypotheses and the previously issued cut, an incoming-clear proper-live input is served. If a proper-dead or incoming-clear live input has a positive issue, assume its work is computing with the full seed and that the strong positive cut is available. The positive nominal successor is proper live. The two conclusions use the respective end-relief anchors, and the first uses no hypothesis on the successor.*

Proof. Use parent home $[0, P)$, write depth $k = x$ at its left end and $k = P - 1 - x$ at its right end. A "missing end" below concerns the retired genuine front. Provisional B-long probes are never used as donor anchors. Inputs for this deduction are Lemma 5.18 (applied to the rank-$(j-1)$ children), the all-shift lower organization clauses, actual previous issuing cut, Lemma 5.12 for genuine source intervals and postpulse holes, the precommit common full seed of Lemma 5.9 in the known-computing case, the buffered two-wall repair for previously occupied interior regions in Lemma 5.6, and the two distinct relief anchors in Lemma 5.19/Lemma 5.20. The latter supplies its common calendar for births testing sites of depth up to $100q$, including prior stencil rows. Write $W = 3v\ d_0 + 2w^6 + 12d_0$ for one complete search-and-fill time. These are conditional induction inputs; in particular a positive successor is not assumed in proving service of a proper current input.

Starting material, and two different depths For current-input service start at the genuine issued cut of a proper live parent. In $[1, P-2]$ all positions except a span H of length $\leq 2q$ match. Outside K their surviving true addresses are the ordinary addresses and their scratch is the specified positive reset. A run avoiding H and K and their one-site collars has proper lower incumbents by the matching triple. A proper unmarked run of length $\geq q+1$ survives every Grow Purge against arbitrary exterior by its rightmost site's q strong predecessors; all other sites have a strong right neighbor. Incumbent ordinary addresses are exempt from true Grow withdrawal and no incumbent acquires a Temp mark. Lower proper successor and service turn each genuine positive transition into a proper served next row. Thus such initial intact strips are available for the first ordinary-work search at s_0. An old true edge itself need not have been proper at the initial cut. A real earlier gap well inside this near-full input has the preceding genuine issuing geometry: if no raw has yet occurred in the buffered history, the previous ordinary interval and ballot cannot create such an internal gap. The ballot has no births and all strict interior proper occupants of the single interval take its positive decision. If a cut-induced defect actually exists in the interior, it is in the single K wound, which Lemma 5.12 repairs in Grow before normalization unless the actual event is too late. This is where free formatting alone would be insufficient.

For proper successor start instead at the genuine $5d_0$ fully occupied, clean Shrink(0) seed of a *computing* parent whose next issued state is known positive. Choose a Temp-reset instruction in this window and the following clean full Purge. Every strict interior of the home was proper before and after this reset, by lower triples and lower proper-successor clause. Both true endpoints may be improper. Here H is empty and the argument starts at this seed rather than at s_0. Backward disturbance inheritance (Lemma 5.16) gives incoming

clearance of the positive successor when the input was proper dead or undisturbed live, as required by the second relief case.

A search of length W together with waiting at most two extra ordinary rounds for a complete reset has fewer than $4q + 4$ native Shrink inspections: the spacing is qd_0 and $(W + 2d_0)/(qd_0) \le 3q + (w/2 + 14)/q$ since $d_0 \ge 4w^5$. Two such searches and the short pre/postpulse interior cleanup ($\le 2w^6 + 4d_0 + 20$ fine steps) cost at most $n = 8q + 12$ inspections, counting a partial inspection at each of the three interval ends. On a certified connected ordinary interval, an arbitrarily long *intervening* stretch costs no interior losses: at a true restored end the outward face is exempt and every other face through the attached core has positive grade. The distinct reason that at most two searches can intervene is verified below.

For any contiguous old candidate strip of $L = q+1+2n = 17q+25$ proper unmarked sites missing K, each of these n inspections can remove at most one site from each end. Between inspections the $q + 1$-residue proof prevents all native Purge losses, and Temp never acts on those old sites. Induction through the inspections leaves its central $q+1$ as a fixed old backing. If the time interval starts with a partial Purge the old strip is present through that sweep's clean prefix; if crossing normalization start afresh at its actual first initialization. When only one prospective search and K are relevant, the smaller $L1 = q+1+2(4q+6) = 9q+13$ works. In material starting at depth $p + 1$ two such strips separated by $d + 3$ allow a K-free one with closest fixed point of depth at most $p+1+L1+d+3+4q+6 = p+13q+d+23 \le p+14q+d+22$. An H and K hazard together use three candidates separated by $2q + d + 5$.

For the two-search bootstrap select three length-L strips starting at depth 4 and separated by $2q + d + 5$ unused sites. The one-cell-expanded H hits at most one, and the one-cell-expanded K hits at most one. The closest fixed point of a usable third candidate has depth at most $4 + 2(17q + 25 + 2q + d + 5) + (8q + 12) = 46q + 2d + 76 < 47q$. Every site of the selected strip avoids the true endpoints. The same selection at the other end works simultaneously. Before either end's first excessive loss, these retained witnesses prevent any retired deficit from exceeding $47q < 60q$ in the proposed one- or two-search interval, even if no birth worked during its search. This justifies the $< 60q$ depth invocation prospectively from real old material, rather than assuming an arbitrary $60q$ deficit cannot increase at all during W.

A second selection separates the endpoint used for the *relief calendar* from the farther guard that will actually fill it. If K meets depths $0..58q$, its farthest affected depth is $< 59q$. Take two length-L strips starting at $60q$ and separated by $2q + 5$. They both avoid K, one avoids H enlarged by one on each side, and both end before $97q$. A fixed $q+1$ guard therefore lies beyond K and depth $60q$ but before depth $97q$. If K first meets this end beyond $58q$, use two such strips starting at depth 3, separated by $2q+5$. Both end before $37q$, and one avoids H. In this distant-K case the genuine interior around K is handled by the next paragraph. With no K use either selection. All selected runs have room for lower stencil halos in the $100q$ zone. During the clean calendar each surviving chosen guard can grow all the way to the true end in $< 97q < B$ attempts. Before each attempt there is a complete initialized Purge; an additional correctly aligned occupied piece surviving it is backed, and each newly occupied point attaches by an exact ordinary bond. All queried lower inputs were incoming-clear on the preparatory row and are proper by the rank-$(j - 1)$ proper-successor and service clauses of Proposition 5.3(ii). Every legal in-home birth is consequently realized. At completion the new outward face is exempt, its inward face joins the guard, and every interior bond is positive. This produces every position in the interval from the true end through the fixed deep guard. It does not count an isolated provisional extension into an open interior cavity as a secured end.

Why an interior wound costs no extra search The near-K case just considered fills past the affected column to an old ordinary guard beyond $60q$. If K is farther than $58q$ from both ends, the near-end bootstrap witnesses lie at depths $< 47q$. Thus there are $> 11q$ sites between such an extreme witness and K. On the last actual full inner region before the pulse, the two sides of K each have a $q + 9$ backing within $q + 15$ of its corresponding edge; take their nearest fixed interior $q + 1$ sites after trimming three on each face. The sites in these deep guards (more than $58q - q - 15$ from either competing end) were in the initially near-full proper interior, or in the genuine full seed. If K is a new later pulse, no previous unrelated deep interior hole exists; if K gave a delayed cut, Lemma 5.12 supplies its real previously occupied bracketing geometry. The near-end deficit $< 47q$ cannot have reached either guard. An unhit ordinary part between the previous edge and this raw location has no mechanism for a first clean interior split (positive bonds, lower properness, and a lower hold).

Across an actual pulse these guards bound the previously occupied internal wound. They lose at most two faces to scheduled Shrinks before a complete initialized ordinary repair; the elapsed time is $< w^6$, and $w^6/(q\ d_0) \le 1/(4w)$. In the last Grow/first ordinary seam allow one final withdrawal on extended neighbors. The ordinary old positions of these guards cannot be removed by it, and normalization preserves their strict bonds. The guarded cavity has length $\le d + 20$ to its facing old supports. Lemma 5.6's two-wall clearance provides the nine simultaneous stencils after the common alien reset and its ancestry rows; remaining true in-home bids close that cavity in a clean B-pass. If a raw wound was still open across normalization, the same fixed ordinary guards were present before the pulse on the Grow side; reset Cn on ordinary entry and run that same repair. If the initial defect was a prior cut far earlier, the established Grow repair has already completed. A wound within $58q$ of only the other end belongs to that other end's ledger, separated from the present end by $P - 116q - d$. This covers the full interval between the two protected end regions. A distant pulse therefore does not leave a permanent inner break from which a restored near rim could erode. Only the bounded cleanup interval is charged to its neighboring old strip. The local argument uses actual pre-event occupancy of the inner region; it is not a statement about every interval modulo an arbitrary spatial error mask.

How many waits an end can incur On clean ordinary work before its first effective repair, a previously improper original true endpoint can disappear freely one time. If it never disappears and has continuous, properly backed inner material, it causes no lost interior. Otherwise request a relief pass when the existing end deficit could grow at Shrink(0) or when an already missing interior edge position must be restored. At the outset it is also legitimate to restore the old shallow pre-Shrink shortage in the proper-input case. During the requested wait count native losses against the old backing above. The age/deficit and the prospective near/middle walls of Lemma 5.20 are available up to the first claimed excessive loss. That lemma supplies the complete common clear Heal. A repair is considered secured when the end is actually continuously occupied into its old ordinary material and all these preexisting sites are proper or newly born proper. If a newly filled site is still marked, its complete positive inward bonds and the true outward exemption preclude clean retreat. One more clean Temp reset is allowed inside the long clear interval. In the immediate-clear alternative, if the pulse intervenes before that reset, include that case in the interrupted first search below.

From a secured rim to the next actual pulse, no second free disappearance is possible at this end. Its old lower cells were proper on the preparing clear row or were proper births, so a hold remains served and proper by Lemma 5.18. Its outward face has zero depth, its interior is strongly joined through the old backing, and a clean Purge cannot propagate a zero through that backed run. A previously shallow frontage cannot reopen spontaneously. In the near-K case the secured rim extends through the only possible wound to the fixed guard

beyond it. In the distant-K case any actual interior cut is dealt with in its own $\leq w^6$ cleanup; the separating $> 11q$ interval and fixed backing prevent destruction of the near-end source. Beyond these transient regions the occupied interior has no first native split. Simultaneously argue for both ends until the first excessive loss; the opposite extremity cannot traverse the interior without first violating its own $< 60q$ bound.

There is one exceptional physical interval. If it precedes the first search, that search uses a clear pass after the event and there is no further occasion to reopen. If it meets the search, the common calendar skips the event and gives a clear pass, or one can restart once after it if the attempted immediate pass was cut. If it first occurs after a secured repair, only that event can reopen this end: use one new pass after the pulse, with a subsequent initialized check. If it never affects this end, no second search is charged. Thus at most two waits are charged per end, including the case in which they are consecutive. During a healthy intermission the secured rim and its attached interior have no exposed inner face except the bounded distant wound treated above. Consequently the old candidate-strip accounting uses at most the n paid inspections even if the time between the two waits is macroscopically long. In a pre-event intermission all the near/deep candidate strips in the old uninterrupted interior lie on that connected ordinary body. Existing uncharged dates before the very first loss likewise have full or previously certified interior and no exposed face at these strips. This proves the prospective bootstrap: a first $> 60q$ loss cannot occur during one or both waits because the old near witness of depth $< 47q$ is still present; on the secured portions the end is restored. A first deep wound continuing beyond its allowed repair contradicts the distinct buffered-cavity argument. A single temporary raw revisit of K by the final ballot is accounted for below, not counted as a new pre-ballot erosion episode.

Service, final time cuts, and the single common ballot exception In current-input service an end deficit exposed during Shrink(0) has a relevant search starting no later than its last inspection near s_1. Allow $< W$ for that request, another $< W$ for a pulse intervening before its successful completion, and the explicit short alignment/cleanup delays. Their sum is $< 8v\ d_0$ for this w. A newly free original *true endpoint only* after the last Shrink(0) cannot expose another interior position by Shrink(1) or Shrink(2): depth 1 fails their strict $3A$ and $6A$ tests, and the long proper inner interval supplies its own Purge backing. Thus unspent later free losses touch only the two actual endpoint coordinates excluded in properness. If the pulse occurs after the previously needed requests have completed, take that now prepared interior as the pre-event row. In particular when a new event first occurs later than $s_1 + 8v\ d_0$ the earlier work has already established full $[1, P-1)$ using the clean-history case. If it occurred no later than $s_1 + 4v\ d_0$, any last resulting search is complete by $s_1 + 8v\ d_0$ ($4v\ d_0 + W$ plus its cleanup is $< 8v\ d_0$); if it lies in the remaining crossover interval, earlier unhit initial repairs were done by $s_1 + 4v\ d_0$ and the late-event calculation applies. All these initial search dates fall in the horizon $s_1 + 20v\ d_0$ for which the relief comparison was proved. At s_c the proper-current home contains more than $P/2$ matching positions: before the late event its prepared interior has length $P-2$, and under the late event the guard split below loses $< 2q$ positions in that interval (during a deep in-progress repair one can also retain both long pieces of that prior interior apart from K and two possible native Shrink layers on each facing edge). No assumption about that home's future simulated output entered the argument.

For proper successor take the prepared proper $[1, P-1)$ of the full-seed case. Its ultimate positive ballot is known. If the physical pulse was spent during early preparation and fully healed, the only later free losses are endpoints. A delayed ballot effect may still occur inside K. Otherwise split the previously prepared stable old interval $[1, P-1)$ at the late K. During computation each side of length $\geq q+1$ is an unmarked proper ordinary run, so the whole side survives all clean initialized or suitably preinitialized Purges until the ballot. No scheduled

Shrink remains in computation. When the pulse instead occurs in Shrink(1) or Shrink(2) and both sides have $\geq q+40$ positions, the two deep guards have room for the possible two native Shrink losses before buffered repair. They restore the entire interior hole before ballot since computation intervenes for more than w^6 fine ticks (if arbitrarily close to the stage boundary the initial computation segment supplies that time). On the outer faces of the two pieces the end depth ≤ 1 is exempt under both late thresholds. If a piece is shorter than $q+40$, it is between K and a single true edge. Enclose that whole side with K. The other long piece's exposed face towards K has depth $\leq q+40+d+2$ from that same true end, below $3A$; its opposite true face is also exempt. None of its surviving old interior cells can die by later Shrink or Temp. Its rightmost Purge backing and all strict bonds persist. At most one side can be short, since $P-2-d>2(q+40)$. If K misses the home there is no such interior loss.

At the live ballot the strong cut gives the positive decision at every unhit proper preballot incumbent and the correct central payload outside K. Lower proper successor and proper service transfer that live vote, including an incumbent one position inside a missing true endpoint. The same interval I handles occupancy and payload: take K intersected with the home, enlarged by its one possible short adjacent interior piece and at most two discrete boundary positions, for length $\leq d+q+42<2q$. If no short piece remains, take just K; with no exception use the empty set. In particular combining a missing left flank with a separate independent right-flank exception is unnecessary. The two ends excluded in $[1, P-1)$ are outside the support requirement. Together with the previously derived incoming clearance of the positive successor, this gives its properness. A dead successor follows instead from the actual small-group cut and needs no occupancy argument here. □

5.6 Directed advance and initialized descent

A *willing donor* is an immediate proper-live neighbor whose true served state, in the served input triple with the proper-dead target, makes the prescribed rank-j map F_j request a legal live recruitment into that target, with its clock, address and opposite-neighbor tests satisfied. In a left request the proposed address is the donor's address plus one; in a right request it is that address minus one. The nine-input properness and incoming-clear hypotheses are imposed separately at each use.

Lemma 5.22 (Directed clearance and a permanent startup anchor)**.** *For a proper-dead target and a willing proper live neighbor in the nine proper incoming-clear input stencil, the lower hypotheses give the clear preparation rows for directed recruitment outside the common short interruption. The donor has a permanent in-home backing and can secure its startup anchor in at most $6d_0$ clean ticks, even if its original endpoint is absent.*

Proof. Use time measured in child steps and the canonical target $H=[0,P)$, with willing left home $[-P,0)$. A reflected willing donor uses $x \mapsto P-1-x$; Purge still propagates from the physical left, and both proposals, when legal, have the clock needed below. Write $H_0=2q+2$ for the handle length, $d=4w+8$. Choose $w \geq 10^9$. We use the preceding-rank exactness/omission, nominal classification, proper successor, proper service and nine-proper-input birth clauses, under the actual protected stencils including shifts, with the conventions of Lemma 5.4. Assume a single common enlarged exception K×S of both spans $\leq d$ which contains all intersecting failed lower calls after rounding, and the actual previous-cut conclusion that a nonexceptional ordinary vote starts the next Grow with normalized addresses. The nine parent inputs are proper and incoming-clear, including the dead target. No exclusion based solely on vacancy of the target is used.

Directed clearance on the nine-home strip Use $Z=[-2P-8B, 3P+8B]$. The preparation rows omit the first $6q$ rows and the slab $[\inf S-8, \sup S+2L+8]$ of at most

$17q$ rows, with $L = 4q + 12$. Test actual matching occupants of every conflicting child phase or grid, with an additional eight child sites on both sides and four rows of stencil lookback. A genuine Grow occupant belongs to the home recovered by subtracting its extended address. Its distance to that home is $\leq 1.1P + 1$. In ordinary time its current home $[a, a + P)$ accepts no recruits from another ordinary home. At normalization its ancestors can be from $[a + iP, a + (i + 1)P)$, $i = -2, -1, 0, 1, 2$: for $i \leq -3$ the address would be $\geq 3P$ and for $i \geq 3$ it would be $< -2P$. The nearer edge of each of those five intervals is at distance $\leq 2P$ from any x in the current home. All these tests, also after a time-L predecessor trace of speed one with $L = 4q + 12$ defined below, intersect one of the nine parent homes or its $64A$ pad: at worst their nearer edge is $8B + 8 + L + 4$ beyond $-4P$ or $5P$, which is $< 64A$. The four extra units allow the two one-child-unit roundings at each end. Thus if the candidate epoch intersects $[0, T)$, every authentic relevant foreign source in that epoch was proper dead by incoming clearance and classification. All lookups described here take place on the named protected spatial and temporal domains.

A proper-dead source on its normalized starting cut occupies a span $\leq$q/2. No new same-phase member can appear except by a holding predecessor, an old neighboring donor or the event. The first complete native Purge ending after a specified last exceptional fine row finishes within $2H_0$=$4q + 4$ further transitions: wait at most $H_0 - 1$ for its initialization and use $1 + 2q$ transitions for its three stages. Address normalization or an age-zero vote lies on an elementary instruction, never inside that Purge. For a short component with strong breaks at both ends, no same-clock birth occurs during the pass. After h information ticks any surviving Cn=1 requires h immediately consecutive strong predecessors on the coordinates of the old rows. This follows by induction from the propagation rule testing the left bond and old left bit. With at most q original consecutive sites all bits are zero by $h = q$. The rightmost remaining site in each component has a right break and is killed; successive rows erase every survivor in $\leq$q ticks, even with extra omissions. An adjoining different exact owner has no strong link to it.

If the pulse meets the initial clearance, then from the initial row to the first complete Purge after its end there are at most six elementary instructions in this clock: the first four-q wait, d duration, and second four-q wait occupy $\leq 8q+d + 8$ ticks, less than $6H_0$. Honest descendants of the initial span are inside its six-site enlargement on both sides, and those of the raw sites are inside K enlarged by six on both sides. These two regions have combined length at most $q/2+d+24 < q$. Each strong component has at most this length even if the two spanning intervals originated far apart. The initial complete Purge lies strictly inside Grow. If the pulse is later, initial survivors died already; then only the six-site enlargement of K can contain independent descendants on the first complete postpulse Purge. At normalization use the union of the five possible sources, which all died earlier or lie in that same enlarged K. Distinct short groups consequently cannot create one strongly backed union at normalization. If one traces an ordinary occupant back until this normalization, any intervening birth is in its own home; a previous empty home with no raw has only empty subsequent rows. This covers all sources that could have furnished the observed site or its earlier ordinary home.

It follows that at a queried time t a foreign member requires either a pulse within $4q + 8$ prior child units or a cut b with $0 \leq t - b \leq 4q + 8$ (one lower unit of alignment on both boundaries is included). In the latter case trace to $b - 1$ inside the foreign grid. Every new nonexceptional member requires a served predecessor within one step, and the vote cannot create an incumbent at a vacant site. Either this hits S or gives an actual previous terminal occupant in that halo. If $t \geq 6q - 4$ and $\geq$4 rows of lookback are allowed, then $b > 0$, so the preceding epoch also meets the directed-clear work interval. Its age at this previous terminal row is $T - 1$, greater than $L = 4q + 12$. Apply the same preceding five-source calculation there. Only a raw event within L earlier ticks is possible. Consequently in canonical time, discard

the first $6q$ rows and the one interval [inf $S-8$,sup $S+2L+8$], inclusive. Its number of integer rows is $\le d+2L+18<17q$. Outside these slabs there is no conflicting served occupant in Z with the halos just specified. If S is empty the latter slab is empty. The truncation at the upper end of $[0,T)$ in this argument is interior for the dates used, through the service sample s_c.

A conflicting nominal nonproper-dead lower tile is nominal live by lower classification. Its nonzero nominal issue requires a computing predecessor with a genuinely live served donor or incumbent: the transition on three dead arguments returns D and a surviving native continuation or birth inherits the pertinent clock and address residues. Thus a competing nominal blocker has a served predecessor confined to one additional space/time tile. For a shifted cell that meets the pad of a queried canonical child, its near edge can differ by at most one additional child unit for rounding and less than one child unit for the pad $64A/P$. The three predecessor candidates and these shifts fit the spatial/timing halos above. At a remaining birth date require also one previous clear canonical row and its lower stencils, already paid by the four time units. Every canonical lower input at the previous row was either proper-dead or live undisturbed, so lower proper-successor clause makes each of its successors proper, and all nine are incoming-clear at the new row. Lower directed advance therefore supplies each genuine prescribed birth. The same clearance rules out an opposite *served* state of wrong age in the recruitment vacancy test.

A permanent ordinary startup anchor In the last $20q$ donor positions omit the facing endpoint, the first site if needed to avoid a missing neighbor immediately before this interval, the common input exception of length $2q$ and K of length d, and one neighbor on both sides of each deleted piece to ensure matching triples initially. Within the $20q$-window these omissions leave at most three intervals of total length $\ge 18q-d-10$. One has length at least $(18q-d-10)/3 > q+3$. Fix $q+3$ consecutive sites in it. Their initial inputs are proper by lower triples; their genuine survivor fields outside K have the preceding positive ballot's fresh Temp and clock/address reset. Every other clean continuation keeps the prescribed result by lower proper-successor and service. The run endures for all Grow: no true Grow withdrawal acts on an ordinary address, Temp cannot be set on an incumbent, and in every Purge its rightmost site has q predecessors for all propagation ticks while its other sites have a strong right bond. This argument is independent of gaps at either end of the run or clearance at the start of Grow. It is valid for a right donor with the same rightmost-site proof. Its facing site has distance $< 20q$ to the true facing endpoint.

After the initial clear rows use a complete clean Heal starting at its Temp reset. Before the first such Heal, a different true matching source cannot have reached this short rim from a remote home: during at most $6d_0$ fine ticks its total displacement is at most $6d_0 < P-40q$, and a genuine incumbent in the target came only from its initial proper-dead source, cleared by the first complete Purge. If S meets this first startup, discard its $\le 17q$ slab and restart; the new full Heal still finishes by $6d_0$ since two partial/full cycles plus both short slabs use $\le 4d_0+23q+4<6d_0$. A raw-only same-phase source other than the donor dies at its first complete Purge as above (genuine source cores are a distance $P-20q$ away). A forged copy of the donor's own label, if joined to its authentic backing, has the correct structural age/address for this geometric calculation; a fresh reset clears Temp and Cn before growth. A surviving different label in the facing rim could not have q consecutive genuine predecessors there and is removed by the initialized pass. Each of the B bilateral attempts now preserves the connected prefix from the anchor and either finds the next slot already occupied with the right label or recruits it by the nine-site implication. A newly attached site cannot be purged from this prefix: when its growth is rightward the right tip has q strong predecessors, and when leftward the new left tip has a strong right neighbor and the prefix's right edge

still has q predecessors. During attempts no temporary retreat occurs. After $< 20q$ such steps the true endpoint is present. Subsequent temporary retreat can remove the exterior temporary extension, but its first removal from the true endpoint would test either an inward positive bond or the expressly exempt outward face. The same anchor procedure on a genuine opposite donor gives its facing rim. □

Lemma 5.23 (Directed advance through one interruption)**.** *In a protected rank-j stencil of nine proper incoming-clear inputs, under the lower-rank clauses and one common exception $K\times S$ of both spans at most d, let the central target $[0, P)$ be proper dead and an immediate neighbor a willing donor (the left one if both are willing; a right donor is treated by reflection). Then more than $P/2$ positions match at s_c. The pulse may occur at any date, including startup, first contact, the last withdrawal and normalization. The old-site debit is $D = d+4(q+20) < 5q$, with at most six complete Grow cycles charged when free progress must be repaid. Two B-long provisional probes which were never committed progress are excluded from that debit.*

Proof. Let $q = w^2$, $B = 128q$, $v = w^4$, $P = 5{\cdot}2^w$, $d = 4w + 8$, $d_0 = (2w^3 + 4)(2q + 2)$, $w \geq 10^9$. A round means one Main cycle of d_0 fine updates; one Grow cycle has $2v$ attack rounds followed by v withdrawal rounds. There are $N = floor(7P/(5v))$ Grow cycles. On the nine proper incoming-clear inputs assume the startup and directed-clearance statements of Lemma 5.22. Away from the common raw rectangle use preceding-rank proper-successor, service, and nominal-triple. The no-competing-phase clearance and lower nine-site advance give every eligible central birth after the designated clear rows, with the stencil conventions of Lemma 5.4. These are induction hypotheses, not a hypothesis that the injured parent work later computes.

Reference geometry and clean progress Use the permanent anchor of Lemma 5.22 in the willing home $L = -1$. In each potentially opposing genuine right source choose the reflected ordinary anchor avoiding K. The source +1, if genuine, bars +2 and all more remote sources from crossing this anchor. An incumbent never changes its extended owner during Grow; to cross the fixed nonempty anchor a same-phase lineage must at some date be one of its sites, which already has the anchor owner's address. The same argument at the willing anchor bars a rear left source. If +1 is a proper-dead residue its first full Purge removes its at-most-$q/2$ span. Up to $P + 10v$ into the rightward progress coordinate, the range bound permits only +1 or +2 in addition to L: a +3 source has smallest possible position $3P-1.1P = 1.9P > P+10v$. Thus after startup the only genuine labels on the path are L and at most one R from {1,2}. Extend the reference on either side to its persistently supported facing rim when there is an opposing source. The nine-home clearance strip contains these working intervals. In the case R does not reach this region, only L has to be followed. A wrong label due to the subsequent pulse is cleared separately below.

A clean source interval from its anchor is continuous up to its first free front or a meeting. At a free face the B provisional attempts grow at most B sites; their B retreats delete the full unjoined outward extension unless it reaches its own true exempt endpoint. Only a genuine true Grow attack gives the required unit of progress into an empty target beyond the willing true endpoint. It creates an unmarked incumbent. On a withdrawal the permanent L face loses at most one; its remaining long run is backed. Between these instructions a clean unmarked proper strong run of more than q members is not pruned: only its rightmost site can have a right break and it has q strong predecessors in the pure information pass. Every clean small component with strong breaks and length ≤q is erased at the Purge immediately before the elementary instruction. Thus retained detached strong pieces at such a cut have $\geq q + 1$ sites.

For different L,R labels any possible meeting in the target and up to the stated overshoot has a positive grade. At an adjacent L|R junction the extended integer drops by either $2P-1$ or $3P-1$. Both extended values are in $[-1.1P, 2.1P)$, so the left value is $> .9P-2 > P/2$ and the right value $< .1P+2 < P/2$. Their residues are consecutive, with equal clock. Each side has its own q-strong backing. A last-slot tie chooses one donor; the filled position has an exact bond on its donor side and on its other side either another exact bond or precisely this positive junction. During an unbroken clean join the marks cannot withdraw from positive faces, neither true Grow outward face sees zero at the seam, and each separate label withstands Purge. In particular a joined prefix cannot break in its interior. A right-hand terminal point of R at $x < P-1$ has negative extended address relative to either possible right home. Its homeward right face can recruit at true Grow into $x+1$ (still outside that home), and cannot be removed by true Grow in that direction. At $x = P-1$ or slightly farther, a negatively extended or ordinary R end does not withdraw on its right. This argument works even for a backed R piece temporarily cut away from its distant home.

Follow the longest continuous prefix from the left anchor after each completed round. Disregard unjoined provisional tips. Below $P-1$ a free attack extends this prefix by at least one. This follows from the preceding paragraph when its exposed label is R; when it is L its proposed address is outside its home and below its upper clip. If reaching an occupied position, incorporate the whole continuously occupied ordered piece until the next free face. Once an R-ended prefix reaches $P-1$ the target is safe; before a pulse this is an actual connection to the genuine opposing rim. If the prefix remains L, continue up to $S* = P+10v$, inside the legal clip at $1.1P$ ($10v+2 < .1P$). After first attaining $S*$ the next withdrawal can lose at most v, and the next $2v$ attacks restore this loss prior to the next withdrawal. A contact farther out instead pins the attained prefix. The target then has at least $9v$ free-tip reserve or a positive join. Until one of these alternatives, a clean complete Grow cycle gives net progress $\geq$v; the minimum over intermediate committed fronts of that cycle is the initial front, since attack precedes withdrawal. Startup fits in the first Grow cycle. If uninterrupted, by $\lceil 1.05P/v \rceil$ full cycles the target has the established alternatives: $(\lceil 1.05P/v \rceil - 1)v > P+10v+2$ for this w. By $\lceil 1.1P/v \rceil$ there has also been a fresh Temp reset far past first contact within the target. First contact beyond an attained overshoot has no bearing inside the target.

What is lost near a pulse For a pulse before completion of the initial recovery use the startup case below. Otherwise take as reference the latest clean Temp-reset row preceding the pulse, at most d_0 steps before it. The designated instruction sets every incumbent mark to zero and the following fine instruction begins a fresh Purge. At this reference the committed L and possible R intervals run from their ordinary anchors to their free tips, or meet with a positive seam. A first temporary contact made since this reference can have up to B additional marked sites per side. It costs zero previously certified sites to discard both of those tails. If the junction existed at the reference it is already unmarked; use the established-junction case.

After the $\leq 17q$ interruption slab, allow one whole clean Purge and one fresh complete clean bilateral Heal. If it can finish within Grow, at most four true Grow instructions occur between the reference and the completion of its repairs: distance to pulse $\leq d_0$, its slab $\leq 17q$, a complete clean Purge costs $\leq 2(2q+2)$, waiting for a new round costs $\leq d_0$, and that round costs $\leq d_0$. Their sum $< 4d_0$. If it cannot finish before normalization, that normalization is at most $17q + 2(2q+2) + 2d_0$ beyond the pulse. Continue to the first complete clean ordinary Heal. This adds at most two ordinary Shrink inspections before the end of the recovery, since $17q + 2(2q+2) + 3d_0 < qd_0$. Altogether there are fewer than ten true Grow/active-Shrink inspections in either case; allow ten exposed layers *on each face*.

Consider any one contiguous reference stretch of a single label disjoint from K, initially

unmarked with all its sites proper. In the normalization case restrict attention to stretches lying in one future ordinary home, so that its interior does not acquire a normalized weak wrap. If it has length $\geq q + 21$, excise the outer ten sites on each side. The remainder has at least $q + 1$ sites and endures to the clearance and recovery. Indeed, at each true Grow withdrawal or Shrink lose at most one from each exposed face, at all other instructions a reference incumbent is unmarked, and between such losses the rightmost surviving point of any run of length $\geq q + 1$ has q strong predecessors in that Purge's information pass. During killing all sites left of it have a strong right neighbor. At normalization the incoming Purge initializes Cn before its first propagation, and inside this home the consecutive integer relation of the stretch persists. The preceding-rank proper-successor clause transfers every live native hold outside K, including at the two ends. No "one layer per fine tick" loss is needed for a genuinely proper unmarked incumbent. Initial anchors were chosen through actual earlier triples and are handled by the permanent ordinary-run proof independently of this trimming.

Restrict the reference interval of interest either to the region between the two anchors at an established join, to a free committed arm ending at its old front, or, for finishing across normalization, to the previously completely occupied target $[0, P)$. There are at most two reference labels and one excluded K. Each label minus one interval K has at most two pieces; hence at most four. For a short piece of length $< q + 21$ debit its whole length, at most $q + 20$. For a long piece debit its twenty excised margins. The total reference-occupied positions outside the resulting guard runs are at most $D = d + 4(q + 20) < 5q < B - 10$. This counts missing old positions across all affected pieces even when the defect lies next to an old weak join or a new true wrap. It excludes a genuine old vacancy between originally unjoined arms. The bound for a single arm is no larger. In the two-anchor case the invariant anchors themselves may be added to the guard union; both were chosen outside K. At no time can a short detached end of the former permanent progress cost more than the same D, since any extra reference stretch longer than $q + 20$ would contain a retained guard.

After the slab and before measuring fresh progress take the first whole initialized clean Purge. At most 12 hardware elementary instructions can elapse from the first pulse row up to the end of this pass, since $17q + 2(2q + 2) + 4 < 12(2q + 2)$. A completely new matching raw label has its descendants within K enlarged by twelve per side until that pass, a span $\leq d + 24 <$q, unless their strong component contains genuine older sites. Different labels cannot provide strong backing. A surviving component on that Purge must have started with $> q$ sites and, in clean rows where all lower slots are proper, all such sites survive. Thus a surviving exact-label component intersects the complement of the enlarged column and has a witness of its label descended from before the event. The genuine labels in our strip are L and R in that order. This order is inherited by different surviving components: choose their pre-event genuine witnesses. Paths following chosen actual parents cannot cross. Equality at a vertex assigns the same label, and a simultaneous swap of two adjacent occupied positions would require both positions to have been vacant on that old row. For witnesses outside the twelve-site enlargement, their ancestry never intersects a raw assignment. Consequently an R run strictly to the left of a distinct L run cannot both survive this pass. A genuine rear source cannot send a bypassing witness through the permanent facing anchor. Additional honest steps preserve the order by the same argument. The sole lower raw event cannot supply another independent label afterward.

Recovering a free prefix, an old bridge, or an ordinary target In a fresh complete Grow Heal, after its Temp reset and the Purge before its first attempt, surviving unobstructed label components of length $\geq q + 1$ are backed. They stay occupied throughout the B attempts. A new same-label site has an exact bond to its component. Its right tip, if exposed, has q strong predecessors, and every other site has a strong right continuation. When two

components meet, the joined link is exact for equal labels and has the preceding positive grade for L|R; the long separate backings persist. Therefore no old backed component shrinks during attempts. Any gap between consecutive surviving guards in a previously continuous reference region has length at most D. If both sides have the same label all intermediate addresses of that label are legal because the legal addresses form an interval. If sides are L and R, the physical intervals of legal addresses overlap: their origins differ by $2P$ or $3P$ and the legal length is $3.2P$ up to rounding, with overlap $\geq .2P - 2$. If both inward bids were clipped there would be a positive gap between those legal intervals; hence at least one advancing face is legal whenever the gap is nonempty. With foreign phases absent the age test passes, and lower advance supplies a birth. Thus vacancy count falls by at least one per attempt until the guards join, within $D < B$ attempts. Extra correctly ordered backed pieces in the gap only shorten this count. In a reached join every recruited site has positive bonds on both faces and its label has $\geq q + 1$ contiguous backing. It survives the B temporary removals. This argument handles several subgaps in parallel (each has length bounded by the *total* D); it does not assume both sides reclaim their original labels.

For the old free L arm this restores its prefix from the fixed facing anchor to within D of its pre-event tip: apply the preceding gap argument to consecutive surviving guards within its reference interval. At most D positions of the reference interval lie to the right of the last guard in that interval. If its whole new exterior extension was temporary only those old D positions are charged, and the originally occupied ordinary anchor still starts the prefix. For a formerly joined arm apply the argument between the two permanent facing anchors. This restores the entire target. Once the join is stable, on the R-labelled right side the homeward face does not withdraw, and the connected target is protected. For a lone prior overshoot the restored prefix reaches at least $P + 9v - D > P + 8v$; even a whole first remaining withdrawal cannot expose the target. The clean prefix argument then applies to later cycles.

For a pulse meeting the final cycle too late for a whole Grow Heal the prior late geometry already covered the target. Compute its guards *restricted to H* with the bound D above. These survive including normalization; each is in-H, unmarked before any true ordinary loss, of length $\geq q + 1$ after trimming, and their union contains $\geq$P$-$D sites. Thus there is at least one in-home guard. All positions in H to the left of its leftmost guard or right of its rightmost guard have combined length $\leq$D. After a full clean Purge any additional matching fragments are backed, and every occupied position in H then has the unique ordinary local address. During ordinary attempts join adjacent guards by the same argument. The leftmost guard also reaches address zero, and the rightmost guard reaches $P - 1$, each within $D < B$ attempts, by one-sided legal in-home bids. Both outward faces are exempt at their true endpoint. This completes the ordinary in-home repair, permitting several local subgaps with total deficit D. The argument never demands an intact target exactly on the normalization row. If the pulse first occurs in later ordinary work apply this same guard count in H to the previously full ordinary target at the last Temp reset. If the service sample arrives before full healing, the old guards alone supply $\geq$P$-$D$>$P/2. With the event spent, after normalization with its fresh Cn reset or after complete repair, an ordinary full home stays full and proper to s_c by the native full-home argument of Lemma 5.18, using the same mid-period start as in Lemma 5.24. When the sample arrives before the pulse, the same argument applies to the earlier full ordinary home. No subsequent source advance is required.

Exhaustive timing and progress cases The cases below refer to the first possibly contaminated fine update. A pulse spanning a boundary uses the case covering its latest required cleanup. In every case in which the target is covered at the end of Grow and the pulse, if any, is spent within Grow (no pulse, or a completed Grow repair), it is full at normalization and the full-home argument just cited carries it, full and proper, to s_c.

Interruption cases. Appendix B tabulates the associated time and vacancy bounds. **No pulse, or pulse whose initial/cleanup interval meets startup before the first secured reset.** Use the K-avoiding permanent ordinary anchor(s). Restart the fully clear startup Heal; even a pulse in the last initial recovery round costs $\leq 12d_0$ from the cut, below one $3v$-round cycle. A remote rival cannot arrive in $12d_0 < P - 40q$. Drop this first whole cycle. The remaining $N - 1$ clean cycles suffice to attain contact or $P + 10v$.

Free arms; damage during a $2v$ attack, including an unreset first L/R probe join. Reference taken before the event; debit no pre-reference sites for the two provisional tails, and $\leq$D for its old L tip. All internal old gaps bridge. Exclude the cycle of the reference and any cycle intersecting the repair. At most two Grow cycles touched; on resuming mid-cycle allow an additional v of retreat before the next whole cycle.

Free arms; damage in a v withdrawal, including a provisional join first made there. Same reference and D. At most four true Grow steps before a clean repair; the partially withdrawn reference already counts its old state. The same two-cycle discard and one residual v bound cover an interruption that crosses the attack/withdrawal or whole-cycle boundary.

Previously committed L/R contact, in an attack or withdrawal. The adjacent anchors are restored by one full Grow Heal using $\leq$D missing old sites; the established weak seam's two possible short cut-off fragments were both paid in D. If enough Grow remains for that Heal there is no distance repayment. Otherwise use the final-cycle row below.

Lone overshoot already attained; pulse at the tip, the clip, a remote arm, or the true donor border. Old prefix restored to $P + 9v - D$ or farther, or its genuine joined case applies. The legal clip is $> S*$. One whole v retreat still leaves target covered. Subsequent complete attacks make up clean withdrawals.

Pulse in the final withdrawal or straddling normalization, without room for a complete clean Grow Heal. The undamaged interval before it contains more than $\lceil 1.1P/v \rceil$ full cycles. The target was covered by committed geometry. Restrict old guards to H, permit D missing even on the normalized cut itself, complete the ordinary in-home Heal. At most two early Shrink inspections in the specified recovery window.

Pulse starts after Grow. Old full ordinary H, its D-debit guards or the already observed service sample decide the count. With the event spent, a restored ordinary full home stays full to s_c by Lemma 5.18 (with the mid-period reset used in Lemma 5.24); before healing the guards give $\geq P - D$, and before the pulse the old full home gives P.

For the free-arm rows let c be the complete pre-reference cycles after startup. If the old tip is still unsatisfied its attained displacement is at least cv at the reference, since completed clean cycles give $\geq$v and progress never dips below its starting frontier during the following partial clean cycle. After recovery the displacement is at least cv$-$D. The reference-to-completion interval is $< 4d_0 < 3v\ d_0$, meeting at most two whole cycles even if it crosses their boundary. Until the next cycle boundary one additional withdrawal costs at most v. The initial discarded cycle, the two interrupted cycles and this one additional v use four cycle-equivalents. Allow two further complete cycle-equivalents for the two integer endpoints in counting disjoint pre- and post-reference complete cycles. Thus the eventual free displacement lower bound, until permanent contact or overshoot, is $(N - 6)v - D$. It exceeds $P + 10v + 2$: $N \geq 1.4P/v - 1$, $D < 5q < v$, and $.4P > 18v + 2$. The late threshold used by the final row has spare room since $1.1P/v + 3 < N$ for this w. The omitted and the still-open old macroscopic gap never enter the D count; progress into such a gap is paid only from the available honest attack rounds.

Finally if only the right donor is willing, read all prefix statements from the right. At the encounter the label order in physical left-to-right coordinates is still the smaller-origin label followed by the larger, giving the same positive junction and left-priority resolution. A reflected free right worker has legal decreasing addresses beyond its own left edge until it reaches the reflected overshoot. A homeward end of a left source in the target has positive extended address and true Grow cannot kill it on its left. Each Purge proof still uses the physical rightmost site's q physical left predecessors. The clipping bound is again $\geq 1.1P - 2$. These are the only directional tests used by the argument. Combining the service count with the already derived strong issue and lower proper successor gives the requested central nominal birth. It is served under the ordinary proper-live service clause, once the positive successor has been made proper by the end-repair lemma.

This discharges the induction step using only preceding-rank protected holds, cuts, service and mandatory recruitment. For rank one the child history is the actual radius-one data rule off its flags, with physical service equal to identity. There are no distinct geographic rank-zero grids. Thus each finite rank is treated from its base, with the zero-top-exception variant whenever the entire window is already sparse one level lower. □

Lemma 5.24 (Two periods of served children)**.** *Let owners 2 and 3 be the central pair of homes $k, k+1$, and let h be the rank-j period whose work interval is translated to $[0, T)$ below. Let W_0 be the union, over the six predecessor tiles of owners 0,...,5 in period $h-1$, of the tile's m-collar enlarged by one owner unit, which contains all queries at that tile (Section 5), together with the lower protected footprints of all child tiles of these homes in periods $h-1$, h and $h+1$. It contains every physical query of the two-period argument below, including the preceding seed and work, service lookahead, the protection collars and all protected lower calls, and it lies in $[k-m-3, k+m+5) \times [h-m-2, h+m+1)$ in rank-j owner units. Assume that the deviations restricted to W_0 are $(j-1)$-sparse, that the lower layers are ordinary and unhalted, and that the addressed loads supply the true read-only operands. Use lower coordinates and translate the work interval of period h to $[0, T)$. Number the six consecutive served inputs at its bottom 0,...,5, and assume that they are live and belong to the specified common tribe. Suppose the two prescribed transitions for owners 2 and 3, using these actual served inputs and the fixed environment, are live with the stipulated next tribe residues. Then every lower site in those two homes is served live, with the lower structural residues stipulated by the schedule, namely the clock and address residues of its own work interval, on $[0, 2T)$, hence on the requested $[T/2, 2T)$; their two upper successors at T also serve the prescribed values. The endpoint $2T$ is excluded. If T is odd the left endpoint of the requested integer interval is its ceiling.*

Proof. **Earlier premises.** The rank-j precommit dichotomy and simultaneous full ordinary seed for up to six computing predecessor homes, with J empty, are supplied by Lemma 5.9. A noncomputing predecessor returns nominal D. These are rank-j past-epoch conclusions; neither rank-j proper-live service nor rank-j directed advance is assumed. At rank $j-1$ use its protected exact-computation, nominal-triple, proper-successor and service clauses, including Lemma 5.18, with the conventions of Lemma 5.4. At $j = 1$ the exact unflagged native update replaces the lower clauses. Empty J here follows from the stronger $(j-1)$-sparsity of the entire test; lower-rank sparse defects are allowed. Without the earlier seed lemma this result is a conditional induction implication.

1. **The four full predecessor homes.** The six served positive bottom values imply six positive preceding nominal issues. A noncomputing predecessor issues D, so all six preceding homes computed. Select their simultaneous length-$5d_0$ full stage-0 interval before the first read. Each handle's Purge reset is a legal starting time for the full-home argument; the distance to the next such time is at most $2q + 2 < 5d_0$. At that time

the four inner homes 1,...,4 have all P occupants, and every one of these $4P$ sites has two matching neighbors in the six-home stretch, also across the aligned wraps. The lower nominal-triple rule makes each proper; a served matching triple is nominal matching since a served live value equals its nominal value. The next reset gives Cn=1 simultaneously. During the reset all four homes still have positive neighbors. Now apply the full-home specialization of Lemma 5.18 up to the first ballot at time zero. Its proof requires only ordinary addresses, common genuine clock and a Cn reset for a mid-period start: the Temp-removal clause exempts the outward true face even for an initially marked endpoint and every other face is positively bonded. All $4P$ sites stay proper and served.

2. **All four ballots at zero retain their endpoints.** In each of these four homes the three clean reset evaluations give, separately for each repetition, one packet string on the central output fifth: a serialized live result with every letter live, or all message-empty. This is the total evaluator's property for arbitrary received input including a malformed supplier. Fixed shifts to the owner's five fifths are wholly inside that full home; induction on the shift length transports each intended letter, since every crossed interior bond is strong, every required overwrite occurs and no site is born or dies. Thus every site of that home has the same triple of bank liveness bits (payloads may differ by offset). If at least two bits are live, every site starts verification Active, stays Active and votes live. Otherwise all start inactive with cleared scratch. During verification strict in-owner edges forbid entry from outside, no births occur in a full home, and each epoch starting with no Active bits has zero arrivals everywhere. No counter exceeds $4d$, all sites therefore vote D. Lower exact computation at proper sites makes this nominal result D at every one of the P sites; their served states would all be D, so their parent could not have the assumed positive nominal issue. This excludes the second case individually in each of homes 1,...,4. The first case and lower proper-successor clause now give a proper served reset at every site of those homes, including the two endpoints of each. The argument never supposed that a lower triple survived the vote.
3. **The first work interval** $[0, T)$**.** Starting at zero, the four reset homes stay full and proper to the pre-ballot row by Lemma 5.18. Every old carrier retains its X until the ballot. For both computing centers 2 and 3, all three contributing supplier homes are among those four. For each supplier and fixed offset, the three stored homologous coordinates are spaced P_0 apart. Nominal acceptance at time zero certifies that all three carry the appropriate encoded letter off one interval of at most $2q < P_0$ positions. Hence at least two are correct at every offset. Each inject/shift/receive between these consecutive full homes crosses only strong in-home links or the synchronized positive $P - 1$ to 0 wrap. Induct on unit shifts from injection to the receiving full site; no shift can substitute another source. The majority gives the exact packet at every offset, all three parsers recover the respective served input, and evaluation with the frozen operands gives the prescribed live transition at the two centers in each repetition. Readout inside each full center is exact at every site by step 2, so all sites there start verification Active. They remain Active to the ballot, return their appropriate live letters with zero age and cleared scratch, and lower proper-successor clause followed by lower proper-live service gives proper service of every endpoint across time T. The exterior neighbors at that vote may die.
4. **The second work interval** $[T, 2T)$**.** Apply the reset-home proof to homes 2,3, independently, now without using occupancy of homes 1 and 4. On every pre-ballot structural step each proper site prescribes a live result by the instruction calculation and gets a proper served successor from the lower clauses. Their clocks and address residues are those of their own work interval by the true schedule. Thus both central spatial

homes are full at every date T through $2T-1$. At $T+s_c$ each has exactly P matching occupants, exceeding $P/2$, and at T the full correct payload gives each prescribed nominal successor. They therefore serve those upper values. On the physical grid recursive lookahead obeys $b_0 = 0$ and $b_k \le s_c\, T^{k-1} + b_{k-1} \le (T-1)T^{k-1} + T^{k-1} - 1 = T^k - 1$, so every last lower service observation fits its own tile.

□

Assembly. Lemmas 5.10 and 5.11 give the issue and classification, Lemma 5.17 gives separation and the matching middle, Lemma 5.21 gives service and proper successors, and Lemma 5.23 gives the central requested birth. The dead successor uses the small-group cut. The structural hold and ancestry conclusions apply wherever the child step is protected. At rank one the calculation starts with the native step. This closes the induction in Proposition 5.3.

6 The dynamical alternatives

The initialized nested descent adapts Gács [10, Section 11, Lemmas 11.1–11.2]; the halting cap and noise-selected external operands are used as specified here.

6.1 Local comparison for the external operands

Lemma 6.1 (Locking the read-only parameter)**.** *At a rate in I_n, a protected comparison contained in a good stripe box uses SEARCH through rank n and locks at rank $b = n+1$ on the true initial tape $U(n)$. For each eligible protected test at these ranks, the conclusion of the organization induction holds for its actual data, with deviations covered by E from Section 4.4. Above rank b the read-only transition is the uniform locked transition. If $U(n)$ halts, some first finite rank has a transition identically* D*. If it never halts, all ranks are ordinary.*

Proof. In a frozen reference box the distinguished vector of a rank-j tile with left coordinate x is the base vector at

$$x + u(1 + P + \cdots + P^{j-1}),$$

zero extended. The formula at rank zero is the original vector at x. At the next promotion, the three samples have coordinates obtained from the three child positions $u, u+1, u+2$; after the common lower-rank displacement they are spaced by P^j. When $j < n$, their total span $2P^j$ is less than P^n. There is at most one flip in each good channel. At $j = n$, both adjacent samples differ in every good channel, with no restriction on its phase. Thus the strict-majority test first succeeds at this promotion, with counter n, giving the true lock at $b = n+1$. Minority channels can be arbitrary spatial functions. The reads on the two protected transactions are the actual loads on their continuously served corridors, so the same argument applies to the read-only operands of a represented forged instruction. This use of the lower claims is part of the rank induction: the actual loads consume established lower observations.

To justify the comparison without global phases, take a fixed eligible protected invocation of rank $j \le b$, including its calls, and choose an anchor at its tile corner. If the anchor were bad, its paint would contain a full V^j inside the tested rectangle. By Definition 4.1 and the calculation following it, this is incompatible even with j-sparsity. Choose the fixed D_0 so that D_0^{n+2} dominates the largest required width or height of every clipped domain and all its sum-of-offset lookups through rank b. This is possible since each such bound is a fixed multiple of $\max(P,T)^{n+1}$. Extend the majority stripes by their phases and the static minority functions by arbitrary values outside the good box, all constant in time. On the entire clipped

domain keep the original physical data history. Every local proposal, every required load and every x, y observation then has exactly the external values of this comparison; lower service lookahead stays inside its tile. Both the invocation and its lower calls after one excision in the sparsity argument use the same frozen extension. Only the errors already flagged in E need be charged. For ranks above b, apply the comparison locally wherever a lower invocation through b is needed. The true tape is the same at all such locks, and the external word at higher ranks is zero. Thus no common phase across separated boxes is needed. The last claims follow from the exact true step and halted guard of Section 3.5. □

6.2 An empty cut after halting

In both lemmas of this subsection, K denotes the first halted rank; here it does not denote the exceptional column of Section 5. All layers strictly below K are ordinary and unhalted; no ordinary induction is continued above it. A protected rank-K work rectangle is a rank-K tile with its protected test (Definition 5.2), and its buffered test is that protected test, to which the deviations are restricted as in Section 5.

Lemma 6.2 (All-start cut at a halted rank)**.** *Let $K \geq 1$ be the first halted rank. In every protected rank-K work rectangle whose buffered test is $(K-1)$-sparse, the output cut has no served live child of the tested clock and address group. This holds for arbitrary initial physical rows when the work and its preceding history lie after initialization.*

Proof. Use the strong issue of Lemma 5.10 with an empty rank-K exception. The lower calls still see their actual lower-rank defects. If the tile computes, both unaffected repetitions freshly initialize the evaluator and use the genuine halted literal at the covered loads (Lemma 6.1 for the single rule). The evaluator returns D on every parsed triple, including the case of a forged nonhalted work tape and all malformed inputs. The dead half of the issuing cut confines matching survivors to the exceptional column, which is empty. If the tile fails to compute, the noncomputing clause gives the same empty cut. These calculations assume no form of the starting row. They apply separately on each translated grid. □

Lemma 6.3 (Ascent from a live physical site)**.** *At fixed $K \geq 1$ there are centered tests of fixed multiples of V^{j+1}, $0 \leq j < K$, with multipliers independent of j and K, and a finite common burn-in such that a live physical site after burn-in has at least one failed j-sparsity test centered there. This holds for all physical starts and the tested field of deviations.*

Proof. Assume all tests pass. Begin at rank zero and whenever there is a served live site at rank j, trace it to a parent computation. Its clock and residue determine a parent grid. Let c be the latest cut date of that grid not after the site's date t, and put $a = c - T$ in child time units. For $j = 0$ each update in the trace is unflagged by the rank-zero test and uses the native rule. For $j > 0$, at each protected trace query a served live output equals its live nominal output. Proposition 5.3(iii) thus gives a computing owner and the native transition on the actual served triple, as in Lemma 5.4. The live result has a served incumbent or donor at distance at most one on the preceding row. Its clock and address residues preserve the parent phase. Trace $t - a < 2T$ such steps. The ancestor at date a is a matching child at a parent cut, at spatial displacement less than $2T$ in child units. At that cut addresses have normalized and the canonical ballot did not make any births.

Use the immediately preceding parent work on the unique spatial home of this child. A D issue would leave that child absent by the empty-exception strong cut. At rank K this contradicts Lemma 6.2. At a smaller rank the issue equals the native map on three actual served inputs. That map on three empty inputs gives empty. Choose a served live operand and continue the ascent from its site, rather than from the issue itself.

The displacement from one served site at rank j to the next ascent input is less than $3T\,T^j$ in time and less than $(2T+2P)P^j$ in space. The former pays for the two periods of ancestry and one of work; the latter pays for ancestry and the offsets to the home and its neighboring input. Since $P, T \geq 2$, geometric summation over the ranks $0, \ldots, j$ bounds the total displacement from the physical site to the rank-$(j+1)$ ascent input by $6T^{j+1}$ backward in time and by $(4T/P+4)P^{j+1}$ in either spatial direction. Use these bounds as centered half-extents before padding. They bound the intermediate trace positions and the preceding parent work as well: the bound for each step already includes two periods of ancestry, one period of parent work and the horizontal home and operand offsets. Add $(m+2)P^{j+1}+4P^j$ on each spatial side and $(m+2)T^{j+1}+4T^j$ on each temporal side. The padding includes one parent tile extent and its m-collar enlarged by one parent unit, with the four child rounding units. It therefore contains every protected rank-j trace query and every selected rank-$(j+1)$ work test by the localization bound in Section 5. The rank-j sparsity of the centered rectangle passes to each of these restricted tests. In particular the rank-$(j+1)$ cut has an empty exception. The rectangle contains the original site and date, and its multipliers relative to V^{j+1} depend only on P, T, m. Observations are within the service lookahead established in Section 5. Choose the burn-in greater than the past extent of these finitely many tests, including their required helper innovation prefixes. The observations can include later updates from the actual history. This proves the claim from every start. □

6.3 Exponential forgetting in a halting window

Couple any two starts using common whole-symbol innovations. At rank-zero halting the data proposal is always D, hence it couples in one tick. Otherwise let $K \geq 1$ be the first halted rank (Lemma 6.1 for the single rule; the hardwired literal for the family) and apply Lemma 6.3 to both trajectories starting after the common burn-in. For an output site at time t consider each necessary helper read (x, s) in its finite physical light cone at dates $s \geq t/2$. The individual disagreement bound has exponent at least $ct/2$, by Lemma 4.5. A union bound over the polynomially many sites, times and slots in this part of the cone gives a bound of the form $\mathrm{poly}(t)e^{-ct/2}$. In the family no helpers are needed.

On agreement of these helper reads, any remaining differing data site has a preceding differing data input, since a common replacement erases the discrepancy and a common deterministic input yields the same proposal. Such a chain stays in one logical lane. At every disputed vertex some data value is live, as D is unique. Lemma 6.3 places every vertex beyond burn-in into the same noise-measurable set of failed tests, independently of which trajectory supplies the live site. A chain lasting to the middle has temporal extent at least $t/2 - O(1)$. Lemma 4.6 bounds its probability by $C'e^{-c'(t/2-O(1))}$, which is at most $C''e^{-c't/2}$; the parent steps are nearest or diagonal neighbors in the lane. The helper coupling and the cluster bound use no independence between the two failure events. The fixed finite number of physical slots is covered by a union bound.

Absorb burn-in and polynomial factors into the constants. The resulting one-site total disagreement is at most Ce^{-ct} for every pair of physical starts. A union bound gives $C|F|e^{-ct}$ for finite F. Compare to a start sampled from a stationary law, whose existence follows as in Section 2. This proves (2), including the full physical symbol.

6.4 The initialized hierarchy

Lemma 6.4 (An initialized negative-time comparison)**.** *For a nonhalting input there is an initial physical law with a compatible fully occupied aligned comparison before time zero. On its homogeneous ranks, all ranks for the family and ranks $j \geq b$ for the single rule, the rank-j site i at each rank-j date $s \leq 0$ is live with clock $s \bmod T$ and address $i \bmod P$. In*

the fixed-rate family this comparison has no negative-time faults, and the initial law is the point mass at its time-zero row. For the single rule its deviations on negative times lie in the paint, and its time-zero helper row comes from the stationary autonomous realization. For the single rule the comparison is set in every residue lane of Section 3.6, in that lane's logical coordinates and with the same aligned skeleton. In both cases every physical data site has age 0 *at date* 0*, the time-zero data row is deterministic, and positive-time updates are those of the PCA with fresh innovations.*

Proof. First construct a clean all-positive history on the site-homogeneous ordinary ranks, all ranks for the family and ranks $j \geq b$ for the single rule. At a finite top rank h, take the fixed live padding letter of the packet convention of Section 3.3. At rank-h date -1, that is at physical time $-T^h$, put the aligned live row in which site i has ordinary address $i \bmod P$, age $T-1$, that letter as retained letter, active bit one, three copies of that letter in the output banks and default values in the other fields. Its ballot gives an aligned live reset row at rank-h date 0. Apply the native initialization tick to that row to obtain the aligned live row at date 1. Use the rows at dates $-1, 0, 1$ as the top data. Encode each upper row and unfold each work period between successive upper rows. Lemma 3.1 gives the encoded successor at the next boundary and full occupancy at every intervening date. In each unfolding E_j gives age zero and ordinary address a at block position a, and the clock advances by one per tick within a period. At a fixed homogeneous rank $j < h$ the resulting rows cover physical dates from $-T^h$ to T^h on its rank-j grid. Diagonal compactness over the finite state spaces at individual sites and dates gives a compatible two-sided history on the homogeneous ranks, with clock $s \bmod T$ and address $i \bmod P$ at rank-j site i and date s. At time zero each fixed spatial coordinate eventually has upper coordinate 0 or -1 after iterating $i \mapsto \lfloor i/P \rfloor$, where the next digit is an endpoint. The complete encoded endpoint is independent of the particular live state. Hence the time-zero boundary is determined at each coordinate by finitely many encoding substitutions. In the family the initial law is the point mass at this time-zero row, followed by the PCA with fresh innovations.

For the single rule continue below rank b on negative dates. Impose at every aligned boundary the complete E_j row of the already prescribed upper row. Between boundaries run each ordinary work with its true searching parameter and the environment word from the distinguished representative coordinate and date in the autonomous helper history. Use the stipulated full positive following boundary even where this replaces the last ballot. Before that ballot the ordinary evolution of the full row stays fully occupied. When the parent box is good, its representative operands and the actual addressed loads agree with one frozen environment. The clean identity then implies that any mismatch at the lower forced boundary points to a mismatch between the stipulated upper step and its own true update. Writing D_j for the sites of deviation of the stipulated rank-j history from its read-only rule, this gives the recursion: every point of D_j whose parent is good has a parent discrepancy in D_{j+1}. At the locked rank b of Lemma 6.1 we have $D_b = \varnothing$. Induction backward through these at most b promotions shows that each physical mismatch must have a bad parent anchor in this finite ancestry. Its space and time displacement lie in the fixed-multiple H_n paint by the sum of the bounded lookup displacements through b. At time zero the data row is the aligned skeleton just obtained, with every scratch field default. Its endpoints at successive levels make it independent of the helper phases. Use the time-zero auxiliary row of Lemma 4.5; future innovations are independent of it. Carry out this construction in every residue lane of Section 3.6, in that lane's logical coordinates and with the same aligned skeleton. The time-zero helper row serves every lane at once, and each lane's negative-time comparison is a measurable function of the one helper history. The time-zero data rows of all lanes together form one deterministic physical data row with age zero at every site. Thus the negative-time comparison and the actual future join at an allowed initial law. □

6.5 The nonconvergent clock observable

Lemma 6.5 (Nested tests and the clock witness)**.** *On a nonhalting input, for every physical site x and every integer $t \geq 0$, the initialized law of Lemma 6.4 gives probability greater than $3/4$ to the event that the data field at x is live with age $t \bmod T$ at date t.*

Proof. Fix x and t. Construct deterministic nested tests from the observation upward in the initialized aligned grid of the lane of x; in the family this is the aligned grid of the line. For the time span $[a, b')$ of a smaller rectangle of length less than $T^j/2$, choose $h = \lfloor a/T^j - 1/2 \rfloor$. Then the span is contained in $[(h+1/2)T^j, (h+2)T^j)$. For spatial span $[c, e)$ of length less than P^j, choose $k = \lfloor c/P^j \rfloor$. It fits in the adjacent two rank-j homes $k, k+1$. Put $M_0 = \lfloor 2w/5 \rfloor$. In rank-$j$ owner units take the new test with spatial span $[k - M_0, k+2+M_0)$ and temporal span $[h - M_0, h+2+M_0)$, enlarged by four child units at each end: $4P^{j-1}$ rank-zero sites of the lane at each spatial end and $4T^{j-1}$ ticks at each temporal end. This includes the six inputs around the central pair, the preceding work, the two forward periods and the collar of Proposition 5.3. In particular it contains the region W_0 of Lemma 5.24 for the central pair $k, k+1$ and the work interval in period h, which lies in $[k - m - 3, k + m + 5) \times [h - m - 2, h + m + 1)$ in owner units, because $M_0 \geq 2m \geq m + 3$. Its width and height are less than wP^j, wT^j, also less than $P^{j+1}, T^{j+1}/2$, respectively. Thus each new test admits the next choice of covering homes.

The fixed clock satisfies $T > 8M_0 + 40$, which is equivalent to $(2M_0 + 10)T^{r-1} < T^r/4$ at every rank $r \geq 1$. Choose a sufficiently large finite top rank $r \geq 1$, above the lock when applicable, with $4(t+2) < T^r$. The preceding smaller test contains the observation and has height less than $T^r/4$, so it lies in $(-T^r/4, T^r/2)$ in time; for a one-site starting observation the left endpoint is t. In both cases the left endpoint lies in $[-T^r/2, T^r/2)$, so the floor prescription gives $h = -1$; use the same spatial covering rule. All six top inputs at $-T^r$, even with lookahead at most $T^r - 1$, are read on strictly negative dates. At each aligned child work boundary the comparison row is the exact encoding of the corresponding upper row, and throughout each intervening work interval its child rows have full occupancy and the scheduled residues. Apply Definition 5.1 inductively from rank zero within the negative-time observation cones. The three central packet copies at the bottom of each tile decode the comparison symbol, and its service sample has all P matching children by the lower-rank assertion. Thus the six actual top inputs are the prescribed live comparison symbols in one tribe. At rank r the rule on their triples is the homogeneous comparison rule, so its two prescribed results are the live comparison symbols at rank-r date zero, with the next clock and address residues.

Require $(j-1)$-sparsity of the rank-j test at all these finitely many ranks. Apply Lemma 5.24 to the top test. At a rank $j > 1$ the conclusion supplies the six served live inputs of the next test down and the served live successors of its two central owners, with the stipulated residues. These sites and dates lie in the covered pair of homes and time interval by the nesting construction. Apply Proposition 5.3(iii) at rank $j-1$ to those two owners. Their protected tests lie in the rank-$(j-1)$ test since $M_0 \geq m + 1$. Each served live successor equals its live nominal value, so each owner computes, and that nominal issue is the prescribed transition on its actual served triple. This supplies the live-output premise for the next application of Lemma 5.24. At rank 1 the conclusion gives the root physical site's live value with the initialized age directly. The tests depend only on the coordinates of the observation in its lane's aligned grid, and the descent applies to any homes of that grid. Apply Lemma 4.6 to these deterministic tests. For the fixed-rate family use the raw flags, which dominate all deviations of the comparison. For the single rule require sparsity of the set E of Section 4.4. It carries paint at every date and dominates the deviations of the joined history, including those of the negative-time comparison of Lemma 6.4. At each eligible protected

invocation through rank $b = n + 1$ in the descent and its lower calls, apply Lemma 6.1 with that invocation's protected test and its sparsity hypothesis. Its anchor is good and its addressed loads agree with the local frozen comparison. Above b the locked transition is uniform, and any needed comparisons through b are made locally at their own protected invocations. In both constructions Lemma 4.6 bounds the sum of the failure probabilities of the prescribed tests by a number strictly less than $1/4$, uniformly in their translations. For the single rule this estimate applies in each residue lane by translation invariance of the physical rule and the uniform bounds for the helper events. This proves the assertion. □

For the clopen event that the origin is live of age zero, the probability is greater than $3/4$ at dates congruent to zero modulo T and less than $1/4$ at dates congruent to one. Hence this initial law gives failure of weak convergence. Write ν_t for the row law at date t. For $i \in \{0, \dots, T-1\}$, take the phase means $N^{-1} \sum_{n=0}^{N-1} \nu_{i+nT}$. Along one common subsequence let their weak limits be λ_i. The Feller property and telescoping of the phase means give $\lambda_i P_\varepsilon^T = \lambda_i$ and $\lambda_i P_\varepsilon = \lambda_{(i+1) \bmod T}$. The clopen event separates λ_0 and λ_1. Their full equal phase average is stationary for P_ε. This leaves open multiplicity of stationary row laws for P_ε.

One rule is used for the entire computable sequence of rational windows. The estimates above hold at every $\varepsilon \in I_n$, proving Theorem A. For the family the radius-one table is effectively enumerable from the initial tape on n. The geometry and both raw-flag thresholds in Lemma 4.6 have no word-width dependence. Fix one rational $\varepsilon_* > 0$ below those thresholds, for instance a sufficiently small inverse power of two. The halted rank enters the convergence constants only, and the nonhalting witness uses the same threshold for every family member. This proves Theorem C.

7 The square-lattice interaction

Write f for the rule of Theorem A, r_0 for its physical radius and $N = |\mathcal{A}| \geq 2$ for its full alphabet size. An interaction term is anchored at $v = (x, t) \in \mathbb{Z}^2$. Write $U(v) = (x + [-r_0, r_0] \cap \mathbb{Z}, t - 1)$ for its ordered parent neighborhood, and put $S_v = \{v\} \cup U(v)$. The term on S_v has energy

$$\Phi_v(\omega) = \mathbf{1}_{\{\omega_v \neq f(\omega_{U(v)})\}}.$$

This single finite-range interaction $\Phi = (\Phi_v)_{v \in \mathbb{Z}^2}$ is translation invariant. For each parent pattern the sum of the child weights is the same number

$$c_\beta = 1 + (N-1)e^{-\beta}.$$

In a finite volume the factors touching the volume are indexed by $v : S_v \cap \Lambda \neq \varnothing$. Dividing each by c_β cancels out of the conditional normalization. The resulting specification is

$$\gamma_\Lambda^\tau(\omega) = \frac{1}{Z_\Lambda^\tau} \prod_{v : S_v \cap \Lambda \neq \varnothing} p_\varepsilon\big((\omega\tau)_v \mid (\omega\tau)_{U(v)}\big). \tag{6}$$

Indeed, the normalized weight of a child unequal to the rule's proposal is $e^{-\beta}/c_\beta = \varepsilon/N$ and the normalized weight of a child equal to it is $1/c_\beta = 1 - \varepsilon + \varepsilon/N$, where

$$\varepsilon = \varepsilon(\beta) = \frac{N}{e^\beta + N - 1}. \tag{7}$$

For $I_n = (a_n, b_n)$, enumerate positive rational pairs $\beta_- < \beta_+$ and dovetail increasing precision enclosures for $\varepsilon(\beta_-)$ and $\varepsilon(\beta_+)$. Return $J_n = (\beta_-, \beta_+)$ on the first certified pair with

$$a_n < \varepsilon(\beta_+) < \varepsilon(\beta_-) < b_n.$$

The map (7) is a continuous strictly decreasing bijection from $(0,\infty)$ onto $(0,1)$. The inverse image of each I_n is therefore a nonempty open interval, so the search terminates. Images of the chosen windows lie in the disjoint I_n, hence the J_n are disjoint. The search uses only the rule's alphabet size and the endpoints of I_n.

Subsections 7.1 and 7.2 prove the dichotomy for Φ itself: exactly one DLR measure in each halting window and at least two in each nonhalting window. Subsection 7.3 blocks an interaction by a time period and doubles it, and Subsection 7.4 applies both steps to Φ with the clock period to prove Theorem B.

7.1 Halting windows

We use the following theorem from the preprint [8]. The equation references below refer to the corresponding formulas in this paper; in (6), replace p_ε by the generic strictly positive table p. The source uses the sup norm, while (3) uses graph distance; since $\|z\|_\infty \le \|z\|_1 \le 2\|z\|_\infty$, the source estimate implies our convention after replacing its exponent c by c/2. Its forward implication uses Bérard's triangular coupling, as explained in [8].

Theorem 7.1 ([8, Theorem 3.1]: one-dimensional dynamical–equilibrium equivalence)**.** Let p be a finite-alphabet, finite-radius, strictly positive PCA on $\mathbb{Z}$, and let γ be the specification (6) of its normalized space–time interaction on $\mathbb{Z}^2$. The following are equivalent:

(i) the PCA is uniformly exponentially ergodic, (2);

(ii) γ has arbitrary-volume exponential weak spatial mixing, (3).

Under either condition γ has exactly one DLR state, $\bar\mu$, the two-sided stationary trajectory law of the unique stationary row law μ; there are $C_*, c_* > 0$ such that for every finite Λ, every $E \subset \Lambda$ and every exterior τ,

$$d_{\mathrm{TV}}\big((\gamma_\Lambda^\tau)_E, \bar\mu_E\big) \;\le\; C_* \sum_{v\in E} e^{-c_*\,\mathrm{dist}(v,\Lambda^c)}; \tag{8}$$

and for bounded observables f, g supported on finite sets E, F,

$$\big|\mathrm{Cov}_{\bar\mu}(f,g)\big| \;\le\; C_*\|f\|_\infty\|g\|_\infty\,(|E|+|F|)\,e^{-c_*\,\mathrm{dist}(E,F)}, \tag{9}$$

after enlarging the constants if necessary.

For every β in a halting window, (1) and (7) give a strictly positive finite-radius PCA on the line, and Theorem A gives (2) from every start. Thus (6) has exactly one DLR measure and satisfies (3) by Theorem 7.1. The constants in the weak mixing estimate may depend on β, including through the halting rank.

7.2 Nonhalting windows

Let P denote the PCA kernel at the rate $\varepsilon(\beta)$ for a fixed β in a nonhalting window. Let ν be the initialized law of Lemma 6.4, $T \ge 2$ the clock period and E_s the one-site event that the data field at the origin is live with age $s \in \{0,\dots,T-1\}$. Write θ_x for the spatial shift by x, acting on rows, events and laws, so that $\theta_x E_s$ is the same event at the site x. By Lemma 6.5, for every $x \in \mathbb{Z}$ and every integer $t \ge 0$,

$$(\nu P^t)(\theta_x E_{t \bmod T}) > \tfrac{3}{4}.$$

For $R \ge 1$ consider the joint averages over times and spatial shifts

$$\lambda_R = \frac{1}{R(2R+1)} \sum_{k=0}^{R-1} \sum_{|x|\le R} \theta_x(\nu P^{kT}).$$

The difference $\lambda_R P^T - \lambda_R$ telescopes to $\big(R(2R+1)\big)^{-1} \sum_{|x|\le R} \theta_x(\nu P^{RT} - \nu)$, so its total variation is at most $1/R$. The sum over x in $\theta_1\lambda_R - \lambda_R$ telescopes to two terms, so its total variation is at most $1/(2R+1)$. The kernel commutes with every θ_x. Hence, for $0 \le s < T$ and every y, the number $(\lambda_R P^s)(\theta_y E_s)$ is an average of numbers $(\nu P^{kT+s})(\theta_{y'} E_s)$, each greater than $3/4$. Take a weak subsequential limit ρ of λ_R as $R \to \infty$; it exists by compactness. Since the two discrepancies tend to zero, the Feller property gives $\rho P^T = \rho$, and continuity of the shift gives $\theta_1 \rho = \rho$. The events $\theta_y E_s$ are clopen and P^s is Feller, hence $(\rho P^s)(\theta_y E_s) \ge 3/4$ for every y and every $0 \le s < T$.

For each k, start at date $-kT$ with law ρ and evolve by P. The resulting histories are consistent on their overlapping domains since $\rho P^T = \rho$. Their extension defines a two-sided PCA history with row laws $\rho P^{t \bmod T}$. This history satisfies the DLR equations for (6). To see this, fix a finite region Λ and a finite exterior cylinder. Choose a date strictly before both and a finite collection of sites by tracing their parents back to that date. Include every S_v meeting Λ, then close the collection under parents down to the same date. Condition on the bottom row. The joint law on the finite collection is a product of the update factors, as each new row is conditionally independent given its parents. Conditioning on the other variables of this collection, the factors affected by the assignment in Λ are exactly those in (6). This proves the integrated DLR identity against the exterior cylinder after marginalization, and the monotone class theorem extends it to the full exterior sigma algebra. This is the finite-cone argument underlying [8, Lemma 2.4] and also works with these periodic row laws.

Call this history Q. Since $\theta_1 \rho = \rho$ and P commutes with θ_1, Q is invariant under the translation by $(1, 0)$. Its row laws and kernel are T-periodic in time, so it is invariant under the translation by $(0, T)$.

The one-time translate Q_1 of Q is a DLR measure by translation invariance of Φ, and it has the same two invariances. The original measure gives E_0 probability at least $3/4$ at date zero. The translate with row law ρP at date zero gives it probability at most $1/4$, since $E_0 \cap E_1 = \varnothing$. The two measures are distinct. This argument exhibits the nonstationary alternative behind the general implication in [8, Theorem 9.2 and Corollary 9.3]. It proves the nonhalting alternative for Φ.

7.3 Blocking and doubling

Fix an integer $T \ge 1$. The blocking map $\Pi_T : \mathcal{A}^{\mathbb{Z}^2} \to (\mathcal{A}^T)^{\mathbb{Z}^2}$ is

$$(\Pi_T \omega)_{(x,k)} = \big(\omega_{(x,kT)}, \omega_{(x,kT+1)}, \dots, \omega_{(x,kT+T-1)}\big).$$

The block of the coarse site $u = (x, k)$ is $\{x\} \times [kT, kT + T)$. The map is a homeomorphism, and it acts in the same way on configurations on any union of blocks; every configuration of blocks is allowed. For $v = (x, kT + j)$ with $0 \le j < T$, anchor at $u = (x, k)$ the term $\Phi^{[T]}_{u,j}(\eta) = \Phi_v(\Pi_T^{-1}\eta)$. Its support is the set of coarse sites whose blocks meet S_v: $(x + [-r_0, r_0]) \times \{k\}$ if $j \ge 1$, and $\{(x, k)\} \cup (x + [-r_0, r_0]) \times \{k-1\}$ if $j = 0$. These T terms at each coarse site form the blocked interaction $\Phi^{[T]}$ on $(\mathcal{A}^T)^{\mathbb{Z}^2}$. The decomposition $t = kT + j$ is unique, so every term of Φ appears exactly once. The coarse translation by (a, b) is the original translation by (a, bT), so $\Phi^{[T]}$ is translation invariant. It has horizontal range r_0 and vertical range one, and each of its anchored terms takes values in $\{0, 1\}$.

Lemma 7.2 (Exact vertical blocking)**.** *Fix $T \ge 1$ and $\beta > 0$, and write γ and $\hat\gamma$ for the specifications of Φ and $\Phi^{[T]}$ at inverse temperature β.*

(i) *For every finite $\hat\Lambda \subset \mathbb{Z}^2$, with U the union of its blocks, and every exterior τ on U^c, the image of γ^τ_U under Π_T is $\hat\gamma^{\Pi_T \tau}_{\hat\Lambda}$.*

(ii) *The map $\mu \mapsto \mu \circ \Pi_T^{-1}$ is a bijection from the DLR measures of γ onto those of $\hat{\gamma}$. The image is invariant under all coarse translations exactly when μ is invariant under the translations by $(1,0)$ and $(0,T)$.*

(iii) *If γ satisfies (3) with constants C, c, then $\hat{\gamma}$ satisfies (3) on the coarse lattice with constants CT, c.*

Proof. (i) A term Φ_v meets U exactly when some site of S_v lies in a block of $\hat{\Lambda}$, that is, exactly when the support of the corresponding coarse term meets $\hat{\Lambda}$. Hence $H_U(\omega \mid \tau) = \hat{H}_{\hat{\Lambda}}(\Pi_T\omega \mid \Pi_T\tau)$ term by term, for every interior ω and every exterior τ. This includes the terms of an exterior child with an interior parent. The two normalizations are the same finite sum under the bijection $\mathcal{A}^U \to (\mathcal{A}^T)^{\hat{\Lambda}}$.

(ii) Let μ be a DLR measure of γ and let $\hat{\Lambda}$ be finite. The sigma algebra of the spins off U is the preimage under Π_T of the sigma algebra of the coarse spins off $\hat{\Lambda}$. Hence (i) carries the DLR identity of μ on U to the DLR identity of $\mu \circ \Pi_T^{-1}$ on $\hat{\Lambda}$. Conversely, let $\hat{\mu}$ be a DLR measure of $\hat{\gamma}$ and put $\mu = \hat{\mu} \circ \Pi_T$. Given a finite $\Lambda \subset \mathbb{Z}^2$, let U be the union of the blocks meeting Λ. By (i) the conditional law of ω_U given the spins off U is γ_U^τ. Conditioning further on $\omega_{U\setminus\Lambda}$ gives the conditional law of ω_Λ given the spins off Λ. All weights are positive, and the terms of H_U that miss Λ cancel, so this law is $\gamma_\Lambda^{\omega_{U\setminus\Lambda}\tau}$. Neither direction uses invariance, and holes or disconnected volumes change nothing. The statement on invariance holds because the coarse translation by (a,b) is the original translation by (a,bT).

(iii) Let v and z lie in the blocks of distinct coarse sites $u = (x_1,k_1)$ and $u' = (x_2,k_2)$. If $k_1 \neq k_2$, the times of v and z differ by at least $T(|k_1-k_2|-1)+1 \geq |k_1-k_2|$. Hence the graph distance from v to z is at least that from u to u', and $\operatorname{dist}(v,U^c) \geq \operatorname{dist}(u,\hat{\Lambda}^c)$ for every site v in the block of $u \in \hat{\Lambda}$. Let $\hat{E} \subset \hat{\Lambda}$ and let E be the union of its blocks. By (i), the $\hat{E}$-marginal of $\hat{\gamma}_{\hat{\Lambda}}^{\Pi_T\tau}$ is the image of the E-marginal of γ_U^τ under a bijection. Since each block has T sites, for all exteriors τ, τ' on U^c

$$d_{\mathrm{TV}}\big((\hat{\gamma}_{\hat{\Lambda}}^{\Pi_T\tau})_{\hat{E}}, (\hat{\gamma}_{\hat{\Lambda}}^{\Pi_T\tau'})_{\hat{E}}\big) \leq C \sum_{v\in E} e^{-c\operatorname{dist}(v,U^c)} \leq CT \sum_{u\in\hat{E}} e^{-c\operatorname{dist}(u,\hat{\Lambda}^c)}.$$

□

The normalized factors of (6) differ from $e^{-\beta\Phi_v}$ by the constant c_β, so (i) holds equally in the normalized form. The terms with $1 \leq j < T$ at a coarse site share a within-row support when this family is nonempty. After merging by support, the energy on that support is an integer between 0 and $T-1$. The offset-zero term has its own support across two coarse rows and takes values in $\{0,1\}$. Thus the bound for all merged terms is $\max\{1, T-1\}$. We keep the anchored terms separate, as in Section 2.

Let $\mathcal{A}'$ be a finite alphabet and Φ' a translation-invariant finite-range interaction on $\mathcal{A}'^{\mathbb{Z}^2}$, given by finitely many anchored terms at each site. On $\mathcal{B} = \mathcal{A}' \times \mathcal{A}' \times \{0,1\}$ the doubled interaction $\Phi'^{(2)}$ has, for each anchored term Φ'_α of Φ', the two anchored terms $(s,s',r) \mapsto \Phi'_\alpha(s)$ and $(s,s',r) \mapsto \Phi'_\alpha(s')$ with the same support; no term reads r. Define $\sigma(s,s',r) = (s',s,1-r)$ at every site. Write υ for the uniform product law on $\{0,1\}^{\mathbb{Z}^2}$ and υ_Λ for the uniform law on $\{0,1\}^\Lambda$.

Lemma 7.3 (Doubling with an exact symmetry)**.** *Fix $\beta > 0$, and write γ' and $\gamma^{(2)}$ for the specifications of Φ' and $\Phi'^{(2)}$.*

(i) *For every finite Λ and every exterior (τ_1,τ_2,τ_r), $\gamma_\Lambda^{(2),(\tau_1,\tau_2,\tau_r)} = \gamma_\Lambda'^{\tau_1} \otimes \gamma_\Lambda'^{\tau_2} \otimes \upsilon_\Lambda$.*

(ii) *The map σ is an involution of $\mathcal{B}$ without fixed points, and it permutes the terms of $\Phi'^{(2)}$.*

(iii) *If γ' satisfies (3) with constants C, c, then $\gamma^{(2)}$ satisfies (3) with constants $2C, c$.*

(iv) *If κ_1 and κ_2 are DLR measures of γ', then $\kappa_1 \otimes \kappa_2 \otimes \upsilon$ is a DLR measure of $\gamma^{(2)}$, and σ maps it to $\kappa_2 \otimes \kappa_1 \otimes \upsilon$. It is translation invariant if κ_1 and κ_2 are.*

Proof. (i) The Hamiltonian of $\Phi'^{(2)}$ is the sum of the Hamiltonians of Φ' in s and in s', and every assignment of r has the same weight, so the Boltzmann weight factors. This holds in every finite volume, with or without holes. (ii) Applying σ twice gives the identity, and $r \neq 1 - r$. The map σ exchanges the two terms built from each Φ'_α, so every H_Λ is invariant. (iii) By (i), for two exteriors the E-marginals are products of the E-marginals of the two copies and of the same law of the free bits. Total variation of products satisfies $d_{\mathrm{TV}}(\mu_1 \otimes \mu_2, \nu_1 \otimes \nu_2) \le d_{\mathrm{TV}}(\mu_1, \nu_1) + d_{\mathrm{TV}}(\mu_2, \nu_2)$, by adding and subtracting $\nu_1 \otimes \mu_2$. Each of the two terms is bounded by (3) for γ'. (iv) Under the product measure the three coordinate fields are independent, so the conditional law of the spins in Λ given those outside is the product of the three conditional laws. These are $\gamma'^{\tau_1}_\Lambda$, $\gamma'^{\tau_2}_\Lambda$ and υ_Λ, and (i) identifies their product with the kernel of $\gamma^{(2)}$. The map σ exchanges the factors of the two copies and preserves υ. Products of translation-invariant measures are translation invariant. □

7.4 Proof of Theorem B

Let $T \ge 2$ be the clock period of Section 3.2. It is fixed with the rule and does not depend on n or β. Let Ψ be the doubled interaction of $\Phi^{[T]}$, on the alphabet $\mathcal{B} = \mathcal{A}^T \times \mathcal{A}^T \times \{0,1\}$ with $2N^{2T}$ letters. It is one translation-invariant finite-range interaction on $\mathbb{Z}^2$, independent of n and β, with horizontal range r_0 and vertical range one. At each site it has $2T$ anchored terms, two copies of each term of $\Phi^{[T]}$, and each takes values in $\{0,1\}$. Merging the terms with a common support gives integer energies at most $2T - 2$, since the $2(T-1)$ terms with $j \ge 1$ of the two copies share one support. By Lemma 7.3(ii) the involution σ of $\mathcal{B}$ has no fixed points and leaves Ψ invariant. Let $\mathcal{A}_0 \subset \mathcal{A}$ be the set of symbols whose data field is live with age zero, and define the one-site observable

$$M(s, s', r) = \tfrac{1}{2}(\mathbf{1}_{\mathcal{A}_0}(s_0) - \mathbf{1}_{\mathcal{A}_0}(s'_0)),$$

where s_0 is the component of $s \in \mathcal{A}^T$ at time offset zero. Then $M \circ \sigma = -M$. The windows J_n are those constructed above.

Let $\beta \in J_n$ with $U(n)$ halting. By Subsection 7.1, the specification (6) of Φ satisfies (3) with some constants $C, c > 0$. Lemmas 7.2(iii) and 7.3(iii) give (3) for the specification of Ψ, with constants $2CT, c$. If ϱ, ϱ' are DLR measures of Ψ and G is an event depending on the spins in a finite set E, then for every finite $\Lambda \supset E$

$$|\varrho(G) - \varrho'(G)| \le \sup_{\tau,\tau'} |\gamma^{(2),\tau}_\Lambda(G) - \gamma^{(2),\tau'}_\Lambda(G)| \le 2CT \sum_{v \in E} e^{-c\,\mathrm{dist}(v, \Lambda^c)},$$

where $\gamma^{(2)}$ is the specification of Ψ. The right side tends to zero as Λ grows, so Ψ has at most one DLR measure. It has one: with $\hat{\mu} = \bar{\mu} \circ \Pi_T^{-1}$, where $\bar{\mu}$ is the unique DLR measure of Φ, Lemmas 7.2(ii) and 7.3(iv) show that $\hat{\mu} \otimes \hat{\mu} \otimes \upsilon$ is a DLR measure of Ψ.

Let $\beta \in J_n$ with $U(n)$ running forever, and let Q, Q_1 be the two DLR measures of Φ from Subsection 7.2. By Lemma 7.2(ii), $\hat{Q} = Q \circ \Pi_T^{-1}$ and $\hat{Q}_1 = Q_1 \circ \Pi_T^{-1}$ are DLR measures of $\Phi^{[T]}$, invariant under all coarse translations. The indicator $\chi(s) = \mathbf{1}_{\mathcal{A}_0}(s_0)$ at the coarse origin satisfies $\hat{Q}(\chi) = \rho(E_0) \ge 3/4$ and $\hat{Q}_1(\chi) = (\rho P)(E_0) \le 1 - (\rho P)(E_1) \le 1/4$. By Lemma 7.3(iv),

$$\mu^+_\beta = \hat{Q} \otimes \hat{Q}_1 \otimes \upsilon, \qquad \mu^-_\beta = \hat{Q}_1 \otimes \hat{Q} \otimes \upsilon$$

are translation-invariant DLR measures of Ψ, and σ exchanges them. Moreover $\mu^+_\beta(M) = \frac{1}{2}(\hat{Q}(\chi) - \hat{Q}_1(\chi)) \ge \frac{1}{4}$ and $\mu^-_\beta(M) = -\mu^+_\beta(M) \le -\frac{1}{4}$.

Nothing is asserted at inverse temperatures outside the union of the J_n. The space–time interaction Φ itself has exactly one DLR measure in each halting window and at least two in each nonhalting window, by Subsections 7.1 and 7.2. This completes the proof of Theorem B.

8 Effective openness of planar mixing sets

We prove Theorem D for every translation-invariant finite-range rational interaction with finite energies; the symmetry and observable play no role in this direction. Distances in the three sections on mixing sets are maximum-norm distances. Since $\mathrm{dist}_\infty \le \mathrm{dist}_1 \le 2\,\mathrm{dist}_\infty$, changing between this convention and (3) changes only the exponential rate.

8.1 The external weak-to-strong mixing implication

The one external implication needed for the upper bound is the following; it is also used, through the equivalence below, to state Theorem F in the convention (3). Suppose a finite-alphabet, translation-invariant, finite-range, finite real-valued interaction on $\mathbb{Z}^2$, with no hard constraints, satisfies

$$d_{\mathrm{TV}}\big((\gamma^{\tau}_{\Lambda})_D, (\gamma^{\tau'}_{\Lambda})_D\big) \le C_w |D| e^{-m_w \,\mathrm{dist}_\infty(D,\Lambda^{\mathrm{c}})} \tag{10}$$

for all finite Λ, $D \subset \Lambda$, and all exteriors, with $C_w, m_w > 0$. Then there are $L_0, C_s, m_s > 0$ such that, for every square B of odd side $L \ge L_0$, every $D \subset B$, and exteriors differing at one site j,

$$d_{\mathrm{TV}}\big((\gamma^{\tau}_{B})_D, (\gamma^{\tau'}_{B})_D\big) \le C_s |D| e^{-m_s \,\mathrm{dist}_\infty(D,j)}. \tag{11}$$

This is the two-dimensional theorem of Martinelli, Olivieri and Schonmann [29, Theorem 1.1], in the cardinality-times-distance formulation recorded by Spinka [30, equations (2)–(3), Section 1.2]; his Remark 2 gives the finite-range extension. The original displayed weak-mixing convention in [29] is a sum over pairs, so we use Spinka's explicit formulation bridge. The conclusion here is restricted to sufficiently large squares; both sources state it for the centered boxes $[-n,n]^2 \cap \mathbb{Z}^2$, hence, by translation invariance, for squares of odd side. It asserts neither strong mixing on arbitrary volumes nor effectivity of L_0, C_s, m_s. Multi-site potentials are allowed; ferromagnetism and pair interactions are not hypotheses.

Finite energies give full support. EWM gives uniqueness by letting a surrounding box grow in the DLR equation, and the unique state is translation invariant. Thus the field formulation in [30] applies to our specification. Estimate (3) implies (10). All remaining steps are finite-volume arguments proved below.

8.2 A finite block certificate

Choose an integer $R \ge 1$ bounding the maximum-norm diameter of every interaction support. For $B = \{0,\dots,L-1\}^2$, let

$$J_B = \{j \notin B : \mathrm{dist}_\infty(j,B) \le R\}, \qquad N_B = L^2.$$

Every exterior coordinate read by its Hamiltonian belongs to J_B. On laws on $\mathcal{A}^B$, let W_H denote optimal transport with the Hamming cost, the number of different coordinates. Define

$$c_j(\beta) = \max_{\substack{\tau,\tau' \text{ on } J_B \\ \tau=\tau' \text{ off } j}} W_H(\gamma^{\tau}_{B,\beta}, \gamma^{\tau'}_{B,\beta}), \qquad C_B(\beta) = \sum_{j \in J_B} c_j(\beta). \tag{12}$$

The certificate is $C_B(\beta) < N_B$. It is the Dobrushin–Shlosman finite-size uniqueness condition with the Hamming metric [40], in the form stated by Martinelli, Olivieri and Schonmann [29, Section 3] and the block-coupling form developed by Weitz [41]. Dobrushin and Shlosman showed that it implies weak mixing in all finite volumes, and Martinelli, Olivieri and Schonmann showed that in two dimensions it holds on some finite set exactly when weak mixing holds in all finite volumes [29, Theorems 3.1 and 3.2]; the same chain is used for

space–time specifications in [8, Theorem 12.1]. The two lemmas below prove both directions in the conventions and with the constants used here. What Theorem D adds is the computable modulus and the uniform enumeration in the inverse temperature. For the zero-field nearest-neighbor ferromagnetic Ising model on $\mathbb{Z}^d$, Shlosman proved $\sup_V \beta_V(d) = \beta_{cr}(d)$, where the finite-volume Dobrushin–Shlosman condition on V holds for $\beta < \beta_V(d)$ [42, Theorem 4].

Lemma 8.1 (Sufficiency of a block certificate)**.** *If $C_B < N_B$, the specification has arbitrary-volume EWM. More precisely, put $D_B = L + 2R + 1$. If $C_B > 0$, set*

$$\kappa = \frac{\log(N_B/C_B)}{2D_B}, \quad h = e^{\kappa D_B} C_B = \sqrt{N_B C_B}, \quad g = N_B - h.$$

If $C_B = 0$, set $\kappa = 1, h = 0, g = N_B$. With

$$m = \kappa/2, \qquad r = e^{-\kappa/2}, \qquad K = 1 + \frac{8r}{(1-r)^2}, \qquad C = 2e^{mL} \max\{1, hK/g\},$$

the maximum-norm version of (3) *holds with constants C, m.*

Proof. Fix a finite volume Ω and two exteriors. Update each full translate $t + B \subset \Omega$ at rate one by its heat bath. For every pair of block-boundary assignments differing at one site, choose an optimal Hamming coupling. For a pair differing at several sites, interpolate in a fixed order and glue the one-site couplings. All intermediate assignments are legal by positivity. If the boundaries agree, use identical outputs. These choices give a finite-state coupled chain, each of whose marginals preserves its prescribed Gibbs law. Start with those marginals and take a limit of time averages to obtain an invariant coupling Q. Irreducibility is unnecessary. Write q_x for its disagreement probability, including the fixed exterior disagreements.

The choices of coupling are independent of the weights that follow. Fix $v \in \Omega$, put $w_x = e^{-\kappa|x-v|_\infty}$, and let n_x be the number of full translates containing x. On changing a boundary site j of $t + B$, each output coordinate has $w_x \le e^{\kappa D_B} w_j$; its expected total weighted cost is therefore at most $e^{\kappa D_B} w_j c_{j-t}$. Gluing and the triangle inequality give the sum of these bounds for several boundary disagreements. For any source site, summing over all affected translates costs at most h. On

$$G = \{x \in \Omega : \operatorname{dist}_\infty(x, \Omega^c) \ge L\}$$

we have $n_x = N_B$. Stationary drift of weighted disagreements consequently gives

$$0 \le -\sum_{x \in \Omega} n_x w_x q_x + h \sum_{x \in \mathbb{Z}^2} w_x q_x, \qquad g \sum_{x \in G} w_x q_x \le h \sum_{x \notin G} w_x.$$

If $d_v = \operatorname{dist}_\infty(v, \Omega^c) \ge L$, every $x \notin G$ has $|x - v|_\infty \ge d_v - L$. The shell count $|\{x : |x|_\infty = n\}| = 8n$ yields

$$q_v \le (hK/g) e^{-m(d_v - L)}.$$

For $d_v < L$, use $q_v \le 1$. The same invariant coupling works for every v, so the union bound on disagreements in any $E \subset \Omega$ proves (3) with the displayed constants. This also covers $h = 0$ and volumes containing no full block. □

Lemma 8.2 (Completeness of square certificates)**.** *If the specification has EWM, or already just* (10)*, some square has $C_B < N_B$. With the constants of* (11)*, it suffices to choose an odd integer*

$$L \ge \max\{1, L_0\}, \quad L > H_B^2, \qquad H_B = 4R(1+R)(6400/m_s^2 + 18 + C_s).$$

Proof. For one changed exterior site j, put $s_L = \lceil 5\log(L+1)/m_s \rceil$ and $F = \{x \in B : |x-j|_\infty > s_L\}$. By (11) its two marginal laws have TV distance at most $C_s L^2 (L+1)^{-5}$. Maximally couple those laws and then complete the conditional distributions on $B \setminus F$. The expected Hamming cost is at most

$$(2s_L+1)^2 + C_s L^4/(L+1)^5 \le (2s_L+1)^2 + C_s.$$

Since $|J_B| = 4RL+4R^2$, this bounds C_B by $(4RL+4R^2)[(2s_L+1)^2+C_s]$. Using $\log(L+1) \le 4(L+1)^{1/4}$ gives $C_B \le H_B L^{3/2} < L^2$. □

In particular, (10) and (3) are equivalent for these planar finite-energy specifications (compare [29, Theorem 3.2]): the first gives a square certificate and Lemma 8.1 gives the sum-over-sites estimate. We will use this equivalence for the transfer-matrix construction below.

8.3 Effective continuity and enumeration

Lemma 8.3 (Computable certificate and modulus)**.** *The number $C_B(\beta_0)$ is a computable real, uniformly in a rational $\beta_0 > 0$ and the finite rational table. If E_B is the range of the finite Hamiltonian over all configurations on $B \cup J_B$, then*

$$|C_B(\beta) - C_B(\beta')| \le K_B|\beta-\beta'|, \qquad K_B = |J_B| N_B E_B.$$

Proof. All energies in this finite Hamiltonian are rational. Differentiation gives $\partial_\beta \gamma(\omega) = \gamma(\omega)[\gamma(H) - H(\omega)]$, so the TV distance between two temperatures is at most $E_B|\beta-\beta'|/2$. The Hamming metric has diameter N_B; the transport triangle inequality and maximal coupling give

$$|W_H(\mu,\nu) - W_H(\mu',\nu')| \le N_B[d_{\mathrm{TV}}(\mu,\mu') + d_{\mathrm{TV}}(\nu,\nu')].$$

Finite maxima and the sum in (12) give the modulus. At rational temperature every Gibbs probability is computable by exponential enclosures and finite normalization. Approximate its probability vector within certified TV error ϵ by a rational vector. The rational transport optimum is computable by enumerating feasible bases of its finite linear program; the last inequality bounds its error by $2N_B\epsilon$. This produces rational upper bounds $u_n(\beta_0) \to C_B(\beta_0)$. Boltzmann probabilities need not be rational. □

Proof of Theorem D. Dovetail over $L \ge 1$, rational $\beta_0 > 0$, rational $0 < \rho < \beta_0$, and precision n. Enumerate $(\beta_0-\rho, \beta_0+\rho)$ whenever

$$u_n(\beta_0) + K_B\rho < N_B.$$

Every enumerated closed interval satisfies the certificate, hence has EWM by Lemma 8.1. Conversely, an EWM point has a strict square certificate by Lemma 8.2. Continuity, rational density and sufficiently accurate upper bounds supply an enumerated interval containing it. This remains true when $K_B = 0$. Intersect the union with $[a,b]$; a mixing endpoint is covered by an ordinary neighborhood before restriction. The enumeration uses only the finite table and range, never the external mixing constants. □

9 Sequential realization of effectively open sets

Fix the radius-one member of the native family with a literal program that runs forever. Write $\mathcal{H}$ for its alphabet, F for its deterministic map, and D for its single empty state. Its ordinary simulated maps satisfy $F_j(\mathsf{D},\mathsf{D},\mathsf{D}) = \mathsf{D}$. We use the deterministic organization, the initialized occupied comparison of Lemma 6.4, and the descent in Lemma 6.5. The two additional lemmas below make the probabilistic and empty-state arguments explicit.

9.1 Local alarms and sparse deviations

An innovation vector at a logical space–time vertex may have dependent coordinates; vectors at different vertices are independent. A *paint* is simply a set of vertices declared possible deviations. The following estimate permits overlapping local alarms computed from the same vectors as the raw errors.

Lemma 9.1 (Sparsity with local alarm paint)**.** *Use the native P, T, w, put $S = \max(P, T)$, and fix positive integers M, M_0. (In Lemmas 9.1 and 9.2, M and M_0 are rectangle multipliers local to these lemmas; they differ from the constants of the same names in the appendix tables and from the observable M.) Raw flags are iid with probability at most u. For every $k \geq 1$ and center z, suppose a potential alarm $A_k(z)$ has probability at most $e^{-H_{\mathrm{al}}\lambda^k}$, with a fixed $\lambda > 1$, depends only on innovations in the rectangle about z of half-axes MP^k, MT^k, and paints a subset of that rectangle. Let E be any subset of the raw flags and all such paints. For each $i \geq 0$, prescribe a rectangle of half-axes M_0P^{i+1}, M_0T^{i+1}, with deterministic center. Then*

$$\sum_{i \geq 0} \mathbb{P}(E \text{ fails } i\text{-sparsity in the prescribed rectangle}) \leq \tfrac{1}{8}$$

under the following explicit choices:

$$1 < \nu < \min(2, \lambda), \quad B_* = 8(P+2)^2(T+2)^2, \quad D_* = (10 + 8w + 8M)^2,$$
$$K_* = \min\{k \geq 1 : 2^k/D_* \geq 2\nu^k\}, \quad C_t = 2^{K_*} B_*^{2^{K_*}},$$
$$A_* = (2M_0P + 10)(2M_0T + 10), \quad C_0 = (2w + 2M + 3)^2,$$
$$C_h = (2M_0 + 2M + 3)^2 S^{2K_*},$$
$$H_{\mathrm{raw}} \geq \max\left\{1, 2^{K_*}\log B_* + \log 2, \frac{\log 2}{\nu^{K_*} - 1}, \log(32K_*A_*C_t)\right\},$$
$$H_{\mathrm{al}} \geq \max\left\{1, \lambda^{K_*}\left[H_{\mathrm{raw}} + \log(2K_*C_0) + \frac{2\log S}{\log\lambda}\right], \frac{2\log S + \log 4}{\lambda - 1}, \log[16(K_* + 1)C_h]\right\},$$
$$u \leq e^{-H_{\mathrm{raw}}}/(4w^2).$$

All lower bounds can be rounded strictly upward to rationals. Alarms at different ranks need not be independent. A nonlocal actual alarm is also allowed if it implies an event with the stated local support and bound.

Proof. For a multiple r of K_*, let b_r be the supremum failure probability in a full rank-r box of sides $2wP^r, 2wT^r$, using raw flags and paints through rank r. Initially $b_0 \leq 4w^2u$. If no paint of ranks $r - K_* + 1, \ldots, r$ meets a failed box, Lemma 4.2, applied K_* times to this same restricted error field, gives 2^{K_*} disjoint full leaves at rank $r - K_*$. There are at most $B_*^{2^{K_*}-1}$ trees.

A lower-aggregate failure at a leaf of rank l depends only on that rectangle enlarged by $2MP^l, 2MT^l$ on each side, once for paint and once for its support. The intersection graph of enlarged disjoint leaves has degree plus one at most D_*: after scaling the axes, the relevant original disjoint boxes fit in a square of side at most $6w + 8M + 2$, and division by their area gives this conservative bound. A greedy independent set leaves at least $2^{K_*}/D_*$ events with independent innovation supports. This estimate is unconditional; we do not condition on absence of the larger paints. A union bound for the possible new centers gives

$$b_r \leq B_*^{2^{K_*}} b_{r-K_*}^{2^{K_*}/D_*} + K_*C_0S^{2r}e^{-H_{\mathrm{al}}\lambda^{r-K_*+1}}. \tag{13}$$

The displayed choices imply by induction $b_r \leq e^{-H_{\mathrm{raw}}\nu^r}$. For the first term, $(2^{K_*}/D_*)\nu^{r-K_*} \geq 2\nu^r$ pays the tree count and half the budget. For the second, use $\nu^r \leq \lambda^r$, $1 \leq \lambda^r$, and $r \leq \lambda^r/\log\lambda$ in the bound on H_{al}.

For test rank i, put $r = K_*\lfloor i/K_* \rfloor$. Cover its moving rank-$i$ windows by full rank-i boxes at steps wP^i, wT^i, including endpoint overlaps, as in Lemma 4.2. There are at most A_* boxes. A failure with paints only through r descends at most $K_* - 1$ times to a failed rank-r leaf, with at most $A_* C_t$ choices. Paints of rank $k > r$ meeting the original prescribed rectangle have at most $C_h S^{2k}$ possible centers, since $i + 1 \le k + K_*$. Thus

$$\mathbb{P}(\text{test } i \text{ fails}) \le A_* C_t e^{-H_{\mathrm{raw}}\nu^r} + C_h \sum_{k>r} S^{2k} e^{-H_{\mathrm{al}}\lambda^k}. \tag{14}$$

The sum of the first terms over i is at most $2K_* A_* C_t e^{-H_{\mathrm{raw}}} \le 1/16$, using $(\nu^{K_*})^j \ge 1 + j(\nu^{K_*} - 1)$. Each rank k in the second term is counted at most $k + K_*$ times. Since $\lambda^k \ge 1 + k(\lambda - 1)$, their total is at most $C_h e^{-H_{\mathrm{al}}} \sum_{k\ge1}(k + K_*)4^{-k} \le 1/16$. Heredity of sparsity handles every subset E. □

Lemma 9.2 (Empty-sector ancestry)**.** *Use the fixed nonhalting native rule above. Join an actual history at an entrance date τ to the deterministic all-*D *history before τ, and suppose its deviations are contained in E. Put $m = \lfloor w/5 \rfloor$ and $M_0 = 100(T + P + m + 100)$. There are deterministic centered tests, one at each rank i, contained in rectangles of half-axes $M_0 P^{i+1}, M_0 T^{i+1}$, such that a live physical site implies failure of at least one test. Consequently, with the hypotheses of Lemma 9.1, the all-empty entrance has live probability at most $1/8$ at every site and later date. From the occupied entrance of Lemma 6.4, under the same bounds, the live probability is at least $7/8$.*

Proof. The all-empty deterministic history is valid because $F_j(\mathsf{D}, \mathsf{D}, \mathsf{D}) = \mathsf{D}$. Suppose the enclosing tests pass and a rank-j served site is live. Use the intermediate ascent in the proof of Lemma 6.3, before its halted-rank conclusion: the site's clock and address choose a translated parent grid. Trace incumbents or donors to the cut one full period before its latest cut, and then to the inputs of the preceding parent work. Proposition 5.3(iii) and the empty-exception cut exclude a dead issue. The positive issue is the native map on three actual served inputs, at least one of which must therefore be live. Continue from that rank-$(j+1)$ input, reselecting its grid as necessary.

The accumulated displacement and protection padding through rank j, as calculated in Lemma 6.3, have horizontal half-extent $(4T/P + 4 + m + 2 + 4/P)P^{j+1}$ and temporal half-extent $(6 + m + 2 + 4/T)T^{j+1}$. They fit inside the stated M_0 rectangles. They contain the protected queries and service lookahead, so passing the centered test suffices on every selected grid. In each step the move to the earlier cut and the preceding work's input moves the bottom date back by at least two parent periods, $2T^{j+1}$ in physical units. The rank-j service lookahead is less than T^j. Eventually the entire reached observation is before τ, where every grid observes only D. This contradicts its being served live. This argument needs no halted top rank.

The union bound and Lemma 9.1 give the empty estimate. For the occupied estimate use the deterministic nested descent of Lemma 6.5, with the present dominating deviation set. Its tests fit the same rectangles. The compatible negative-time occupied comparison supplies the top inputs, and passing all tests gives the live clock at the observation. The same sum of failure probabilities is at most $1/8$. Flags may be extended independently before entry for these estimates; actual deviations there remain absent. □

9.2 A finite thermal counter

Fix $M = 2$, M_0 as in Lemma 9.2, $\lambda = 3/2$ and $\nu = 5/4$ in Lemma 9.1. Choose its rational constants, then rational $u_c > 0$ with $u_c < e^{-H_{\mathrm{raw}}}/(4w^2)$. Set

$$Q = 2|\mathcal{H}|, \qquad J = 1 + \left\lceil \frac{Q}{au_c} \right\rceil, \qquad u(\beta) = \frac{Q}{e^{J\beta} + Q - 1}, \qquad p(\beta) = u(\beta)/2. \tag{15}$$

Thus $u(\beta) < u_c$ on $[a, b]$. These choices precede the enumerator of W. The score map $p : (0, \infty) \to (0, 1/2)$ is a computable strictly decreasing homeomorphism, with inverse

$$p^{-1}(v) = J^{-1} \log[Q/(2v) - (Q - 1)].$$

For a positive integer N and finite rational union $D \subset (0, 1/2)$ of open intervals, use the physical alphabet $\mathcal{H} \times \{0, 1\}$, with coordinates (H, Z). At site x the raw bit proposes zero. On the old row compute

$$V_x = \frac{1}{N} \sum_{i=0}^{N-1} Z_{x+1+2i}, \qquad H \text{ proposes } \begin{cases} \mathsf{D}, & V_x \in D, \\ F(H_{x-2}, H_x, H_{x+2}), & V_x \notin D. \end{cases}$$

Apply one independent uniform whole-symbol replacement of probability $u(\beta)$. This is a literal finite table: the sum and comparisons can be carried out by integer arithmetic in its definition. Its space–time interaction assigns energy zero when the complete child equals its proposal and J otherwise. Its normalized child sum is parent-independent, and gives exactly (15).

For one parity lane, all its reporter sites lie in the opposite parity. After one update the reporter rows are iid Bernoulli(p) in space and time, independently of that lane's entire native innovation tape. Simultaneous tests may overlap. Pack the innovations of one site of each parity into a single independent logical vector. Take all initial reporter bits zero; the first detector cannot fire because $0 \notin D$.

Lemma 9.3 (Finite counter away from its collar)**.** *Fix a rational margin $\delta > 0$. Put*

$$\eta_N = 2e^{-2N\delta^2}, \qquad i_N = \max\{1, \min\{i \geq 0 : P^i \geq N + 1\}\}.$$

Choose N effectively large enough that

$$\eta_N < \min\{1/12, e^{-H_{\mathrm{al}}(3/2)^{i_N}}\}. \tag{16}$$

If $(p(\beta) - \delta, p(\beta) + \delta) \subset D$, the complete PCA is uniformly exponentially ergodic, with one-site common-innovation disagreement at most $(3\eta_N)^{t-1}$ for $t \geq 2$. If $\beta \geq a$ and $\mathrm{dist}(p(\beta), D) \geq \delta$, its mismatch interaction has two translation-invariant DLR measures whose live probabilities are at least $7/8$ and at most $1/8$, respectively. Its doubled magnet has magnetic order with gap $3/8$ in the latter case and arbitrary-volume EWM in the former.

Proof. For N independent reporter bits, the centered log moment-generating function is at most $h^2/8$: the second derivative of the log moment-generating function is a variance on an interval of length one, at most $1/4$, and its value and first derivative at zero vanish. Multiplication and optimization at $h = 4\delta$ give $\mathbb{P}(|V_x - p| \geq \delta) \leq \eta_N$. This bounds a false alarm in the exterior case and a missed alarm in the interior case. The search (16) terminates because $i_N \leq 1 + \log_P(N + 1)$, $P > 3/2$, and $N\delta^2$ eventually dominates $(N + 1)^{\log_P(3/2)}$.

On the cold side every actual native step is F on its actual inputs except at a native raw flag or an alarm that writes D. A past alarm has no further charged effect: later non-alarm steps again use F on the current, possibly empty, inputs. Each potential false alarm has a determining footprint inside the rank-i_N box of half-axes $2P^{i_N}, 2T^{i_N}$, and paints only its center. Other ranks have no alarms. The packed logical innovations therefore satisfy Lemma 9.1, even for overlapping detectors. Lemma 9.2 gives the two live bounds from the explicit occupied and empty entrances, uniformly in position and date.

Average each forward history over growing positive space–time rectangles and take a weak limit. Boundary fractions vanish under every fixed translation. The finite-cone factorization in Section 7.2 gives the DLR identities for every volume strictly above the entrance; these

local identities survive the averages and the limit. The resulting translation-invariant DLR measures retain their respective live bounds. Double the interaction as in Lemma 7.3, add its free bit, and use

$$M(h, h', r) = \tfrac{1}{2}(\mathbf{1}_{\{h \text{ live}\}} - \mathbf{1}_{\{h' \text{ live}\}}).$$

The two orders of the factors give means at least $(7/8 - 1/8)/2 = 3/8$ and at most $-3/8$, exchanged by the fixed-point-free involution. Here no time blocking is needed because the live indicator survives the space–time average.

On the warm side couple any two full starting rows by the same physical innovations. Their entire raw fields agree after one tick. A native discrepancy at a later date requires a missed common detector and a disputed one of its three lane-neighbor parents. A fixed backward path of length $t-1$ uses distinct positive reporter dates; its misses are independent and cost at most η_N^{t-1}. There are at most 3^{t-1} such paths. The raw bit has already coupled, giving the stated bound for the full symbol. A finite-set union bound and comparison with a stationary start prove (2). The PCA is finite-alphabet, finite-radius and strictly positive, so Theorem 7.1, imported from [8, Theorem 3.1], gives arbitrary-volume EWM of the normalized space–time specification. It has the same specification as the mismatch energy, differing only by a constant per child. Lemma 7.3 preserves EWM. □

9.3 Nesting the actual mixing sets

Proof of Theorem E. Write $W = G \cap [a, b]$ for the supplied effectively open ambient set, and put $O = p(G \cap (0, \infty))$. Its rational interval enumeration is computable from the homeomorphism p. Enumerate rational pairs (c_i, ρ_i), $\rho_i > 0$, with $[c_i - 3\rho_i, c_i + 3\rho_i] \subset O$, so that their cores $[c_i - \rho_i, c_i + \rho_i]$ cover O. Strict enclosure in an enumerated image interval is semidecidable and suffices to produce such a cover.

At stage s take the finite list issued in the first s enumeration steps. For a pair issued at step $j_i \le s$, put

$$r_i(s) = 2\rho_i - \rho_i 2^{-(s-j_i)}, \qquad d_i(s) = \rho_i 2^{-(s-j_i)-3}.$$

Let K_s be the closed union of radii $r_i(s)$, let D_s be the open union of radii $r_i(s) + d_i(s)$, and let O_s be the closed union of radii $r_i(s) + 2d_i(s)$, all about c_i. They lie in O, and

$$O_s \subset \operatorname{int} K_{s+1}, \qquad r_i(s+1) - r_i(s) = 4d_i(s).$$

For a nonempty list take $\delta_s = \min_i d_i(s)$ and choose N_s by (16). It may also be required to exceed every preceding batch size. Use the doubled magnet of Lemma 9.3. It has EWM on $p^{-1}(K_s)$ and magnetic order off $p^{-1}(O_s)$. Order excludes EWM because EWM implies uniqueness. Therefore, with every set restricted to $[a, b]$,

$$p^{-1}(K_s) \subseteq \mathcal{M}_{\Phi_s} \subseteq p^{-1}(O_s) \subseteq p^{-1}(K_{s+1}) \subseteq \mathcal{M}_{\Phi_{s+1}}. \tag{17}$$

The strict inclusion in the interior gives the asserted relative compact containment, including window endpoints. Every point of W eventually belongs to an issued core. Empty stages have $D = \varnothing$, with $N = 1$; the two clean-sector estimates use raw flags alone and give order everywhere, hence an empty EWM set. The displayed containments prove the union and all nesting claims. Every $\beta \notin W$ is outside every $p^{-1}(O_s)$, so every stage orders there.

The alphabet is the fixed $(\mathcal{H} \times \{0, 1\})^2 \times \{0, 1\}$; J, the involution and M are common. Each anchored term has energy zero or J, vertical range is one, and the finite horizontal range grows with the batch size. All searches and finite tables are uniform in the enumerator and stage. □

Together Theorems D and E give an equivalence: a subset of the window is relatively effectively open if and only if it is a union of actual EWM sets of a uniformly computable sequence of finite rational magnets. The realizing sequence may moreover have the nesting and common data of Theorem E. The uncertain part of a stage lies inside its collars in W; no order claim is made there.

10 Exact explicit diagrams by line decorations

The construction in this section uses auxiliary Markov chains shared by consecutive edges of each row and column. Summing finite decorations into effective couplings is classical [31, 32]; sharing a line lets a selected positive-matrix eigenvalue distinguish real algebraic conjugates.

10.1 The decorated magnet and its Potts input

At a vertex put a color in $\{0, \dots, q-1\}$, $q = 65536$, horizontal and vertical internal states in $\{1, \dots, k\}$, and an index for each outgoing east and north edge in $\{1, \dots, L\}$. For each ordered internal pair and each equality class $e \in \{0, 1\}$ of the endpoint colors, specify a list of L rational energies. The index at the tail selects its edge energy. Every alphabet combination is allowed. There are just two translation-invariant nearest-neighbor terms per site. The color half-turn $s \mapsto s + q/2 \pmod q$ fixes the other fields and is an invariant fixed-point-free involution. Let $M = 1$ on the first color half and $M = -1$ on the second.

At inverse temperature β, summing the edge index gives strictly positive transfer matrices T_1, T_0, used identically in the two orientations. The ordinary square-lattice Potts input has coupling $\log R$, where $R \geq 1$, and threshold

$$c = 1 + \sqrt{q} = 257, \qquad p_c = \frac{\sqrt{q}}{1 + \sqrt{q}} = \frac{256}{257}.$$

Theorems 1.2–1.3 of Duminil-Copin, Raoufi and Tassion [36] give the critical value (proved for every $q \geq 1$ by Beffara and Duminil-Copin [45]) and the exponential *finite-volume wired* connection bound below it: in a graph-distance ball of radius n, the wired probability of a center-to-boundary connection is at most $e^{-c_P n}$, for some $c_P > 0$. Exploring the dominating wired boundary cluster couples any two fixed-color Potts exteriors: discrepancies on a set V require a wired connection from a site of V to the boundary. Indeed a colored random-cluster edge has conditional open probability zero for incompatible pinned clusters, p when no free color is lost, and $p/[p + q(1-p)]$ otherwise; all are dominated by the fully wired exploration. At each adaptive step retain each marginal's own partial-assignment conditional kernel. Off the explored cluster the untested interior laws are the same free law and can be coupled. The wired-ball bound thus gives arbitrary-domain EWM, including domains with holes, when $1 \leq R < c$.

For order at equality we use the free critical diagonal inverse correlation length from [37, Theorem 1.2(3)] (there $x_n = (\lfloor n/2 \rfloor, \lfloor n/2 \rfloor)$):

$$\alpha = \ell + 2 \sum_{n \geq 1} \frac{(-1)^n \tanh(n\ell)}{n}, \qquad \cosh \ell = 128.$$

Since $\tanh x / x$ decreases for $x > 0$, the alternating-series estimate gives $\alpha \geq \ell - 2 > \log 8$. If τ denotes free critical connectivity, FKG concatenation gives $\tau(ns, ns) \geq \tau(s, 0)^{2n}$; the diagonal limit therefore gives $\tau(s, 0) \leq e^{-s\alpha}$. A finite wired cluster at the origin has a surrounding dual free open circuit crossing both horizontal rays. Its crossing endpoints are separated by

$i + j + 1$, for some $i, j \geq 0$. At self-duality the component/face identity exchanges wired and free boundary conditions. A union bound consequently gives

$$\phi^w_{p_c}(0 \text{ finite}) \leq \sum_{i,j \geq 0} e^{-(i+j+1)\alpha} = \frac{e^{-\alpha}}{(1 - e^{-\alpha})^2} < \frac{8}{49}.$$

Monotonicity extends this to $p \geq p_c$. Coloring the infinite wired cluster with a fixed color and all finite clusters uniformly gives a translation-invariant Potts DLR state whose M-mean exceeds $41/49$. The paired color gives the opposite state. We retain the uniform gap $1/4$. This uses the discontinuous critical order in [37], rather than treating equality as an unclassified temperature.

10.2 Shared Perron directions and arbitrary exteriors

Lemma 10.1 (Line decoration criterion)**.** *Suppose at a fixed temperature the two positive matrices have common strictly positive right and left Perron vectors r, l, with eigenvalues λ_1, λ_0 and $R = \lambda_1/\lambda_0 \geq 1$. The full decorated magnet has arbitrary-volume EWM if $R < 257$, and has opposite translation-invariant DLR states, exchanged by the color half-turn, with M-means of magnitude at least $1/4$ if $R \geq 257$.*

Proof. Put $\pi_i = l_i r_i/(lr)$ and

$$P_e(i,j) = \frac{T_e(i,j) r_j}{\lambda_e r_i}, \quad \eta = \min_{e,i,j} P_e(i,j) > 0, \quad \rho = 1 - k\eta/2 \in [1/2, 1).$$

Both stochastic matrices preserve π. Coupling rows by their common mass gives, for every product P of n transitions,

$$|P(i,j)/\pi_j - 1| \leq \rho^n/\pi_{\min}. \tag{18}$$

Given a Potts Gibbs state, condition on all colors and independently put the two-sided chain with marginal π and transitions P_e on each row and column. Consistency follows from $\pi P_e = \pi$. Draw the outgoing indices with their conditional normalized Boltzmann weights. If $w_e(i,j,s)$ is an indexed weight, then

$$\lambda_e P_e(i,j) \frac{w_e(i,j,s)}{T_e(i,j)} = w_e(i,j,s) \frac{r_j}{r_i}.$$

For a finite change the r-ratios telescope along each line. The Potts factors, chain factors and index densities therefore give exactly the decorated Hamiltonian's finite-change ratios, also for a fixed exterior tail index. Positivity and finite normalization give the DLR equations. The lift commutes with translations and keeps the color observable. This proves the order assertion.

For mixing, let E_0 be the tested full spins in a finite Λ, with distance d to Λ^c. Set

$$E = E_0 \cup (E_0 + \text{east}) \cup (E_0 + \text{north}), \qquad h = \lfloor (d-2)/4 \rfloor,$$
$$U = E + [-h, h]^2, \qquad D = E + [-2h, 2h]^2.$$

Take d large; small distances use TV at most one. The domain D lies inside Λ, every maximal row or column run has length at least $4h + 1$, and there are at most $2|D|$ runs. Condition on a full exterior for D. Summing indices and internal states along one run of length n, and dividing by its interior Perron factors, gives $u^{\mathsf{T}} P_{e_1} \cdots P_{e_{n-1}} v$, with positive end vectors u, v. By (18) its ratio to $(\sum u_i)(\sum \pi_j v_j)$ lies in $1 \pm \epsilon$, where $\epsilon = \rho^{4h}/\pi_{\min}$. These reference factors depend only on end colors. Their color law is ordinary Potts on interior edges with arbitrary positive weights on the one-site inner collar.

For $\epsilon \le 1/2$, replacing all run factors changes the normalized density in TV by at most $60N_{\mathrm{run}}\epsilon$. This follows by bounding its log ratio by $4N_{\mathrm{run}}\epsilon$ and using $\min(1, e^{4y}-1) \le 60y$. For two exteriors the total run error is at most $240|D|\epsilon$. Condition the two reference laws on their inner collars. The remaining color kernels are ordinary Potts with fixed boundaries, so the Potts estimate couples the colors in U with error at most $C_P|U|e^{-c_P h}$, uniformly over both collars.

On agreement in U, compare hidden states on the runs of U containing tested sites of E, conditioning their outside states first. There are at least h agreed transitions at either end. The law of the tested segment differs from the common stationary-start Markov segment only by two end-weight ratios in $1 \pm \epsilon'$, where $\epsilon' = \rho^h/\pi_{\min}$. The same density estimate, summed over at most $2|E|$ runs in both orientations and both exteriors, gives error at most $480|E|\epsilon'$. Mix over the conditioned end states. The indices in E_0 couple identically once their colors and internal endpoints agree, because their heads belong to E. Thus the full-spin TV bound is at most

$$\frac{240(4h+1)^2|E|\rho^{4h}}{\pi_{\min}} + C_P(2h+1)^2|E|e^{-c_P h} + \frac{480|E|\rho^h}{\pi_{\min}}. \tag{19}$$

Polynomial factors absorb into an exponential; for example a rate $\min(c_P, -\log\rho)/16$ works after enlarging the constant. Conditioning on the complement of D and then mixing proves the cardinality form (10) for every Λ. Its equivalence to (3), established in Section 8, gives our stated EWM convention. The inflated domains ensure that short runs or holes in the original volume cause no loss. □

10.3 From polynomial matrices to finite energy lists

Write $z = e^{-\beta/d}$, $t = z/(1+z)$, and $t_0 = 1/2$. Since $a > 0$, the window image lies in a rational compact interval $I \subset (0, 1/2)$. The endpoints of its image are transcendental, so comparisons with algebraic cuts terminate. The following conversion will be used without increasing matrix dimension.

Lemma 10.2 (Positive matrix conversion)**.** *Suppose rational polynomial matrices D, F commute, D is strictly positive on $[0,1]$, and $D(t_0)$ has all entries equal and positive. Suppose F has scalar $f(t)$ on the positive Perron direction of D, and $f(t_0) \neq 0$. One can compute two positive polynomial matrices A, B sharing both positive Perron directions, with ratio $R > 1$ on I and*

$$\operatorname{sign}(R - 257) = \operatorname{sign} f \quad \text{on } I.$$

After a common scalar change they are finite Boltzmann sums in z with nonnegative integer coefficients and identical total multiplicity for every entry of both matrices. Their energy lists use rational energies i/d.

Proof. Commutation and simplicity imply that the positive left Perron vector also has eigenvalue f under F. In particular $f(t_0)$ is rational. Choose a rational polynomial a_0 positive on I, with $a_0(t_0)f(t_0) = 1 - c$, $c = 257$. A constant works when its required value is positive; otherwise use a positive multiple of $h_0 - t$, with $\max I < h_0 < t_0$. Put

$$d_K = 1 + K(t - t_0)^2, \qquad B = d_K D, \qquad A = cd_K D + a_0 DF. \tag{20}$$

At t_0 the two matrices coincide and have equal positive entries. In a fixed neighborhood of that point positivity holds already by continuity; on its compact complement the K-term dominates all other entries for large K. Increasing K also makes $A > B$ entrywise on I. Strict polynomial positivity tests on compact intervals find such an integer. The shared Perron roots have ratio $c + a_0 f/d_K$, proving the sign assertion and $R > 1$.

Choose a common Bernstein degree n making all coefficients of all entries positive. For an entry $\sum_j c_j t^j$, its coefficient at position i is

$$\sum_{j \le i} c_j \frac{\binom{i}{j}}{\binom{n}{j}}.$$

These converge uniformly to the polynomial at i/n: sampling j indices with and without replacement has discrepancy at most $j(j-1)/(2n)$. Strict positivity on the compact interval therefore makes degree elevation terminate. This is the univariate Bernstein form of Pólya's 1928 positivity theorem [35]. Substitute $t = z/(1+z)$ and multiply every entry by the same $L_0(1+z)^n$, with integer L_0 clearing the rational coefficients. The resulting positive integer coefficients encode multiplicities of energies i/d. At $z = 1$ all entries of both matrices have the same integer sum, since $t = t_0$. Hence the common index alphabet realizes every pair and equality class by a list of exactly that length. Common multiplication changes no Perron ratio. □

10.4 Sign operations and individual conjugates

For orientation, we first give a direct finite-dimensional sign construction. A selected scalar f is represented by a pair (D, F) satisfying Lemma 10.2. Tensoring carriers and lifting value matrices supports addition, negation, rational scaling and multiplication of selected scalars. Rational polynomials have one-dimensional carriers $D = 1$.

Suppose selected scalars a, b have no common zero on $[0, 1]$ and are nonzero at t_0. Rescale them positively to have absolute values $3, 4$ at t_0, and use their commuting lifts F, G on a tensor carrier D. Put

$$H = F^2 + G^2, \qquad S_0 = \begin{pmatrix} 0 & H/5 \\ 5\,\mathrm{Id} & 0 \end{pmatrix},$$
$$D' = \mathrm{diag}(D^m, D^m)(5\,\mathrm{Id} + S_0),$$
$$F' = \mathrm{diag}(F + G, F + G) \pm S_0.$$

For sufficiently large m, D' is positive on $[0, 1]$: normalized powers of the positive D converge uniformly to its Perron projection, and $D^m H$ has the strictly positive selected scalar $a^2 + b^2$. Compactness makes the convergence and lower bound uniform. At t_0 all blocks coincide with $5D^m$, because $a^2 + b^2 = 25$, so all entries of $D'(t_0)$ agree. The matrices commute, and the new Perron scalar is

$$a + b \pm \sqrt{a^2 + b^2}.$$

The plus operation is positive exactly when at least one input is positive; the minus operation is positive exactly when both are positive. Zeros can occur only at input zeros. These are Rvachev's R-function disjunction and conjunction [38], here implemented on a positive rational polynomial carrier.

For an algebraic cut $r \in I$, let p_r be its irreducible rational polynomial, oriented positive just to its right. Choose rational $l < r < s < 1/2$, with no other root in $[l, s]$, $p_r(l) < 0 < p_r(s)$. Take the minus operation on $t - l, p_r(t)$, and the plus operation on its result and $t - s$. Both pairs have no common zero. The selected result g_r is strictly negative for $t < r$, zero at r, and strictly positive for $t > r$, throughout $[0, 1]$. Below l both outer inputs are negative; between l and s the only relevant polynomial zero is r; at and above s the first, then the second, outer input is positive. Other real conjugates are therefore not forced to be zeros on the selected branch.

After decomposing the explicit open diagram on I, multiply one such factor at each crossing, squared at an isolated excluded point, and choose an overall sign so that $f < 0$ on

the desired warm set. Constants handle a full or empty diagram. This constructs a finite carrier with $f(t_0) \neq 0$, hence a physical magnet by Lemmas 10.2 and 10.1. The next argument reduces its line dimension to at most three.

10.5 Three states suffice

Lemma 10.3 (Dimension-three synthesis)**.** *For every input of Theorem F there is a pair (D, F) as in Lemma 10.2, of dimension at most three, whose selected scalar is negative exactly on the required warm portion of I.*

Proof. For each irreducible rational polynomial with a boundary of the diagram in I, classify its real roots in I by the required order of a warm-negative analytic scalar: a crossing has positive odd order (type 1); an isolated cold point between warm intervals has positive even order (type 2); a warm interior point has order zero (type 0); a cold interior point may use type 0 or 2. Choose permitted types and let $k \leq 3$ be the maximum number of types needed for any polynomial, with $k \geq 1$. Choose I slightly larger than the window image without adding boundary roots near its ends. If there are no boundaries, take a constant one-dimensional sign.

Let P be the product of the distinct irreducible boundary polynomials. For each factor p, assign exponents $e_{j,p} \in \{0, 1, 2\}$, $1 \leq j \leq k$, covering its chosen types, repeating an exponent when needed. At each real root ξ of P in $[0, 1]$, choose a label $j(\xi)$ whose exponent is the required type in I; outside I choose any available label. Put

$$B_j = \prod_p p^{e_{j,p}}.$$

There are no roots at $0, 1, t_0$, since a boundary root is in $(0, 1/2)$ and a rational root has a linear minimal polynomial.

Choose nonnegative rational polynomials h_j, zero at t_0, with $h_{j(\xi)}(\xi) > 0$ and strictly larger than all other $h_i(\xi)$. Approximate separated continuous bump values at these finitely many roots by rational polynomials, square them, and multiply by $(t - t_0)^2$. Strict algebraic comparisons certify a terminating search. The symmetric polynomial matrix

$$L = \operatorname{diag}(h_1, \ldots, h_k) + P^4 \mathbf{J}$$

is positive semidefinite. Off the roots it is entrywise positive. At a root it has a unique positive top diagonal entry. Its top eigenvalue λ is simple everywhere on $[0, 1]$, and its continuous unit top vector v is nonnegative, strictly positive off the roots, and equal to the designated coordinate vector at a root. All other eigenvalues have a uniform ratio less than one to λ. Near a root with label j, the Rayleigh principle gives $0 \leq \lambda - h_j \leq kP^4$.

Define commuting value matrices and their selected scalars by

$$W_j = B_j \prod_{i \neq j} (L - h_i \operatorname{Id})^2, \qquad w_j = B_j \prod_{i \neq j} (\lambda - h_i)^2.$$

At ξ of type e and label j, the scalar w_j has exact order e; every other w_i is $O(|t - \xi|^8)$. Off the roots no w_j vanishes. Let f_0 be the real polynomial formed from the assigned factors $(t - \xi)^e$, with sign giving the desired diagram on I. It is an existence device, not an entry of the rational table. The coefficients

$$a_j = \frac{f_0 w_j}{\sum_i w_i^2}$$

extend continuously through every root, and $f_0 = \sum_j a_j w_j$. Moreover $\sum_j |w_j|/|f_0|$ is bounded on the punctured interval, including near the roots. Sufficiently close uniform rational

polynomial approximations b_j to a_j therefore give $f = \sum_j b_j w_j$ with the same signs, zeros and orders as f_0. The polynomial matrix $F = \sum_j b_j W_j$ commutes with L and selects f. A search for rational tuples can test these properties by real quantifier elimination, specifying λ as the simple top eigenvalue; existence proves termination. In particular $f(t_0) \neq 0$.

It remains to make the top eigenvectors strictly positive at the roots. Choose a rational skew-symmetric polynomial matrix K, zero at t_0, with $K_{i,j(\xi)}(\xi) > 0$ for $i \neq j(\xi)$. Polynomial approximation of the finitely many consistent assignments supplies it. For a sufficiently small positive rational ϵ, set $S = \mathrm{Id} + \epsilon K$. Skew symmetry gives $\det S > 0$ for real t. Both Sv and $S^{-\mathsf{T}}v$ are strictly positive on $[0,1]$: near a formerly zero coordinate their first-order term is $\epsilon(Kv)_i > 0$, with a uniform $O(\epsilon^2)$ remainder for the second vector; away from these neighborhoods the original coordinates have a positive minimum. At $k = 1$ use $K = 0$.

Set $L' = SL\,\mathrm{adj}(S)$ and $F' = SF\,\mathrm{adj}(S)$. Their selected eigenvalues are scaled by $\det S > 0$, and they commute. The top projection of L' has strictly positive entries and a uniform spectral gap. Some power $D = (L')^m$ is therefore entrywise positive throughout $[0,1]$, by uniform convergence of normalized powers. At t_0, $S = \mathrm{Id}$ and $L = P(t_0)^4\mathbf{J}$, so all entries of $D(t_0)$ agree. Searches for ϵ, m use strict polynomial positivity tests. The pair (D, F') has the required dimension and sign. □

Realization part of Theorem F. Apply Lemma 10.3, convert the matrices by Lemma 10.2, and use their equal-length lists in the nearest-neighbor magnet. Lemma 10.1 gives EWM exactly at the negative selected scalar and order with gap $1/4$ at every other point, including its zeros. The color half-turn and observable are those already specified. All data are finite and rational and every search is uniform from the explicit semialgebraic description. Only the finite edge-index length and the energies, rather than the bound of three line states, depend without a fixed size bound on the input. □

10.6 Exact periodic weights and the dimension lower bound

The sharpness statement allows a wider line class than common Perron directions. For positive matrices T_e, put $\lambda_e = \rho(T_e)$ and $Q_e = T_e/\lambda_e$. The *periodic-word condition* is

$$\rho(Q_{e_1} \cdots Q_{e_n}) = 1 \quad \text{for every nonempty word } e_1 \cdots e_n. \tag{21}$$

It says that every periodic color word has exactly the product of its individual Perron weights in its bulk free energy. With identical horizontal and vertical matrices this condition is necessary for those exact weights on all periodic base color configurations: realize any binary equality word along a row by toggling two colors at disagreements, repeat the word twice if needed, and repeat the row vertically. The uniform-color input fixes the pattern-independent pressure correction to zero. The condition concerns exact periodic weights, rather than phase boundaries alone.

Lemma 10.4 (Periodic weights without common directions)**.** *For positive k-dimensional generators,* (21) *is equivalent to*

$$l^{\mathsf{T}} Q_w r = 1 \quad \textit{for all words } w,$$

where r, l are the Perron vectors of one generator, normalized by $l^{\mathsf{T}} r = 1$. It suffices to test words of length at most k. For equality/inequality Potts line decorations satisfying this condition and $R = \lambda_1/\lambda_0 \geq 1$, the conclusion of Lemma 10.1 still holds. Moreover $\lambda_1 - c\lambda_0$ is an eigenvalue of the rational matrix $T_1 - cT_0$, for every rational c.

Proof. From (21), let the power of the reference generator in $Q_0^n Q_w$ tend to infinity. Its limit is $r l^{\mathsf{T}} Q_w$, giving the scalar identity by continuity of spectral radius. Conversely apply the

identity to powers of Q_w and take roots; positive r, l, Q_w identify its Perron growth. For the finite test use $\mathrm{diag}(Q_e, 1)$ on $(r, 1)$. Its word span, starting at dimension one, stabilizes by length k in dimension $k+1$. The output $(l^{\mathsf{T}}, -1)$ annihilates the full span if it annihilates all these short words.

Here is the boundary estimate needed to retain the physical conclusion. Let m_1, M_1 be the minimum and maximum normalized entries, and $\theta = 1 - m_1/M_1 < 1$. For a product $C = Q_1 \cdots Q_n$, $2h \le n$, let a be the normalized column $Q_1 \cdots Q_h \mathbf{1}$ and b^{T} the normalized row $\mathbf{1}^{\mathsf{T}} Q_{n-h+1} \cdots Q_n$, each with coordinate sum one. Uniformly over words,

$$C_{ij} = \frac{a_i b_j}{b^{\mathsf{T}} a} \, [1 + O(\theta^h)]. \tag{22}$$

To verify it, set $\delta = m_1/(kM_1)$. All components of a, b are at least δ. Reverse transitions associated with a positive row message u have probabilities $u_i Q(i,j)/\sum_s u_s Q(s,j)$, minorized by $(m_1/M_1) u_i / \sum_s u_s$. Thus the first h transitions contract arbitrary end laws by θ^h; reversing directions gives the analogous row bound. With $S_C = \sum C_{ij}$, these estimates give $C_{ij}/S_C = a_i b_j + O(\theta^h)$. Apply this matrix to a and use the positive-vector bounds on its spectral radius. By (21) that radius is $1/S_C$; hence $1/S_C = b^{\mathsf{T}} a + O(\theta^h)$. The lower bound δ converts the estimates to the relative form displayed. If $m_1 = M_1$, the normalized matrices are the same rank-one matrix and the formula is immediate.

The same two-ended contraction defines a unique hidden conditional law K_σ for any fixed color configuration. It excludes hidden correlations between distinct lines at infinity by uniform convergence on every finite collection of line segments. For a finite color change $\sigma \to \sigma'$, let G be the product of changed indexed-weight ratios. Enclosing each changed segment in longer intervals, (22) cancels their identical word collars, and gives

$$\mathbb{E}_{K_\sigma} G = \prod \lambda_{\mathrm{new}}/\lambda_{\mathrm{old}}, \qquad K_{\sigma'} = G K_\sigma / \mathbb{E}_{K_\sigma} G.$$

These identities give Gibbs lifting by finite-change ratios as before. Disintegration and uniqueness of K_σ give the converse projection onto the Potts specification.

For mixing, use the inflated domain D of the proof of Lemma 10.1. Apply (22) to its long runs with collar length $\lfloor h/3 \rfloor$. The leading factor depends only on the two word collars; condition their colors. The remaining reference color law is ordinary Potts. Its EWM couples the inner dilation, while the two-ended contraction couples all tested hidden segments. The summed run error is $O(|E_0|(4h+1)^2 \theta^{\lfloor h/3 \rfloor})$; outgoing indices then couple exactly. This is exponential in the distance, uniformly over arbitrary exteriors and holes. As in Section 8, its cardinality form implies (3). Ordered Potts states lift with the same observable. This proves the phase criterion for the wider class. Constant-spectral-radius semigroups are studied in [39]; the short-word test above is the positive-generator specialization needed here.

Finally put $D_0 = \mathrm{span}\{Q_w r - r\}$ and $V = \mathrm{span}\{Q_w r\}$. The scalar identity gives $D_0 \subset \ker l^{\mathsf{T}}$, $V = D_0 \oplus \mathrm{span}(r)$. Both spaces are invariant under every generator, which acts as the identity on V/D_0. On this quotient $T_1 - cT_0$ acts by $\lambda_1 - c\lambda_0$. This proves the eigenvalue assertion even when the invariant spaces vary with temperature. □

Proposition 10.5 (A diagram requiring three line states)**.** *In the inverse-temperature window* $[1/2, 3]$*, put* $z = e^{-\beta}$ *and let* $z_1 < z_2 < z_3$ *be the three positive roots of*

$$p(z) = (31z - 5)(31z - 10)(31z - 15) + 5.$$

The open diagram $z > z_1$*,* $z \neq z_2$*, requires at least three internal line states in the identical-orientation periodic-word class with ferromagnetic Perron ratio, allowing arbitrary finite edge-index lists and any rational energy denominator.*

Proof. For $H = T_1 - cT_0$, the scalar $f = \lambda_1 - c\lambda_0$ is analytic because positive matrices have simple Perron roots, and is an eigenvalue by Lemma 10.4. For a nonconstant diagram it is not identically zero. In a rational Boltzmann coordinate clear negative powers and write

$$\det(Y\,\mathrm{Id} - H(t)) = \sum_{i=0}^{k} a_i(t) Y^i.$$

At real algebraic conjugate points ξ, the finite orders v_i of the nonzero rational coefficients a_i are the same: they are multiplicities of the common irreducible polynomial. If the selected eigenvalue has finite order s, substitution of its leading term forces the minimum of $v_i + is$ to be attained at least twice. The lower envelope of lines of slopes $0, \dots, k$ has at most k distinct breakpoints. Thus there are at most k possible distinct finite orders across a conjugacy class, even if unused eigenvalues vanish identically. This is the elementary Newton-polygon argument.

For the prescribed diagram, z_1 is a crossing and requires positive odd order, z_2 is an isolated cold point and requires positive even order, and z_3 is warm interior and requires order zero. These are three different orders. The roots lie respectively in $(4/31, 6/31)$, $(9/31, 11/31)$, and $(14/31, 16/31)$, all in the window image. For every integer $d' \geq 1$, $p(X^{d'})$ is Eisenstein at 5: its leading coefficient is 31^3, the other coefficients are divisible by 5, and the constant is -745, not divisible by 25. The three positive d'-th roots remain conjugate. The order argument therefore excludes $k \leq 2$ at every rational energy denominator. Lemma 10.3 supplies $k = 3$, completing the sharpness part of Theorem F. □

10.7 An explicit two-state example

The following small list illustrates conjugate selection without the general synthesis. Use $k = 2, L = 3$, with the same table on the two orientations:

Internal tail, head	Equal colors	Different colors
$1,1$ or $2,2$	$(0,0,0)$	$(1,1,1)$
$1,2$	$(1,1,1)$	$(1,1,1)$
$2,1$	$(1,1,2)$	$(1,1,2)$

With $z = e^{-\beta}$, its unequal and equal matrices are

$$B = \begin{pmatrix} 3z & 3z \\ 2z + z^2 & 3z \end{pmatrix}, \qquad A = B + 3(1-z)\,\mathrm{Id}, \qquad R = 1 + \frac{3(1-z)}{3z + z\sqrt{3(2+z)}}.$$

The ratio is strictly decreasing in z. Its equation $R = 257$ is $1 - 257z = 256z\sqrt{(2+z)/3}$. Squaring gives $(1-z)(3 - 1539z + 65536z^2) = 0$. For the two roots $z_- < z_+$ of $65536z^2 - 1539z + 3$, only $z_- < 1/257$ satisfies the unsquared equation. Both roots lie in (e^{-7}, e^{-3}). This one specified magnet therefore has EWM on $[3, -\log z_-)$ and order with gap $1/4$ on $[-\log z_-, 7]$, including the transition; the warmer conjugate z_+ is mixing. Its alphabet has $65536 \cdot 2^2 \cdot 3^2$ symbols and its edge energies are $0, 1, 2$.

10.8 Sizes and the ceiling of the mechanism

The synthesis is effective, with all searches terminating; no elementary size bound is claimed. A bound of three concerns only the internal state on a line. The polynomial degrees, rational energy data and edge-index lists can be extremely large. Their size cannot stay bounded even for a single rational cut. Indeed, with k line states and N alternatives per edge,

every entry of $H = T_1 - 257T_0$, after a common energy shift, has coefficient one-norm at most $258N$ and at most $2N$ monomials. If ℓ is the least index of a nonzero coefficient in $\det(Y\,\mathrm{Id} - H) = \sum a_i(X)Y^i$, each boundary zero of f is a zero of a_ℓ: divide the characteristic identity by f^ℓ and extend across its isolated zeros. The principal-minor expansion bounds the monomial count of a_ℓ by $k!(2N)^k$ and its coefficient one-norm by $k!(258N)^k$. Descartes' rule bounds the number of positive boundary roots, counting an isolated cold point at least twice, by $k!(2N)^k - 1$.

For example, for $u = 4j+2$, $j \geq 1$, prescribe the cut $z = u/(3u+1)$ in the window $[1/2, 2]$. At denominator d', $(3u+1)X^{d'} - u$ is Eisenstein at 2. It divides a_ℓ over the integers, so its leading coefficient gives the lower bound

$$3u + 1 \leq k!(258N)^k.$$

Thus the finite local data grow without bound even when only one cut is requested. Large energies alone cannot remove this cost in this class.

For any fixed rational energy lists, their transfers are Laurent polynomials over $\mathbb{Q}$ in a common $X = e^{-\beta/d'}$. Positive Perron roots are selected by equations with strictly positive eigenvectors. Consequently $R(X) < 257$ is a rational semialgebraic condition by real quantifier elimination. It is open by continuity. In the line class of Lemma 10.4, this strict inequality is exactly the EWM region, and all its finite interior boundary values of X are algebraic.

We only sketch how finite line tables also implement algebraic temperature substitution; the sketch is not used in Theorems D–F. For a rational base energy table, write $n_h = d_0(E_h - \min E_h) \in \{0, \ldots, N_h\}$. Along one row chain for each table type, use transfers $A^{n_h} B^{N_h - n_h}$, with common positive Perron directions. On a range with $0 < \lambda_A/\lambda_B < 1$, their selected weights give the positive base inverse temperature

$$\beta' = -d_0 \log(\lambda_A/\lambda_B).$$

The finite-change cancellation in Lemma 10.1 identifies its Gibbs states with the base states at β'. The inflated-domain proof, with a fixed enlargement for each base stencil, transfers base EWM to the full spins. Conversely a base finite-volume kernel with fixed base exterior, extended by its hidden conditional law, is a mixture of full decorated kernels; full EWM bounds the projected mixture uniformly and hence gives base EWM. Thus this operation pulls the diagram back through an algebraic map in its exponential coordinate. Applied to the semialgebraic diagrams just constructed, it leaves them semialgebraic.

The resulting exp-algebraic transition values are computable. This ceiling concerns the finite rational line-transfer construction and algebraic substitutions of its diagrams; it is not an upper bound on diagrams of general magnets. A substitution into an arbitrary base magnet does not itself supply an algebraic description of that base magnet's mixing set. In particular, these constructions do not resolve the single-magnet problem stated in the discussion.

11 A single magnet with a prescribed response-sign set

Theorems D–F concern the set where the whole specification mixes. We now prescribe a different set while keeping that specification mixing throughout the window: the inverse temperatures where a positive field at one site induces a negative moment somewhere else. The construction uses a small signed Fourier perturbation of an independent update. The computation takes place in its exact moment expansion; the stationary hidden row is independent and uniform.

11.1 Polynomial signs from an enumeration

Put $z = e^{-\beta}$ and $\mathcal{K} = [e^{-b}, e^{-a}]$. A computable change of coordinates transports an enumeration of W to an enumeration of rational open intervals in $(0,1)$, with union V satisfying

$$z \in V \cap \mathcal{K} \quad \Longleftrightarrow \quad -\log z \in W.$$

For example, enumerate rational intervals whose closures are certified inside the image of an enumerated ambient inverse-temperature interval. These cover that image, including neighborhoods of either endpoint of the window. Fix a tape enumerator for V. In its first j steps it completes at most j intervals, with all printed integers of binary length at most j, after a fixed change of the time unit.

Merge the intervals completed by date j, but keep a missing point when two open intervals merely touch. If the components are $(p/q, r/s)$, with positive denominators, define

$$P_j(z) = \prod_{(p/q,r/s)} (qz - p)(sz - r), \qquad Q_j(z) = 1 + 2^{j^2} P_j(z), \qquad j \geq 1. \tag{23}$$

The empty product is one. At an interior point of a component exactly one pair of factors has negative product; outside the finite union all pairs have nonnegative product. Thus $P_j < 0$ exactly on that union. In particular

$$z \notin V \Rightarrow Q_j(z) \geq 1 \text{ (all } j), \qquad z \in V \Rightarrow Q_j(z) < 0 \text{ (all sufficiently large } j). \tag{24}$$

For the second assertion, choose an enumerated neighborhood of z of radius $d > 0$. Every endpoint of every later component is at least d from z. There are at most $2j$ integer linear factors, each of absolute value at least $\min\{1, d\}$. Hence eventually

$$P_j(z) < 0, \qquad |P_j(z)| \geq \min\{1, d\}^{2j},$$

which proves (24). Also, $Q_j(z) < 0$ at any stage implies $z \in V$.

11.2 The finite dual table

We give a finite specification by message types and by the instructions of an ordinary one-ended tape machine. Its finite tape alphabet and control are effectively obtained from the given enumerator. This specifies every entry of the table, including entries not reached from the initial message.

Let $G = (\mathbb{F}_2)^w$, with real characters $\chi_c(y) = (-1)^{c \cdot y}$. Partition the coordinates of a dual index c into three equal bundles. A message sent to offset $d \in \{-1, 0, 1\}$ uses only the bundle reserved for arrival from offset $-d$. Within each bundle reserve separate resident, head and special-code fields, and one parity coordinate. Choose enough further coordinates for the following finite payloads:

- resident symbols $\mathsf{R}(\gamma)$, for γ in a finite tape alphabet Γ, including an allocated blank, counter tracks, and left and right delimiter flags;
- heads $\mathsf{H}(q, i)$, where q is a common control state and $i \in \{\mathsf{b}, \mathsf{v}, \mathsf{u}\}$ is a linear index;
- the distinct special codes A, B_0, I, and stationary inert codes $\mathsf{J}_0, \mathsf{J}_1$.

Every resident has parity zero and is nonzero; every head and each of A, B_0, I, J_1 has parity one. The nonzero code J_0 has parity zero. The codes $A, B_0, I, \mathsf{J}_0, \mathsf{J}_1$ use the self bundle. A resident and a stationary head use their separate fields in the same self bundle, so their sum decodes uniquely. The special-code field distinguishes $A, B_0, I, \mathsf{J}_0, \mathsf{J}_1$ from all tape payloads. The parity of an arbitrary index is the sum of its three parity coordinates.

An option consists of a weight $k(z) \in \mathbb{Z}[z]$ and three outgoing indices. An absent option has coefficient zero. The empty index has the unique empty output with coefficient one. For nonzero indices use the following rules; an output to self always uses the self bundle.

Input type	Scattering options before the common factor α
A	Output to self the initial head $\mathsf{H}(q_{\mathrm{init}}, \mathsf{b})$ and the allocated origin resident, with both delimiter flags; weight one.
Resident only	Send the same resident to self, with weight one.
One head and its resident	Read the resident as the tape symbol. For an instruction $(q, \gamma) \mapsto (q', \gamma', d, L)$, and each nonzero entry $L_{i'i}(z)$, write $\mathsf{R}(\gamma')$ to self and send $\mathsf{H}(q', i')$ to offset d, with weight $L_{i'i}(z)$. If $d = 0$, combine those two payloads in the self bundle.
An allocating head alone	Read an unallocated blank, install its resident, and apply the same instruction rule. Only the initial and right-extension control states permit this case.
Stage fork	For index b, retain the resident and enter the next-stage control, with weight one. For index v, retain it and enter the delay control, with weight one. For index u, there is no option.
Delay control	Either perform the next counter instruction by the ordinary rule with $L = \mathrm{Id}$, retaining an allocated resident, or retain the current resident and enter a fixed quit-setup state; all weights are one. At the last counter state only the quit option is present.
Quit and seek-right control	Retain the resident and move the head right until it reads the right delimiter, then enter erase control; all weights are one.
Erase control, away from the left delimiter	Delete the current resident and send only the erase head to the left neighbor, with weight one.
Erase control, at the left delimiter	Output B_0 to self immediately, or retain the resident with the wait head; both weights are one.
Wait head with its resident	Send B_0 to self, with weight one.
B_0, I	Send I to self, with weight one.
J_p, $p = 0, 1$	Send J_p to self, with weight one.
Every other nonzero tuple	Send J_p to self, where p is its total parity, with weight one.

The instruction rule applies to a head arriving in any one bundle, with a resident in the self bundle, and also to their stationary combination. All other combinations, including multiple heads, have the last rule. The delimiter tests in the specialized controls take precedence over the ordinary instruction rule. The wait state is not an erase state, so the delimiter test does not apply to it. A zero column of L gives an empty option list, rather than an option with empty output.

Here is the common tape program and its finite collection of linear instructions. It computes the first j steps of the fixed enumerator from scratch, merges the completed intervals by integer arithmetic, and issues the factors in (23). The linear amplitudes are the column vector (b, v, u). Besides the identity, the required maps are

$$\begin{aligned} &(b, v, u) \mapsto (b, b, 0), \quad (b, v, u) \mapsto (b, v, 0), \\ &(b, v, u) \mapsto (b, v, u + zv), \quad (b, v, u) \mapsto (b, v, u - v), \\ &(b, v, u) \mapsto (b, u, 0), \quad (b, v, u) \mapsto (b, 2v, u), \quad (b, v, u) \mapsto (b, v + b, u). \end{aligned} \tag{25}$$

At the start of each stage apply $(b, v, u) \mapsto (b, b, 0)$. For each factor $qz - p$, apply $(b, v, u) \mapsto (b, v, 0)$, issue q additions of zv and p subtractions of v, and then apply $(b, v, u) \mapsto (b, u, 0)$. After the last factor perform j^2 doublings $(b, v, u) \mapsto (b, 2v, u)$ and apply $(b, v, u) \mapsto (b, v+b, u)$. Starting from the baseline alone, $(1, 0, 0)$, this gives $(1, Q_j, 0)$. The coefficients used by the finite table belong to $\{1, -1, z, 2\}$. The long integer loops are tape computations, and do not enlarge the physical radius.

All choices of linear index perform the same common instruction, write the same resident and move the same way. Before the fork they therefore occupy the same sites at each date. At the fork the baseline alone continues, so the next stage has no signed carry from the preceding polynomial. Arrange each fork with the head at the right delimiter of one contiguous allocated segment $[0, D_j]$. The baseline visits both ends and takes two extra idle steps before the next fork. Extension of the right end removes the old right flag and uses an allocating head to install the next blank with the new flag. No instruction writes the left flag; only the origin resident carries it, until cleanup. These operations are ordinary read, write and unit-move instructions on finite tracks; they are independent of the linear index.

Before the stage-j fork prepare a counter track of 2^j bits. The delay branch counts through all its binary values, without allocating any further cells, and offers the quit option at every machine tick. The counter can be implemented by a binary increment scan between marked ends, and stops on overflow. The segment includes all computation and counter tracks. For constants depending only on the finite enumerator, its length and the time between successive baseline forks are at most

$$D_j + 1 \le C 2^{cj}, \qquad F_{j+1} - F_j \le C 2^{cj}. \tag{26}$$

Indeed, a j-step simulation has at most j completed outputs, each printed integer is at most 2^j, binary sorting and merging are polynomial in j, and the unary factor loops cost at most a polynomial in 2^j. Doubling the allocated counter track and sweeping the tape have the same bound on a one-tape machine. Each stage terminates even if the enumerator prints nothing. This gives an effective finite program, using only the message types in the table.

Lemma 11.1 (Support, parity and terminal coefficients)**.** *The finite table above has the following properties.*

(i) *Every option of a nonzero index has nonempty output, and every option of every nonzero index, reachable or not, preserves total parity.*

(ii) *Starting from B_0, no positive-time expansion contains an A or B_0 singleton. Starting from A, the only singleton that can pair with an old A or B_0 is B_0 at the origin, produced by cleanup at a stage fork.*

(iii) *The possible terminal dates of stage j form a finite integer interval $\mathcal{T}_j = [f_j, \ell_j] \cap \mathbb{N}$. Its contribution at date $n \in \mathcal{T}_j$ is*

$$Q_j(z) c_{j,n}(z), \qquad c_{j,n}(z) = \sum_{\gamma} d_\gamma \alpha(z)^{m_\gamma} > 0 \quad (0 < z < 1), \tag{27}$$

where the sum is finite, d_γ are positive integers, and m_γ counts active scatterings along the common computation and the specified quit and wait path. Outside $\mathcal{T}_j$ that stage contributes zero. The intervals cover every sufficiently large integer, and $f_j \to \infty$.

Proof. Every ordinary moving head leaves a nonzero resident or sends a nonzero head, and a stationary head remains nonzero. Erasure sends its head until the final nonzero B_0. Resident and inert rules are nonempty. An empty linear-index list contributes zero, not a disappearing branch. Residents have parity zero, heads have parity one, and every listed tape operation replaces one head by one head or by B_0. The other rules preserve parity by their definition. This proves (i).

At a predecessor site, incoming indices from distinct senders occupy disjoint direction bundles. Their sum is a concatenation, so a nonzero message cannot cancel. In particular, a nonempty shape remains nonempty at every contributing step. Induction from A gives one head on a contiguous allocated segment until cleanup, or one finished inert token afterwards. There are no malformed tuples on such a path. A deleted rightmost cell has no other sender and stays empty. All remaining residents retain their position and their payload until visited by the unique head. No output rule produces A; the only rule producing B_0 is the final erasure at the marked origin. The special-code field prevents a tape tuple, or a merger of messages in distinct bundles, from becoming A or B_0. An inert token cannot turn into either. From B_0 there is just I at every subsequent date. Orthogonality against an old A or B_0 thus selects exactly the terminal singleton specified in (ii). Intermediate single-site tape shapes, when present, have different character indices and contribute nothing to that pairing.

For a fixed stage, all linear paths have the same occupancy before the fork. Multiplication by the common α at every active site therefore scales the whole vector equally. Summing those paths gives the maps (25), hence the factor Q_j. The baseline reaching a later stage has only positive weight. After the fork the linear amplitudes have already combined into one v-head, and every delay, quit and erase weight is positive. This proves the coefficient formula, including its sign even when $Q_j = 0$.

For the dates, write t for a quit tick and $p_t \in [0, D_j]$ for its head position. The fixed quit-setup state makes the terminal date

$$d_t = t + (D_j - p_t) + D_j + C_0,$$

or $d_t + 1$ with the private wait. Adjacent ticks satisfy $|p_{t+1} - p_t| \le 1$, so $d_{t+1} - d_t \in \{0, 1, 2\}$. The two available dates at each tick fill every gap. The finite loop has at least 2^{2^j} ticks, giving

$$f_j \le F_j + 2D_j + C_1, \qquad \ell_j - f_j \ge 2^{2^j} - D_j - C_1. \tag{28}$$

Here F_j is the fork date, including its fixed setup. Since every terminal date of stage j is after its fork, $f_j \ge F_j \to \infty$. With $f_j \ge F_j$, equations (26)–(28) give $f_{j+1} \le \ell_j + 1$ for all large j. Since $\ell_j \to \infty$, every integer at least f_{j_0}, for a sufficiently large fixed j_0, lies in some $\mathcal{T}_j$. At a fixed date only finitely many stages, finite shapes and paths occur. This proves (iii) and rules out unintended terminal or merged contributions. □

11.3 A positive PCA and a rational energy table

For each nonzero dual index c, denote its option list by $\mathcal{L}(c)$, and its outgoing messages by $b_d(c, l)$. Define the hidden transition

$$K_z(y \mid h) = \frac{1}{|G|}\left[1 + \alpha(z)\sum_{c\neq 0}\chi_c(y)\sum_{l\in\mathcal{L}(c)}k_l(z)\prod_{d=-1}^{1}\chi_{b_d(c,l)}(h_d)\right], \qquad \alpha(z) = \frac{1-z}{R_0(1+z)^B}. \tag{29}$$

Character orthogonality gives row sum one and gives precisely the prescribed conditional moment for every nontrivial character, with the common factor α. Choose B at least the degree of every $(1-z)k_l(z)$. Increasing the integer R_0 ensures

$$K_z(y \mid h) \ge \frac{1-\epsilon}{|G|}, \qquad 0 \le z \le 1, \qquad 0 < \epsilon < 1/3. \tag{30}$$

It also makes every coefficient in every numerator, over the common denominator $|G|R_0(1+z)^B$, a positive integer: at every degree the background $R_0(1+z)^B$ grows with R_0, whereas the finitely many signed perturbation coefficients are fixed. Thus this choice is effective.

The iid uniform hidden row is invariant. For a nonconstant finite product of characters of a new row, conditional independence multiplies their expansions. Lemma 11.1(i), together with the disjoint arrival bundles, says that each contributing predecessor shape is nonempty. Its expectation under the old iid uniform row is zero. All nonconstant Fourier moments therefore vanish. More generally, the exact n-step expansion is the sum over its finite message diagrams, with weight equal to the product of the local k_l's times α to the number of active scatterings. The weights can be signed; only the physical kernel (29) is a probability.

Adjoin an actual sign $s \in \{-1, 1\}$ with conditional distribution

$$p_z(s \mid y) = \frac{1 + s\,\delta(z)(\chi_A(y) + \chi_{B_0}(y))}{2}, \qquad \delta(z) = \frac{1 - z}{R_1(1 + z)^{B'}}. \tag{31}$$

Choose $B' \geq 1$ and the integer R_1 large enough that $2\delta < 1$ on $[0, 1]$ and the numerator coefficients are positive integers, as above. Earlier signs are ignored. The product $K_z(y \mid h)p_z(s \mid y)$ has a common denominator $g_0(z) \in \mathbb{Z}[z]$ with positive coefficients and strictly positive integer-coefficient numerators. At $z = 1$ every projected symbol (y, s) has probability $1/(2|G|)$, independently of h. Consequently all numerator polynomials have the same sum of coefficients

$$L_0 = \frac{g_0(1)}{2|G|} \in \mathbb{N}.$$

List their monomials with multiplicity in increasing order of exponent, writing $e(h; y, s, \iota)$ for the exponent at position $1 \leq \iota \leq L_0$. The full alphabet is $\mathcal{B} = G \times \{-1, 1\} \times \{1, \ldots, L_0\}$. Its strictly positive radius-one PCA table is

$$T_z(y, s, \iota \mid h) = \frac{z^{e(h;y,s,\iota)}}{g_0(z)}. \tag{32}$$

One anchored interaction on the child and its three old hidden neighbors has energy $e(h; y, s, \iota)$. These are fixed nonnegative integers. The energy sum over a new full symbol is $g_0(e^{-\beta})$, independent of the parents. Dividing each weight by that number gives exactly (32), so the interaction and normalized PCA have the same specification, including factors of exterior children.

Let $g \in G$ be one on each of the three parity coordinates and zero elsewhere. Parity conservation gives

$$K_z(y + g \mid h + g) = K_z(y \mid h), \qquad p_z(-s \mid y + g) = p_z(s \mid y).$$

The corresponding numerator polynomials coincide, and therefore so do their sorted exponent lists. Hence

$$\sigma(y, s, \iota) = (y + g, -s, \iota), \qquad M(y, s, \iota) = S(y, s, \iota) = s$$

give an invariant fixed-point-free involution and an odd Ising-valued one-site observable. We have constructed one finite rational magnet in the sense of the introduction.

Forgetting holds for the full table. Couple two hidden updates through the uniform component of weight $1 - \epsilon$ in (30). A hidden discrepancy has probability at most ϵ times the sum of the three parent discrepancy probabilities, so after t updates its one-site probability is at most $(3\epsilon)^t$, from any two rows. Once the hidden parents agree, couple the full new symbol in (32) identically. A union bound gives (2) on every finite test set for the full PCA, with a harmless one-step change of constants. Theorem 7.1 now gives exactly one DLR measure μ_β, its stationary trajectory law, and arbitrary-volume EWM at every $\beta \in [a, b]$.

11.4 Response, onset and pressure

Conditional on the entire hidden field, the signs at different sites are independent with means $\delta(\chi_A + \chi_{B_0})$. If the lag is (k, n) with $n > 0$, orthogonality in the old iid row identifies $\langle \chi_p(Y_{(0,0)})\chi_c(Y_{(k,n)})\rangle$ with the coefficient of the singleton p at coordinate $-k$ in the n-step expansion from c at the origin. Lemma 11.1 thus gives

$$\langle S_0 S_{(k,n)}\rangle_\beta = \begin{cases} \delta(z)^2 \displaystyle\sum_{j:n\in\mathcal{T}_j} c_{j,n}(z)Q_j(z), & k = 0, \\ 0, & k \neq 0. \end{cases} \tag{33}$$

The common row has zero correlation at distinct sites. Translation invariance and interchange of the two spins give the same formula with $|n|$ for negative lags. In particular every off-column correlation is zero.

If $\beta \notin W$, every summand in (33) is nonnegative, and each available stage contributes strictly positively. The coverage date from Lemma 11.1 is therefore one N_0, independent of β, beyond which all column correlations are strictly positive off W. If $\beta \in W$, only finitely many $Q_j(z)$ are nonnegative. Each such stage has a finite last terminal date. Beyond the maximum of those dates and the coverage date, at least one stage contributes and all contributions are negative. This gives eventual strict negativity along the column. Conversely a negative sum implies a negative Q_j, which implies $\beta \in W$. These assertions all hold for the same magnet and observable.

Symmetry and uniqueness give $\mu_\beta(S_x) = 0$. A point field with dimensionless strength h changes the law to

$$\mu_{\beta,h}(d\omega) = \frac{e^{hS_0(\omega)}}{\mu_\beta(e^{hS_0})}\,\mu_\beta(d\omega), \qquad \mu_{\beta,h}(S_x) = \tanh(h)\langle S_0 S_x\rangle_\beta. \tag{34}$$

The identity follows by expanding $e^{hS_0} = \cosh h + S_0 \sinh h$. A bounded one-site perturbation maps DLR measures bijectively by this tilt, with the inverse tilt e^{-hS_0}; hence this is the unique Gibbs state with the field. Differentiating at zero gives the response $R_\beta(x) = \langle S_0 S_x\rangle_\beta$.

Estimate (9) bounds the absolute responses exponentially in the distance. Their signs distinguish the two sets even though the measured effects can be extremely small.

To make the onset statement precise, let $N(\beta)$ be the least integer $N \geq 1$ such that $R_\beta((0, n)) < 0$ for every $|n| \geq N$, for $\beta \in W$. If β_{b} is a relative boundary point of W, it lies outside W. For any integer L, choose $n \geq \max\{N_0, L\}$. The finite-lag function in (33) is a finite sum of rational functions of z, hence continuous, and is positive at β_{b}. It stays positive in a relative neighborhood of β_{b}. Every β of W in that neighborhood has $N(\beta) > n \geq L$. Thus $N(\beta)$ tends to infinity on approach from W to any relative boundary point; in particular it is unbounded in every such neighborhood.

Finally take square volumes Λ. First retain just the energy terms anchored at children inside Λ, with the exterior fixed. Summing from the top row down, the parent-independent normalization gives exactly $g_0(z)^{|\Lambda|}$. The full Hamiltonian adds only terms anchored at exterior children whose parents meet Λ; there are $O(|\partial\Lambda|)$ of them, with uniformly bounded energy. On the compact inverse-temperature window their weights are between $e^{-C|\partial\Lambda|}$ and $e^{C|\partial\Lambda|}$. Therefore the pressure, with counting measure on the full alphabet, is

$$p(\beta) = \lim_{\Lambda\uparrow\mathbb{Z}^2} \frac{\log Z^\tau_\Lambda(\beta)}{|\Lambda|} = \log g_0(e^{-\beta}). \tag{35}$$

Explicitly $g_0(z) = 2|G|R_0R_1(1 + z)^{B+B'}$. Every local expectation is a rational function of z with denominator a power of g_0: start from the iid hidden row below its finite backward

cone and sum the finitely many transition factors. All fixed-lag correlations are therefore real-analytic; the noncomputable sign boundary comes from the quantifier over lags. The polynomial g_0 is positive on the real window. The pressure is real-analytic there, including on a real neighborhood of either endpoint. There is one DLR measure and EWM everywhere in the window, and the pressure has no singularity. Thus the response-sign boundary gives neither coexistence nor a thermodynamic singularity.

11.5 The effective converse

Let Φ now be any finite rational magnet with exactly one Gibbs measure at every $\beta \in [a, b]$. Its symmetry makes all one-spin means zero. For $x \neq 0$ put $f = M_0 M_x$. For any box Λ containing its support, consider the certificate

$$\max_{\tau} \gamma^{\tau}_{\Lambda,\beta}(f) < 0. \tag{36}$$

Only a finite boundary layer enters, so the maximum has finitely many terms. Every Gibbs measure is a mixture of these finite-volume laws; the certificate is sound. If $\mu_\beta(f) < 0$, uniqueness implies

$$\sup_{\tau} |\gamma^{\tau}_{\Lambda,\beta}(f) - \mu_\beta(f)| \longrightarrow 0$$

along growing boxes. Otherwise a sequence of bad boundary laws has a weak limit; finite range and positivity make that limit a DLR measure, contradicting uniqueness. This proves completeness.

Each finite-volume expectation is a uniformly computable analytic function of β, since the energies and M are rational. It has a computable continuity modulus on rational compact intervals, for example from the covariance derivative and the finite Hamiltonian range. Enumerate rational intervals on which interval enclosures certify (36), dovetailing over $x \neq 0$, boxes and precision, and intersect with $[a, b]$. Soundness and completeness give exactly

$$\{\beta \in [a, b] : R_\beta(x) < 0 \text{ for some } x \neq 0\}.$$

This is uniform in the finite table, its range, the observable and the rational window, under the stated uniqueness promise. Together with the construction it proves Theorem G. The converse concerns the derivative at zero field. It gives no effective-openness assertion for eventual negativity along the column, which has an additional existential–universal quantifier over the lags.

11.6 Disorder lines and fixed linear probes

A *disorder line*, in the correlation sense of Stephenson [46], separates different long-distance sign or oscillation patterns within a disordered regime. Fisher and Widom studied the analogous distinction between monotone and oscillatory decay in linear systems [47]. In a one-parameter window the locus can consist of individual points or a more complicated boundary. Every relative boundary point of W in Theorem G is such a locus for the eventual sign along the measurement column. For instance, taking $W = [a, a + (b - a)\Omega)$ puts this sign boundary at a noncomputable inverse temperature. Formula (35) and uniqueness distinguish this disorder line from the ordered/mixing boundary sought after Theorem F.

A *disorder variety* is instead a locus of Boltzmann weights admitting a reduction of the equilibrium model to lower-dimensional stochastic evolution. Enting studied crystal growth, Ising models and disorder points [48]. Peschel and Emery calculate two-dimensional Ising correlations from one-dimensional kinetic models [49], while Ruján obtains exact free energies and correlations on restricted order/disorder loci of competing-interaction IRF models [50].

The PCA-history/Gibbs correspondence is developed in [4, 14]. The parent-independent normalizer (32) places the entire temperature window of our magnet on such a variety within the ambient family of interactions on the same supports. The response-sign boundary is a subset of that window, rather than the locus of the reduction itself. Bernoulli invariant rows and independent directions in stationary PCA fields were studied by Mairesse and Marcovici [51].

There is also a useful distinction between the family of local responses and a fixed linear probe of it. In Section 11.3 choose R_0 so large that also, on a closed complex neighborhood U of $\mathcal{K}$ avoiding -1,

$$\max_{c\neq 0} \sum_{l\in\mathcal{L}(c)} |\alpha(z)k_l(z)| \leq \theta < 1.$$

For every nonempty dual shape, the absolute sum of outgoing weights for a full step is at most θ: the local sums multiply and at least one site is active. Nonextinction then gives a bound $4\sup_U |\delta|^2\theta^n$ for the absolute column correlation at lag n. Every finite-lag correlation is holomorphic there. This choice preserves (30), coefficient positivity and every sign above. For a fixed computable sequence (c_n) with a given polynomial growth bound, or a given bound $|c_n| \leq B_1 q^n$ with $q\theta < 1$,

$$F(z) = \sum_{n\geq 1} c_n R_{-\log z}((0,n))$$

is computable and holomorphic, with an effective tail bound. For bounded profiles it is the limit of the derivatives for the corresponding finite-support fields. Unless F vanishes identically, it has finitely many real zeros on $\mathcal{K}$, each individually computable by fixing an isolating interval and excluding compact subintervals away from its unique zero. This asserts individual computability, not a uniform zero-isolation algorithm for arbitrary input functions. Thus a fixed such linear response cannot retain a noncomputable boundary. The information lies in the quantification over lags. For a fixed finite transfer matrix the lag sequence instead satisfies a linear recurrence by Cayley–Hamilton; its sign question is a form of the Positivity problem studied by Ouaknine and Worrell [52].

12 Finite hidden Markov rows force exponential forgetting

The response construction has a stationary iid hidden row. Keeping that row while increasing the Fourier amplitudes cannot turn it into an ordered magnet as long as the PCA stays strictly positive. The following proposition gives the obstruction for every finite hidden Markov row on the line, and also covers finite decorations freshly sampled from a bounded hidden neighborhood, ignoring their own earlier values.

Proposition 12.1 (A finite hidden Markov stationary row)**.** *Let P be a translation-invariant synchronous PCA on $\mathbb{Z}$, with a finite alphabet, finite radius and strictly positive local transition probabilities. Suppose some spatially stationary, temporally invariant row law is the output of a stationary finite-state hidden Markov chain. Then P is uniformly exponentially ergodic in the full convention* (2)*. Its normalized space–time specification on $\mathbb{Z}^2$ has exactly one DLR measure and arbitrary-volume EWM* (3)*.*

Hidden transition and emission probabilities may be arbitrary real numbers. Forbidden hidden transitions, periods and several recurrent classes are allowed. The conclusion concerns the full PCA alphabet, rather than a selected readout.

For Bernoulli stationary rows, Marcovici and Taati prove exponential ergodicity in every spatial dimension, with a logarithmic mixing-time bound for finite regions [15]. Earlier

explicit examples include the noisy additive rule in [2, Example 2.3], whose analysis uses [21, Chapters 16–17]. Conditions for Markov invariant laws are studied in [21, 54]; Casse and Marcovici prove ergodicity for positive-rate memory-two PCA with their product-form invariant law [55, Theorem 8 in the preprint]. Our proof uses random erasure and the following external entropy theorem.

12.1 The external block-factorization theorem

For a probability law π on a finite product space write

$$\mathrm{Ent}_\pi(f) = \pi[f \log(f/\pi f)], \qquad f \geq 0.$$

Write $\pi[\mathrm{Ent}_U f]$ for the mean conditional entropy when the coordinates outside U are fixed. Thus, for $\pi f = 1$, it is the part of the relative entropy lost by retaining only the coordinates outside U.

Theorem 12.2 (Caputo–Parisi block factorization, [53, Theorem 2.3, with the non-homogeneous extension of Section 2.2.4])**.** *Consider a finite-spin nearest-neighbor Gibbs specification on $\mathbb{Z}$ with bounded symmetric pair potentials, allowed to depend on the edge, and bounded site potentials. Suppose its marginals satisfy the uniform relative-ratio strong spatial mixing condition*

$$\left\| \frac{d\pi^{\tau'}_{H,\Delta}}{d\pi^{\tau}_{H,\Delta}} - 1 \right\|_\infty \leq K e^{-c\,\mathrm{dist}(x,\Delta)}$$

for every finite $H \subset \mathbb{Z}$, every $\Delta \subset H$ and every $x \notin H$, when τ, τ' differ only at x. There is a constant C, uniform in the finite set H, its exterior, nonnegative weights $(w_U)_{U \subset H}$, and $f \geq 0$, such that

$$\left(\min_{i \in H} \sum_{U \ni i} w_U \right) \mathrm{Ent}_\pi(f) \leq C \sum_{U \subset H} w_U \, \pi[\mathrm{Ent}_U f]. \tag{37}$$

We use this theorem only for a positive nearest-neighbor latent chain. To match its symmetric-potential convention even when that chain is not reversible, we retain the following reduction. For a positive stochastic matrix T, take an independent copy of its stationary chain running backwards in space. If consecutive pair states before relabeling are $(a, b), (c, d)$, their bond weight can be written $T(a, c)T(d, b)$. Interchange the two coordinates at every odd site. The two alternating bond weights are then

$$W_0((a, b), (c, d)) = T(a, d)T(c, b), \qquad W_1((a, b), (c, d)) = T(b, c)T(d, a).$$

Both are symmetric and strictly positive. They are edge-dependent nearest-neighbor potentials of the form allowed in the external theorem.

Here are its mixing and boundary checks. A positive finite matrix has a positive stationary row, and the rows of T^k approach it exponentially, entrywise and in relative ratio. In an interval, changing one endpoint changes a prescribed marginal only through the first prescribed spin on that side and the endpoint normalization. Since $T^k(s, t) = \pi(t)(1+O(e^{-ck}))$ uniformly in s, the stationary weight of the endpoint value cancels between these two factors, and the marginal changes by a factor $1 + O(e^{-ck})$, where k is the distance to that spin. The reversed positive matrix gives the same bound from the other side. For a disconnected finite set only the one or two interval components incident to the changed exterior site are affected; the product of their ratios satisfies the same bound after changing its constant. This is the relative-ratio condition in Theorem 12.2. It holds for the independent pair and survives the alternating on-site relabeling.

The stationary free marginal on a finite interval differs from a fixed-exterior interval law by two strictly positive bounded endpoint factors. Its density ratio therefore lies between constants $l, L > 0$, uniformly in the interval length. Both entropy and averaged conditional entropy compare within these same constants. Indeed, each is an infimum of expectations of the nonnegative integrand

$$f \log(f/g) - f + g,$$

over constants $g \geq 0$, or over nonnegative functions of the retained coordinates, respectively. Limits cover zeros. Thus (37) holds for the free marginal too, with a uniform changed constant. Apply it to the independent pair, with f depending only on the original chain. Independence removes the second chain from both entropy and all averaged conditional entropies. We have obtained (37) for the free marginal of the original positive chain, with no reversibility assumption.

12.2 Erasure contracts one cropped step

First suppose a reference input law μ is a finite sliding-block function of a stationary finite-state chain with positive transition matrix. Allow each latent site to code a physical word of a fixed length b, at one fixed block alignment. A bounded number of neighboring latent coordinates determines each word. Let r be the physical PCA radius. For a nonempty finite output set J, put $I = J + [-r, r]$. We claim that for every input law ν on I,

$$D\big((\nu P)_J \,\|\, (\mu P)_J\big) \leq (1 - \kappa) D(\nu \,\|\, \mu_I), \tag{38}$$

with $\kappa > 0$ independent of J, ν and of translates of the alignment. The assertion is automatic when the right side is infinite.

Take the whole interval H of latent coordinates spanning the coding of I, with its free marginal π_H. Lift ν by the density $f = (d\nu/d\mu_I)$ evaluated on the decoded input. Then $\mathrm{Ent}_{\pi_H}(f) = D(\nu \| \mu_I)$. Strict positivity writes every physical update as a mixture of an input-independent uniform symbol of a fixed weight $u > 0$ and another stochastic kernel. Give the outputs in J independent Bernoulli(u) uniform flags. A latent coordinate influences at most M outputs, where M depends only on the coding and r. Erase that coordinate if all its dependent outputs have uniform flags; coordinates with no such outputs are erased automatically. The random erased set $E \subset H$ is independent of the input, with

$$\mathbb{P}(i \in E) \geq u^M \qquad (i \in H).$$

The erasures of distinct coordinates need not be independent.

Conditional on the flags, the outputs can be sampled knowing only $Z_{H \setminus E}$: every other output uses only retained coordinates, and each flagged output is uniform. Reveal the flags and retained coordinates and use data processing. The entropy chain rule and (37), with $w_U = \mathbb{P}(E = U)$, give

$$\begin{aligned} D\big((\nu P)_J \,\|\, (\mu P)_J\big) &\leq \mathrm{Ent}_{\pi_H}(f) - \sum_U \mathbb{P}(E = U) \pi_H[\mathrm{Ent}_U\, f] \\ &\leq (1 - u^M/C)\, \mathrm{Ent}_{\pi_H}(f). \end{aligned}$$

This proves (38). Taking the whole latent hull causes no diameter cost for disconnected J: the lifted density depends only on I, and its entropy is exactly the entropy on I.

12.3 Hidden chains, classes and time iteration

A primitive hidden transition matrix may have zeros. Choose m with $T^m > 0$ and sample the hidden states at multiples of m. Conditional on this skeleton, the finite bridges are independent. A finite independent mark at each skeleton site can consist of a table containing one random bridge for every ordered pair of possible endpoints, with the appropriate bridge law. Select the entry for the actual pair. Include the emission randomness in these tables, using only each table's finite positive support. The augmented chain of skeleton states and tables has a strictly positive transition matrix. Each physical m-word is a two-site function of it. This is the coding just treated, including deterministic emissions.

For a general stationary finite hidden chain, discard the transient states of zero stationary mass. Take a common multiple p of the periods of its recurrent classes. Conditional on a class and its cyclic subclass at the origin, the p-step skeleton is primitive. Sampling a further common multiple and the bridges gives a finite list of physical reference laws $\mu_1, \dots, \mu_k$, each mixing for a common block shift, and each with the positive latent coding above. Merge identical physical laws. The given stationary row is a positive mixture of these distinct spatially ergodic components.

Each image $\mu_i P$ is ergodic for that block shift: realize the update using iid site randomness and take a factor of a mixing component times that randomness. Invariance of the mixture and uniqueness of spatial ergodic decomposition imply that each image is one of the same finite list. Every component has positive mixture weight, so every one is an image; hence P permutes the list. If $p_{\min} > 0$ is the minimum physical transition probability, conditional independence of the updated sites gives

$$\mu_i(\eta_I) \geq p_{\min}^{|I|} \tag{39}$$

for every finite cylinder, because each μ_i is an image of another component.

Use the worst $\kappa > 0$ from (38) over this finite list and the finitely many alignments. Iterate that inequality through the nested parent sets $J + [-rs, rs]$, for $s = 0, \dots, t$, comparing a deterministic initial row with any initial reference component and its successive images. By (39) the initial relative entropy is at most

$$|J + [-rt, rt]| \log(1/p_{\min}) \leq (2rt+1)|J| \log(1/p_{\min}).$$

Pinsker's inequality yields

$$d_{\mathrm{TV}}\big((\eta P^t)_J, (\mu_i P^t)_J\big) \leq \sqrt{\tfrac{1}{2}(1-\kappa)^t (2rt+1)|J| \log(1/p_{\min})}. \tag{40}$$

At times divisible by the permutation period, the same deterministic-start law approaches every μ_i. The triangle inequality forces all components to coincide on every finite set. Thus the stationary reference is unique, and (40) implies $C|J|e^{-ct}$, uniformly in η, J, t; absorb the polynomial in t into a smaller exponential and use $\sqrt{|J|} \leq |J|$. Integration covers arbitrary initial laws. This is (2). Theorem 7.1 proves the DLR and EWM conclusions of Proposition 12.1.

For Theorem G, any change of Fourier amplitudes that retains the iid stationary hidden row and strict positivity leaves the hidden PCA exponentially ergodic. A full finite decoration freshly sampled from a fixed finite neighborhood in the old hidden row, the site's new hidden symbol and independent site randomness has a finite hidden Markov row: a finite spatial buffer records the overlapping input coordinates. If its old values are ignored and the full update stays strictly positive, the joint stationary row has a finite hidden Markov presentation and the proposition applies to the full PCA. Consequently these amplitude changes and freshly sampled decorations do not create magnetic order. A decoration that reads its own earlier values, or feeds its outputs back into the hidden dynamics, is outside this assertion.

13 Prescribing sampled decay shape on one Ising geometry

The proof of H has three parts. A finite Fourier table turns an enumeration into a countable family of decay modes. A finite Ising decoration realizes the resulting integer energies. A computable fixed point places the prescribed set in the physical inverse-temperature coordinate. The explicit crystal I uses the decoration but a different Fourier table.

13.1 A finite parent for an enumeration

We first describe the parent construction in the coordinate $z \in (0,1)$. All notation in this section is local to this construction. Let O be effectively open in $(0,1)$. Enumerate pairs (K_t, V_t), where K_t is a compact rational interval, V_t is open and rational, and

$$K_t \subset V_t, \qquad \overline{V}_t \subset O,$$

so that the K_t's cover O. A bounded simulation of the enumerator decides what has been printed at a given tick; it need not decide when the next interval will appear.

For each printed pair find finitely many rational (v, s), with $0 < s < v < 1$, such that the caps $2^v z^s$ satisfy

(a) their closed superlevel sets above $1 + z$, on all of $[0,1]$, lie in V_t, and their centers $z_c = s/(1-s)$ lie in V_t;
(b) their strict superlevel sets cover K_t;
(c) $2^v z^s \le (1 + 2^{-t})(1 + z)$ on $[0,1]$.

All tests on a finite proposed list are decidable. A common denominator converts them to rational polynomial sign tests after raising positive quantities to that power.

Here is why the search terminates. In $x = -\log z > 0$, the function $\phi(x) = \log(1 + e^{-x})$ is strictly convex. Its tangent at x_0 has slope $-s_0$, with $s_0 = z_0/(1 + z_0)$, and intercept $c_0 = \phi(x_0) + s_0 x_0$ satisfying

$$s_0 \log 2 < c_0 < \log 2.$$

Raise that tangent slightly and approximate its slope and intercept by $-s$ and $v \log 2$. Strict slack away from its contact keeps its closed superlevel set inside V_t, and the raise can be small enough for (c). As $z \downarrow 0$, the ratio of the cap to $1 + z$ tends uniformly to zero for slopes near $s_0 > 0$; at $z = 1$, strict slack follows from $v < 1$. Compactness supplies finitely many caps covering K_t. Serialize the searches and outputs so that at most one cap is printed per computation tick. There are only finitely many caps with original batch number at most any fixed T.

Lemma 13.1 (Finite parent and its return formula)**.** *From the text of the enumerator one effectively specifies a finite alphabet, an odd one-site spin, and two oriented nearest-neighbor integer energy tables on $\mathbb{Z}^2$: the east table takes values in $\{0, \dots, 4\}$ and the north table values in $\{0, \dots, 11\}$. The alphabet is a product of binary coordinates, so its size is a power of two, and addition of one fixed bit mask is a free involution that preserves both tables and reverses the spin. For $0 < z < 1$, the trajectory law started from the iid uniform hidden row, called the invariant-row Gibbs law, is a Gibbs measure for these tables, and its finite marginals are rational in z, hence analytic. For positive lags its axial correlation is*

$$C_z^{\mathrm{cap}}(n) = \frac{1}{4}\left[\mathfrak{R}_0(z)^n + 2\cos(2\pi n/3)\sum_j w_j(z)\mathfrak{B}_j(n - F_j, z)\right], \tag{41}$$

where $\mathfrak{R}_0 = h(2 + z)$, $\mathfrak{R}_j = h(1 + \gamma_j)$, $\gamma_j = 2^{v_j} z^{s_j}$, and

$$\mathfrak{B}_j(k,z) = \frac{h^k}{\ell_j}\sum_{q=0}^{\ell_j - 1}(1 + \gamma_j e^{2\pi i q/\ell_j})^k = h^k \sum_{\substack{0 \le a \le k \\ \ell_j \mid a}} \binom{k}{a}\gamma_j^a. \tag{42}$$

It is zero for $k < 0$. The entrance weights obey $F_j \geq j + 2$, $0 < w_j \leq (zh/4)^{F_j}$; at a fixed lag only finitely many stages contribute. The scale $h > 0$ can be chosen so small that $3h < 1/4$.

Proof. Use the hidden group $G = \mathbb{F}_2^w$, with characters $\chi_a(y) = (-1)^{a \cdot y}$. In a mask a the low four bits carry a resident label, the next three bits are presence flags for the incoming displacements $-1, 0, 1$, and the remaining $w - 7$ bits are a payload shared by the three directions. The sixteen resident labels include the empty label, a unary tape bit, two interior markers, a right endpoint, and three colors each of the computation origin and of the orbit's left endpoint. A head label consists of a payload, exactly one presence flag and an optional resident. The eight orbit head states use payloads $0, \dots, 7$; the seed, the terminal and the direct character have further reserved payloads with the stationary flag, and the finite control uses the remaining payloads. Residents send only to themselves. An option sends at most one moving head, and may leave a resident; an output to a side neighbor uses only the shared payload and the single flag of that direction, and a center output never carries a side flag. Thus at a receiver distinct senders have distinct flags, and a nonempty character shape cannot cancel to the empty shape. Starting from a seed there is exactly one head; its tape is represented by resident bits. A terminal character has no outgoing option, and reserved majority-triple characters have none either. Unassigned nonzero labels have empty option lists. This means zero Fourier coefficient, rather than an option producing an empty shape.

A finite register program simulates a bounded part of the cap enumeration. Tape stacks in any fixed base, implemented by increment and successful-decrement loops with scratch registers, suffice. Registers are unary prefixes in m disjoint residue classes of sites. Before an orbit is installed all scratch registers are cleared. For a printed cap $v = V/Q$, $s = S/Q$, it deposits markers at

$$0, \quad m(Q - V), \quad m(Q - S), \quad mQ.$$

A round trip has $\ell_j = 2mQ$ advances. The gain and thermal markers toggle independent factors 2 and z, so the complete trip multiplies the advance weight by $2^{2mV} z^{2mS} = \gamma_j^{\ell_j}$. A hold is allowed at each tick. Exiting at the original orbit phase and erasing all markers leaves only the terminal at the origin. The origin resident advances a color modulo three at every tick, including holds; its final multiplier is $2, -1, -1$. This is $2\cos(2\pi n/3)$. The direct mode has weight $h(2 + z)$.

Each option of an input mask has three output masks, for the west, center and east sites of the old row, and a weight

$$\frac{c\, u^e z^s}{4M^e} \quad \text{(resident rule)}, \qquad \frac{c\, u^e z^s}{4M^e(1 + z)} \quad \text{(head rules)},$$

with $c \in \mathbb{Z}$, $1 \leq e \leq 5$, $s \in \{0, 1\}$ ($s = 0$ for the resident rule), $u = ((1 - z)/(1 + z))^2$ and one dyadic divisor $M \geq 2$. A resident pays u/M per tick and the orbit walker $u/[M(1+z)]$; when they occupy one site both factors are charged there. Thus an orbit tick with four markers and one head has weight $h = u^5/[M^5(1 + z)]$, times the gain and thermal factors 2 and z on the marked segments. A work, installation or erasure head has weight $zu^5/[4M^5(1 + z)] = zh/4$ per tick, and every other occupied site has weight $u/M \leq 1$. The only moving $e = 1$ rules are orbit advances, whose report characters lie in the small report block described below. Every remaining moving rule has $e \geq 2$ and uses at most the two blocks of one side report, and stationary rules use no report. Hence $1 \leq e \leq 5$ and $2e > |J|$, where J is the set of nontrivial horizontal report blocks inverted in that rule. At most one candidate orbit is installed per tick. Extra residents cost at most one, proving the entrance estimate.

Here is a fixed pair-energy realization of these rational Fourier coefficients. A site carries two side reports l, r. Its left report is compared with π_- of its west neighbor's hidden word, and its right report with π_+ of its east neighbor's hidden word; here $\pi_\pm(y)$ is the shared

payload of y with the flag of that direction as its lowest bit. Split each report into two blocks of even width: a small block of four bits, holding the flag and the three lowest payload bits that carry every orbit head, and a large block holding the remaining payload bits, whose number is chosen even. For a block of b words put

$$Q(x) = x_0x_1 + x_2x_3 + \cdots \pmod 2, \qquad L_b(z) = \tfrac{1}{2}[b + \sqrt{b} + (b - \sqrt{b})z].$$

The report channel has weight $z^{Q(r\oplus x)}/L_b(z)$. Its nontrivial character multiplier is

$$(-1)^{Q(a)}\frac{1 - z}{\sqrt{b} + 1 + (\sqrt{b} - 1)z}.$$

This follows by multiplying the two-bit Walsh identity $\sum_{x,y}(-1)^{xy+ix+jy} = 2(-1)^{ij}$. Multiply each moving Fourier coefficient by the inverse of its active report multipliers before integrating the reports. For a south context (l, y, r) and a new hidden word y', clearing the common denominator $4M^5(1 + z)^{11}$ in $|G|$ times the resulting hidden transition gives a polynomial of degree at most eleven, with background polynomial $4M^5(1 + z)^{11}$. The contribution of an option has a factor $(1 - z)^{2e-|J|}$, so it vanishes at $z = 1$. Increase M until, at each degree $0 \le \iota \le 11$, the sum of absolute perturbation coefficients is smaller than the background coefficient $4M^5\binom{11}{\iota}$. This terminates because the table is finite and every $e \ge 1$. All coefficient lists are then strictly positive integers with the same sum $I = 2^{13}M^5$ at $z = 1$, a power of two. For a south site (y, l, r, i) and a north site (y', l', r', i'), list the monomials of the polynomial for $(y' \mid l, y, r)$ with multiplicity in increasing order of exponent; the north energy is the exponent at position $i' \in \{0, \dots, I-1\}$, and lies in $\{0, \dots, 11\}$. The east energy from (y, l, r, i) to its east neighbor (y', l', r', i') is the sum of Q over the two blocks of $l' \oplus \pi_-(y)$ and the two blocks of $r \oplus \pi_+(y')$, and lies in $\{0, \dots, 4\}$. The alphabet $G \times \{0, 1\}^{2(w-6)} \times \{0, \dots, I - 1\}$ has a power-of-two size.

The normalized factors have configuration-independent sums. Starting with an iid uniform hidden row and using these factors gives the pair Gibbs specification by finite-slab weight ratios. Every fixed cylinder uses a finite backward cone, hence is a rational function of z with positive denominator. Nonextinction proves iid hidden-row invariance by character orthogonality. The observed spin is the majority of the seed, terminal and direct characters, with Fourier coefficients $1/2, 1/2, 1/2, -1/2$. Only the direct return and seed-to-terminal returns can pair. This gives the factor $1/4$ and (41). Finally, choosing which ticks advance proves (42). Flipping the three presence flags of the hidden word and the flag bit of each report adds one fixed bit mask to every symbol. Every option preserves the parity of presence flags, so both pair tables and the index are invariant, while the seed, terminal and direct characters each change sign, and so does the observed spin. This free involution, with one flipped coordinate as pivot, is the one used by the rail representation in Section 14.5. □

13.2 Which mode dominates?

Each fixed orbit has the simple leading rate $\mathfrak{R}_j = h(1 + \gamma_j)$: all other roots of unity in (42) give strictly smaller modulus. Put $\mathfrak{H}(n) = \sum_j w_j\mathfrak{B}_j(n - F_j)$. At any comparison rate $\rho \ge \sup_j \mathfrak{R}_j$ with $\rho \ge h$,

$$0 \le \mathfrak{H}(n) \le \rho^n \sum_j 4^{-F_j} \le \rho^n/12.$$

Off O, every $\mathfrak{R}_j < \mathfrak{R}_0$. Dominated convergence gives $\mathfrak{H}(n)/\mathfrak{R}_0^n \to 0$. The same bound places $C_z^{\mathrm{cap}}(n)$ between $(11/12)\mathfrak{R}_0^n/4$ and $(14/12)\mathfrak{R}_0^n/4$; the choice of scale makes successive bounds strictly separated, including $C_z^{\mathrm{cap}}(0) = 1$. Thus the correlations are positive and strictly decreasing at every lag.

On O, a cap strictly exceeds $1 + z$. The excesses tend effectively to zero with batch number. The maximum rate above $\mathfrak{R}_0$ is therefore attained by a finite nonempty collection, with a gap to every remaining rate. Combining its positive residues gives

$$C_z^{\mathrm{cap}}(n) = A\rho^n \cos(2\pi n/3) + O(\sigma^n), \qquad A > 0, \quad 0 < \sigma < \rho < 1.$$

Subtracting the cosine leaves a generating function of strictly larger radius, proving precisely two simple nearest poles. Ties between maximizing stages add positive residues.

For context, the parent construction also distinguishes interior and boundary points of the complement. On a compact neighborhood of a complementary point, integrate the positive curvature of ϕ along the segment toward each cap center. Centers a fixed distance away incur a fixed positive deficit, even when outside that neighborhood. For centers in a common compact interval the log excess is at most $\log(1 + 2^{-t}) - \kappa\varrho^2/2$, where ϱ is the distance to the center and $\kappa = \min \phi'' > 0$. At an interior complementary point this gives a strict gap below $\mathfrak{R}_0$ for all late batches; the finite earlier list also has a gap. At a boundary point the supremum is $\mathfrak{R}_0$, unattained. The positive stage series has radius $1/\mathfrak{R}_0$, by a fixed-stage lower bound and the summable upper bound. Pringsheim's theorem [58, Theorem IV.6] makes its positive boundary point singular. Abel averaging of $\mathfrak{H}(n)/\mathfrak{R}_0^n \to 0$ rules out a pole. All orbit eigenvalues lie in the disk with center h and radius $\mathfrak{R}_0 - h$; it meets the outer circle only on the positive axis. Normal convergence of the entrance-weighted resolvents continues the series across every other arc. Adding the baseline and rotating the stage series gives one real simple pole and two non-poles at the boundary. Thus at boundary points the oscillatory term is only subexponentially smaller than the positive pole term. The converse below therefore uses the all-lag sign promise, whose negative lags finite certificates can detect, rather than a condition on eventual oscillation alone.

13.3 Finite Ising realization and the physical coordinate

The finite decoration proved in Section 14.5 realizes an even Boolean output table as a three-site factor $1 + \mathcal{T}$ sign, up to a positive analytic constant. Its transmission $\mathcal{T}$ increases computably and strictly from zero to one. For the integer energy $E \in \{0, \dots, 15\}$, use $\delta = 1/32$ and the identity

$$(1 - \delta\mathcal{T})^E = 1 + \mathcal{T} \sum_{k=1}^{E} (-\delta)^k \binom{E}{k} \mathcal{T}^{k-1}. \tag{43}$$

The absolute coefficient sum is at most $(33/32)^{15} - 1 < 1$. Fair finite selector bits choose the indicated signed monomials with their dyadic probabilities; divide the remaining mass evenly between the two signs. Independent calibration chains supply the powers of $\mathcal{T}$. The Boolean circuit uses formal calibration inputs. Install the chains only after the finite circuit is fixed, each duplicating that circuit's own transmission. This proves the identity also for $E = 0$ and avoids any equation involving an unknown chain length.

The pair table is finite and even under its free involution. Disjunctive normal form, finite addition and multiplexing therefore produce a Boolean circuit for it; exhaustive finite verification is an alternative construction. Fanout-two COPY expansion and the one-contact modules give a finite graph of maximum degree three. This first compilation uses only the program text. The unbounded cap searches occur in its dual computation, at increasing lags, rather than during construction of the finite graph.

Apply the equal-length two-layer routing and 32-site switch decoration of Section 14.6. This uses the same geometric graph for every input, changing only a finite periodic sign pattern on a fixed matching. For candidate enumerator text p, its physical-to-parent coordinate is

$$z_p(\beta) = 1 - \delta\mathcal{T}_p(b_{\mathrm{eff}}(\beta)).$$

It is a computable strictly decreasing homeomorphism from $(0,\infty)$ to $(31/32,1)$, with computable inverse. These assertions follow from the strictly increasing finite transmission factors and the inverse found by monotone search.

The reached parent inverse temperature $B_p = -\log z_p$ is at most $\log(32/31)$. A neighbor changes an energy table of range eleven vertically or four horizontally. The likelihood-ratio bound gives the direct physical Dobrushin row sum

$$2\tanh(11B_p/2) + 2\tanh(2B_p) \le 15\log(32/31) < 0.48. \tag{44}$$

It proves parent uniqueness and arbitrary-boundary exponential mixing on the entire reachable interval. The collar argument of Section 14.7 transfers these properties to the microscopic graph. The configuration-independent finite normalizers give analytic pressure; finite marginals are analytic combinations of the parent's rational marginals. No global hidden-comparison assertion away from this reached interval is required.

Extend the requested $W \subset [a,b]$ to an effectively open $U \subset (0,\infty)$ by the preimage under clipping to $[a,b]$. From p, enumerate $z_p(U)$ by rational inner approximations of images of open intervals. This is a total computable transformation F of enumerator indices. A fixed index can be chosen effectively. To see this directly, let $s(i,x)$ fix the first argument of a two-input program. Choose e so that its computation on (x,t) simulates the t-th tick of $F(s(x,x))$. Then $p = s(e,e)$ enumerates exactly what $F(p)$ enumerates. The substitutions are finite string operations. For this p,

$$z_p(\beta) \in O \quad\Longleftrightarrow\quad z_p(\beta) \in z_p(U) \quad\Longleftrightarrow\quad \beta \in W \quad (\beta \in [a,b]).$$

Four-parent-step sampling gives $C^{\mathrm{cap}}_{z_p(\beta)}(4n)$; because $4 \equiv 1 \pmod 3$, it preserves the cosine's residues and the two simple nearest poles. This proves the realization in H, including the fixed geometry and the one-time computation of signs, site and vector. The square-lattice variant omits the switch decoration: flattening the diluted two-layer routing by $(x,y,z) \mapsto (x-y, x+y+z)$ gives nearest and diagonal exchanges in $\{-1,0,+1\}$. Its coordinate map is $1 - \delta\mathcal{T}_p(\operatorname{atanh}((\tanh\beta)^{L_p}))$, with L_p its wire length, and the fixed-point step is repeated with this map.

13.4 The converse and its sign promise

For any finite-period rational interaction, rational local spin X, site o and vector d, use $f_n = X_o X_{o+nd}$. The finite-box certificate

$$\max_{\tau} \gamma^{\tau}_{\Lambda,\beta}(f_n) < 0$$

is sound by the DLR equation. Under uniqueness, the all-exterior range converges to the unique expectation as Λ grows: otherwise a weak limit of bad boundary laws is a different DLR law. Every strictly negative expectation consequently has a finite certificate. Finite energies make each finite-box expectation a computable analytic function; certified inequalities on rational compact intervals enumerate its strict-negative neighborhoods. Dovetail over boxes, positive lags and precision, as in Section 11.5. This enumerates exactly the set with some negative sampled lag, uniformly in the interaction, observable, sampling data and window.

Under the all-lag sign promise, the latter set equals the oscillatory set $\mathcal{O}$. Uniqueness supplies completeness of the certificates; the promise supplies that identification. In G the object of the converse is already some negative response, so its uniqueness converse needs no additional sign promise. This comparison concerns the sampled-correlation sets; it makes no claim about the crystal's response at every displacement.

14 An explicit crystal and its halting-degree disorder line

We prove I independently of the cap construction. The parent now has 2^{105} symbols and only the two bond-energy values zero and one. A universal finite circuit identity then identifies the microscopic Gibbs marginal exactly with this parent. Notation introduced in this section is local to it; for example η below is a report multiplier, not an initial row.

14.1 The binary-energy parent

The parent table printer and its two register modules define the complete finite tables. We describe their coordinates and the computation relevant to the correlation. A symbol is (y, l, r, i), with $y \in \mathbb{F}_2^{20}$, two 14-bit reports l, r, and $0 \le i < 2^{57}$. Its ordinal is

$$(((y\,2^{14} + l)2^{14} + r)2^{57} + i).$$

For a hidden mask a, $\chi_a(y) = (-1)^{\mathrm{popcount}(a\&y)}$. The observed parent spin is the majority of the three characters with masks $1568, 1056, 1184$. Its four-term Fourier expansion has coefficients $1/2, 1/2, 1/2, -1/2$.

The hidden masks use four low resident bits, three incoming-head flags and thirteen payload bits. The side projection π_- or π_+ retains that direction's flag and the shared payload. With Q the adjacent-pair quadratic on a 28-bit word, the east energy is

$$E_{\mathrm{e}}(s, s') = Q\big(l' \oplus \pi_-(y),\ r \oplus \pi_+(y')\big) \in \{0, 1\}.$$

The printed scattering list consists of tuples (a, l_0, c_0, r_0, c). Put $D_{\mathrm{p}} = 2^{56}$ and

$$N_0(l, y, r, y') = D_{\mathrm{p}} + \sum_{(a,l_0,c_0,r_0,c)} c\,\chi_a(y')\chi_{l_0}(l)\chi_{c_0}(y)\chi_{r_0}(r). \tag{45}$$

The north energy from s to s' is 0 if $i' < N_0$, and 1 otherwise. There are 164 852 terms and

$$\sum |c| = 33\,293\,550\,733\,133\,931 < 2^{56}. \tag{46}$$

Both north multiplicities are therefore strictly positive for every context, including contexts unreachable from a seed. The finite list, rather than an approximation to a real number, defines the Hamiltonian.

Write $\mathsf{t} \in [1/4, 1/2]$ for parent inverse temperature, and set

$$z = e^{-\mathsf{t}}, \qquad \nu = \frac{1 - z}{1 + z}, \qquad \eta = \frac{\nu}{2^{14} + \nu}, \qquad L_{\mathrm{e}}(z) = \tfrac{1}{2}[2^{28} + 2^{14} + (2^{28} - 2^{14})z].$$

The quadratic Walsh identity used in Lemma 13.1 gives this east normalizer. Its nontrivial report multiplier is $(-1)^{Q(a)}\eta$. The printed moving coefficients contain that sign, so every single-head move acquires the positive factor η. The north update has normalizer $2^{20}D_{\mathrm{p}}(1+z)$. Head flags from distinct senders cannot cancel, and residents send only to themselves. Consequently the iid uniform hidden row is invariant. Slab ratios give the pair Gibbs specification and the pressure per parent site

$$p_{\mathrm{p}}(\mathsf{t}) = \log\big(2^{20}D_{\mathrm{p}}(1 + z)\big) + \log L_{\mathrm{e}}(z). \tag{47}$$

Every cylinder probability is rational in z, by its finite backward cone.

For completeness, uniqueness is a direct physical Dobrushin comparison; see [57, Section 2.3]. Changing one neighbor changes each candidate central energy by a number in

$\{-1,0,1\}$. The likelihood-ratio spread is at most $e^{2\mathfrak{t}}$, so its total-variation influence is at most $\tanh(\mathfrak{t}/2)$. Thus the Dobrushin row sum is at most

$$q_{\mathrm{p}} = 4\tanh(1/4) < 0.980 < 1.$$

For a probe V at distance $s \geq 1$ from differing exterior sites, path summation gives the uniform total-variation bound

$$|V|\frac{q_{\mathrm{p}}^{s-1}}{1-q_{\mathrm{p}}}, \qquad \frac{1}{1-q_{\mathrm{p}}} < 50. \tag{48}$$

This is valid in every finite shape and for all boundary symbols. It identifies the constructed trajectory law with the unique DLR law and proves uniform EWM. Formula (47) and the cylinder formulas give analyticity directly.

14.2 The halting real and the finite stage program

An integer encodes a finite instruction list by empty list 0 and cons cell $2^k(2v+1)$, with head k and tail v. Write $k = 4t+c$. For $c = 0, 1$, increment counter c and advance; for $c = 2, 3$, test counter $c-2$, advance on zero, and otherwise decrement and jump to instruction t. Falling off halts. Let K_{h} be the zero-start halting set. This successful-decrement convention is the two-counter model of [56, Definition 2], whose halting problem from an arbitrary configuration is undecidable [56, Theorem 6, after Minsky]. A finite initial configuration is incorporated by an increment prefix and readdressing the original instructions; an increment followed by a successful decrement can jump to any chosen original start while restoring the counter. The effective universal simulation is a many-one reduction. Hence K_{h} is Σ_1^0-complete, not just undecidable.

Let $K_{\mathrm{h},j}$ contain codes $e < j$ that halt within j instructions, and put

$$\vartheta_j = \frac{1}{2} + \sum_{e \in K_{\mathrm{h},j}} 4^{-e-2} = \frac{p_j}{q_j}, \quad q_j = 4^{j+1}, \qquad \vartheta = \frac{1}{2} + \sum_{e \in K_{\mathrm{h}}} 4^{-e-2}. \tag{49}$$

The numerator need not be reduced. We have $\vartheta_j \uparrow \vartheta$, $\vartheta_j < \vartheta$ at every finite stage, and $1/2 \leq \vartheta \leq 7/12$. The gap between a missing digit's possible tail and a present digit is at least $(2/3)4^{-e-2}$. An approximation oracle for ϑ therefore recovers all halting bits successively; a halting oracle sums a finite prefix and bounds the geometric tail. Thus $\vartheta \equiv_T K_{\mathrm{h}}$.

The finite register program computes these bounded approximants using seventeen unary tracks. Register i is a prefix of resident bits at positions $1 + i + 17k$. Increment sets the first missing bit; successful decrement clears the last present bit; every instruction returns to the marked origin. Binary list decoding uses repeated division by two and a bounded-fuel interpreter. The expanded program has 235 register instructions. It clears all scratch registers, subtracts one from each output length, and installs markers on the first block missing the corresponding prefix. The four coordinates are

$$0, \qquad mq_j, \qquad m(2q_j - p_j), \qquad 2mq_j, \qquad m = 17.$$

The one-head transition rules, resident updates, incoming directions and unused labels are the finite list generated by the programs. An independently coded register assembly and scattering interpreter generates exactly the same coefficient multiset, as certified in Section 14.4.

14.3 Return series and the two sides

Residents cost $\mathfrak{r}\nu$, with $\mathfrak{r} = 2^{-6}$. Work heads cost $e_0\nu$, with $e_0 = 2^{-56}$, and moves additionally cost η. At the left endpoint an orbit may hold, start either trip, or begin complete erasure. Its hold weight is

$$\mathfrak{b} = 2^{-7}\nu(\mathfrak{r}\nu)^3.$$

The long trip goes to the right endpoint and back, moving every tick; head weight is $a_0\nu$, $a_0 = 2^{-14}$, and each traversal of the rightmost mp_j edges gains a factor two. The short trip goes to mq_j and back with a stay before every move, head weight $a_1\nu$, $a_1 = 2^{-22}$, and a minus sign on its first stay. Both trips take $\ell_j = 4mq_j$ ticks. A head on a marker pays an extra divisor 4096, omitting that marker's separate resident payment. The net collision factor is $\varkappa = (64\nu)^{-1}$. Their positive full weights are

$$\begin{aligned} P_j^+ &= [a_0\nu(\mathfrak{r}\nu)^4\eta]^{\ell_j} 2^{2mp_j}\varkappa^6, \\ P_j^- &= [a_1\nu(\mathfrak{r}\nu)^4]^{\ell_j}\eta^{\ell_j/2}\varkappa^4, \\ \frac{P_j^+}{P_j^-} &= [2^{16}\eta 2^{\vartheta_j}]^{2mq_j}\varkappa^2. \end{aligned} \tag{50}$$

The long trip has six collision ticks; the short has four and half as many moves. Their elapsed-time color agrees. Every return is a unique sequence of endpoint holds and complete trips. Its series, with the final color omitted, is

$$\sum_{k\ge 0}\mathfrak{B}_j(k)t^k = \frac{1}{1-\mathfrak{b}t-(P_j^+-P_j^-)t^{\ell_j}}. \tag{51}$$

The color $2,-1,-1$ is retained at the origin throughout computation, trips and erasure. Only complete erasure gives a singleton there that pairs with the observed spin.

There is also a positive reference branch. Put $d_0 = 2^{-48}\nu$. It either repeats this direct step or installs four resident bits at $0,1,2,3$, returns, holds with weight $\mathfrak{b}$, and erases them. Its fixed part has thirteen ticks, twelve moves, and head coefficient 2^{-51}. Denote its weight by ω. For $n \ge 1$ its total is

$$\mathfrak{K}(n) = d_0^n + \omega\sum_{a+c=n-13} d_0^a\mathfrak{b}^c,$$

with empty sum zero. The exact majority correlation is

$$C_{\mathrm{t}}^{\mathrm{bin}}(n) = \frac{1}{4}\left[\mathfrak{K}(n) + 2\cos(2\pi n/3)\sum_j w_j\mathfrak{B}_j(n-F_j)\right], \qquad n \ge 1. \tag{52}$$

Set $\mathfrak{B}_j(k) = 0$ for $k < 0$. Orthogonality against the iid old row selects only the direct return and the seed-to-terminal singleton. The other majority character has no rule and the seed is never reproduced. This proves the formula without interpreting signed dual weights as physical probabilities.

Outside its trips and holds, each tick of a stage path has one working head of weight at most $e_0\nu$, every other occupied site has weight at most $\mathfrak{r}\nu < 1$, and moves only add factors $\eta < 1$. Removing the final color factor therefore gives $0 < w_j \le (e_0\nu)^{F_j}$. The stage choice takes at least j ticks, and installation and erasure each travel from the origin to the right endpoint $2mq_j \ge 136$, so $F_j \ge j + 272$. The following window bounds will suffice:

$$\begin{gathered} \ell_j \ge 272, \quad \mathfrak{b} > 2^{-41}, \quad \eta > 2^{-19}, \quad 2^{-1400} < \omega \le d_0^{13}, \quad d_0 < \mathfrak{b}/2, \\ \frac{|P_j^+ - P_j^-|}{\mathfrak{b}^{\ell_j}} \le 2\,16^{-\ell_j}, \qquad \frac{w_j}{\mathfrak{b}^{F_j}} \le 2^{-15F_j}, \qquad \mathfrak{K}(n+1) \le (\mathfrak{b}+2d_0)\mathfrak{K}(n). \end{gathered} \tag{53}$$

They follow from $1/16 < \nu < 1/4$, the thirteen fixed ticks, and the fact that the largest trip tick divided by $\mathfrak{b}$ is at most $2(a_0/2^{-7})\mathfrak{r}\nu < 1/16$. For ω, each of the thirteen fixed ticks has head coefficient 2^{-51}, at most one move and at most three separate residents, so its weight lies between $2^{-51}\nu\,\eta(\mathfrak{r}\nu)^3 > 2^{-104}$ and $2^{-51}\nu < d_0$.

If $P_j^+ < P_j^-$, put $\mathfrak{d}_j = |P_j^+ - P_j^-|/\mathfrak{b}^{\ell_j}$ and choose $\rho_j \in ((\ell_j - 1)/\ell_j, 1)$ solving $\rho_j^{\ell_j - 1}(1-\rho_j) = \mathfrak{d}_j$. In $x = \mathfrak{b}t$ the factorization

$$1 - x + \mathfrak{d}_j x^{\ell_j} = (1 - \rho_j x)\left[1 - (1-\rho_j)(x + \rho_j x^2 + \cdots + \rho_j^{\ell_j - 2} x^{\ell_j - 1})\right]$$

has reciprocals with nonnegative coefficients. The second bracket has no zero in $|x| \le 1/\rho_j$, since $(\ell_j - 1)(1-\rho_j)/\rho_j < 1$. Thus $0 < \mathfrak{B}_j(k) \le \mathfrak{b}^k$, with a simple positive leading rate $\lambda_j = \mathfrak{b}\rho_j$. If $P_j^+ > P_j^-$, the positive recurrence gives a unique simple Perron rate $\lambda_j > \mathfrak{b}$ with

$$\lambda_j^{\ell_j - 1}(\lambda_j - \mathfrak{b}) = P_j^+ - P_j^-.$$

The one-tick hold makes equality in the triangle bound possible only in the positive direction. Induction gives $\mathfrak{B}_j(k) \le \lambda_j^k$ and $0 < \lambda_j/\mathfrak{b} - 1 \le 2\,16^{-\ell_j}$.

Define $\mathfrak{t}_0$ by

$$2^{16}\eta(\mathfrak{t}_0)2^{\vartheta} = 1. \tag{54}$$

The left side increases strictly with $\mathfrak{t}$. At $1/4$ it is below one using $\nu < 1/8$, $\vartheta \le 7/12$; at $1/2$ it is above one using $\nu > 1/5$, $\vartheta \ge 1/2$. Equivalently, $\eta = 2^{-16-\vartheta}$, $\nu = 2^{14}\eta/(1-\eta)$, and $\mathfrak{t}_0 = 2\,\mathrm{atanh}\,\nu$. This is a computable decreasing homeomorphism of ϑ, with computable inverse, so $\mathfrak{t}_0$ is right-c.e. of halting degree.

For $\mathfrak{t} \le \mathfrak{t}_0$, all trips have $P_j^+ < P_j^-$. The estimates in (53) bound the stage sum by $\mathfrak{K}(n)/12$. Together with the last inequality there, they give strict positivity and strict decrease at every lag, including $C(0) = 1$. Each $\lambda_j < \mathfrak{b}$, so dominated convergence gives a positive constant times $\mathfrak{b}^n$. For $\mathfrak{t} > \mathfrak{t}_0$, all sufficiently late trip ratios exceed one. Since their possible excess over $\mathfrak{b}$ tends effectively to zero, the largest $\lambda_j > \mathfrak{b}$ is attained by finitely many stages, with a gap to the rest. Their positive residues add. Equation (52) then has the asserted cosine asymptotic and exactly two simple nearest poles, pointwise in $\mathfrak{t}$.

On the entire monotone side, $\rho_j \to 1$. Let $\mathfrak{H}(n) = \sum_j w_j \mathfrak{B}_j(n - F_j)$. It is nonnegative, $\mathfrak{H}(n)/\mathfrak{b}^n \to 0$, and a fixed-stage lower bound followed by $j \to \infty$ gives radius $1/\mathfrak{b}$ for its series. Pringsheim makes the positive point singular; Abel averaging rules out a pole there. At any other point of $|\mathfrak{b}t| = 1$, the denominators in (51) stay bounded away from zero for all late stages on a common neighborhood, since their perturbations are bounded by $2(|\mathfrak{b}t|/16)^{\ell_j}$. The finitely many exceptional stages have all poles strictly outside the circle. Summable entrance weights give normal convergence. Rotating this series by the two cube roots produces precisely two non-pole singularities; the reference branch supplies the real simple pole. This argument includes $\mathfrak{t} = \mathfrak{t}_0$.

14.4 The universal circuit identity

The parent involution flips word positions $57, 71, 89, 90, 91$. Represent a parent symbol by an orientation spin and 104 rail pairs recording an orbit representative with pivot 57 set to zero. Decoding is

$$d = \mathrm{rep} \oplus [g \oplus \mathrm{maj}(\mathrm{rep})]F, \qquad F = \{57, 71, 89, 90, 91\}.$$

Here maj(rep) is the majority of the three parity bits with masks $1568, 1056, 1184$ in its hidden 20-bit block, and g is the desired negative-spin bit. There are three 105-bit words, seven fair slot bits and one calibration bit, for 323 Boolean inputs. A private orientation per factor reverses the three parent orientations simultaneously, leaving both energies invariant.

The circuit computes the two outgoing energy bits and then the following output bit, where $s \in \{0, \dots, 127\}$ and $c \in \{0, 1\}$:

$$T(E, s, c) = \begin{cases} [s \geq 64], & E = 0, \\ [s \geq 40], & E = 1, \\ c, & E = 2,\ s < 18, \\ 0, & E = 2,\ 18 \leq s < 25, \\ 1, & E = 2,\ s \geq 25. \end{cases} \tag{55}$$

Lemma 14.1 (Circuit identity)**.** *The ordered list of 85 041 gates specified by the accompanying printer computes* (55) *for the decoded binary-parent energies on every one of its* 2^{323} *input assignments.*

Proof. The finite proof uses exact tuple equality, Boolean polynomial coefficients and integer identities. The accompanying certifier implements the following checks; its full object digests are printed in the reproducibility appendix.

First, hash-cons both the returned list and an explanatory raw gate list in a common symbolic term dictionary. Complemented literals, constant folding and commutative interning are justified by their finite truth tables. Their output IDs agree. This proves equality of functions without trusting interning, pruning or fanout expansion.

Second, compare all 315 decoded wires in algebraic normal form with the displayed decode formula. Regard the decoded words as free variables and check that the east output is the quadratic form of the two report differences.

Third, for each two- or three-input compressor exhaust its truth table and verify $\text{low} + 2\,\text{high} = \sum \text{inputs}$. Traverse each width-$w$ block backwards from its output weights $1, 2, \dots, 2^{w-1}$, checking that every substituted carry has twice the low weight modulo 2^w. Discarded carries have weight zero modulo 2^w. The remaining tokens equal the advertised input weights. This proves all 68 blocks, containing 7 574 full adders and 141 half adders. A popcount of m bits uses width m.bit_length(), so its modular identity is an exact count.

Fourth, interpret every population input as the negative-bit indicator of a signed Walsh character. Expand the proposed groups into full decoded coordinates. Their collected integer character dictionary must equal an independently coded reference, checked equal as a multiset to the printed table of (45). The reference has its own 235-instruction assembly and scattering interpreter and reads no circuit. Suggested groups are used only after their identities are verified. Writing each signed population as its size minus twice its bit count, compare the resulting quadratic integer polynomial with the actual incoming tokens of the 58-bit accumulator. The equality is modulo 2^{58}; (46) and $0 \leq i' < 2^{57}$ put $i' - N_0$ strictly between -2^{57} and 2^{57}. Its complemented sign bit is exactly $[i' \geq N_0]$, including equality.

Finally, cut the two energy wires and exhaust all $4 \cdot 128 \cdot 2 = 1024$ remaining assignments. The actual output agrees with (55). The verified regions cover the output cone. Composing their identities proves the lemma for every input. □

14.5 Exact thermal realization of the circuit

We give the finite-temperature identity, including the normalization that makes it valid for a Gibbs specification. Begin with a four-terminal parity core. For $i = 0, \dots, 3$, join terminal t_i to a hub, then the hub to two children $u_{i,0}, u_{i,1}$. Introduce pairs $c_{j,0}, c_{j,1}$, $j = 0, \dots, 3$, with a positive internal bond, and join $u_{i,\lfloor j/2 \rfloor}$ to $c_{j,\lfloor i/2 \rfloor}$ with sign $(-1)^{\text{popcount}(i \& j)}$. All other bonds are positive. It has 20 private spins, 32 bonds, and one contact at each terminal. Every internal degree is three.

Its trace depends only on the terminal product: even terminal flips are symmetries of the four-column Walsh matrix, combined with relabeling internal vertices. At $y = e^{-2b}$, enumerate the 256 center assignments and then the eight choices of the three spins at each port. This gives integer terminal polynomials $P(y), N(y)$, of degree at most 32, for positive and negative product. Their first nonzero degrees are 5 and 4. For $\widetilde{P}(x) = (1+x)^{32} P(1/(1+x))$ and $\widetilde{N}$ similarly, every coefficient of $\widetilde{N}'\widetilde{P} - \widetilde{N}\widetilde{P}'$ is nonnegative, with nonzero coefficients at degrees 13 and 54. These finite integer facts are verified by `cubic_parity.py`. Thus $J(b) = \frac{1}{2}\log(N/P)$ increases strictly from zero to infinity. Reversing one terminal contact gives the orientation with trace proportional to $\exp(J\prod t_i)$, with the same positive analytic normalizer. Put

$$u = \tanh J(b), \qquad c_b = \tanh(2J(b)), \qquad s_b = \tanh b.$$

Concatenating parity cores through a private spin multiplies their transmissions. Repeating a private pad twice gives a two-terminal factor of transmission u. COPY and NOT predicates on rail products use u^2, XOR uses u^3. A constant ZERO or ONE gate uses its output rail pair and herald pair as a four-terminal predicate, so one parity core gives $\alpha_i = u$. A selector realizes AND and OR with transmission $u^3 s_b^5 c_b$. Here is its conditional calculation. Two control rail buffers give a product of mean s_b^2 times the true sign; a separate two-link chain gives a product with mean s_b^2. With these controls U, V, two arms have exponents $J[XA(V+U)+YB(V-U)]$. Exactly one arm is active. Tracing the other arm and a two-link join gives conditional output mean

$$s_b^2 c_b\,[(V+U)A + (V-U)B]/2$$

with input-independent normalizer. Averaging controls selects the appropriate variable or constant branch with gain $s_b^4 c_b$. Additional parity channels multiply by the output and herald signs; closing one positive bond gives $u^3 s_b^5 c_b$. The same transmission is realized by a two-terminal calibration chain. Every raw argument has one contact. Fanout-two expansion gives each rail one producing and at most two consuming contacts.

For gate i, let $G_i \in \{-1, 1\}$ indicate whether its actual output is correct, and use a separate herald product h_i and loop spins t_i. Realize, up to constant analytic factors,

$$(1 + \alpha_i G_i h_i)(1 + \alpha_i h_i)(1 + u h_i t_i t_{i+1}).$$

Summing h_i, and omitting its factor two, gives

$$1 + \alpha_i^2 G_i + \alpha_i u (1 + G_i) t_i t_{i+1}. \tag{56}$$

An incorrect gate has zero transmitted loop coefficient exactly. Close the loop by $1 + u t_{N_g} t_0 (-1)^{\text{output}}$, where N_g is the number of gates. The empty loop product, summed over gate outputs, is 2^{N_g}, independent of the external inputs: gate errors parameterize all output assignments in an acyclic circuit. The full loop product selects its unique fully correct assignment. Their ratio is

$$\mathcal{T}(b) = u \prod_{i=0}^{N_g - 1} (\alpha_i u) = u^{313944} c_b^{23289} s_b^{116445}.$$

The graph has $N_g = 85\,041$ gates and 23 289 nonlinear ones. A separate chain with consecutive transmissions $\alpha_0, u, \ldots, \alpha_{N_g-1}, u, u$ has endpoint product of mean $\mathcal{T}$. Its endpoints are used once as the calibration input. Averaging the seven fair selector bits in (55) gives 1, $1 - 3\mathcal{T}/8$, and $1 - 3\mathcal{T}/4 + 9\mathcal{T}^2/64$, for $E = 0, 1, 2$. Hence the whole three-site factor is

$$A(b)(1 - 3\mathcal{T}(b)/8)^E, \qquad B(b) = -\log(1 - 3\mathcal{T}(b)/8), \tag{57}$$

with $A(b) > 0$ analytic and independent of all visible symbols. This is an exact equality for every boundary input. It is the decoration–iteration principle [32], here with a circuit selecting the finite energy table.

The visible word has one orientation spin and 104 rail pairs, each word having 2^{104} equally weighted lifts. A private orientation per factor has three contacts; a visible spin occurs once in each of its three incident factors. Herald rails have three contacts and the loop and calibration chains have at most three. The finite graph has 23 573 144 spins and 34 543 714 edges per cell, maximum degree three, and no parallel edges. Its pointwise edge printer gives the full finite specification. The map $B(b)$ is computable and strictly increasing from zero to $\log(8/5)$, with computable inverse.

14.6 The fixed geometric graph

Route each edge of the finite trivalent graph on two square layers with disjoint equal-length wire interiors. We record enough of the construction to identify the geometry independently of the circuit. If upper slot counts are $V_{\mathrm{s}}, E_{\mathrm{s}}$, an edge of rank i with owner $(a_{\mathrm{own}}, b_{\mathrm{own}})$ has color

$$j = 16i + (a_{\mathrm{own}} \bmod 4) + 4(b_{\mathrm{own}} \bmod 4).$$

Reserve disjoint port columns for vertices, disjoint horizontal three-lane bands for colors in layer zero, and disjoint trunk columns for colors in layer one. Define

$$J_{\mathrm{col}} = 16E_{\mathrm{s}}, \quad W_{\mathrm{lay}} = 8\lceil J_{\mathrm{col}}/8\rceil + 32, \quad P_y = 8(J_{\mathrm{col}} + 4),$$
$$P_x = 4\big(W_{\mathrm{lay}} + 8(V_{\mathrm{s}} + 1) + P_y + 32\big), \qquad L = 4P_x + 100.$$

Vertex k in that owner cell sits at $(a_{\mathrm{own}}P_x + W_{\mathrm{lay}} + 8k + 4, b_{\mathrm{own}}P_y, 0)$, with distinct west, east and north ports for its incident edges. For port coordinates $(x_1, y_1), (x_2, y_2)$, ordered vertex rows $u \le v$, and color j, set

$$T = P_x\lfloor \min(x_1, x_2)/P_x \rfloor + j, \quad a = u + 8j + 8, \quad b = v + 8j + 12,$$

$$B_{\mathrm{wire}} = 8 + (a - y_1) + (b - y_2) + (b - a) + (x_1 - T) + (x_2 - T), \qquad h_{\mathrm{wire}} = (L - B_{\mathrm{wire}})/2.$$

The path uses its first port, rises to layer one, reaches row a, drops, goes right by h_{wire}, turns one row, goes left to T, rises, follows the trunk to b, drops and goes to the second port column, rises, returns to its row, drops and reaches the vertex. The two endpoint and six via steps account for the eight in B_{wire}. The length is L; parity makes B_{wire} even. The bounds $P_y < P_x/4$, $x_i - T < 5P_x/4$ and total vertical contribution less than $3P_y + 2$ give $h_{\mathrm{wire}} > 0$. Distinct colors have disjoint lane bands and trunks. Each route spans less than $4P_x$ horizontally and $2P_y$ vertically, so equal-color edges, whose owners are separated by four cells, are disjoint. Short port columns are private to their incidences. Thus interiors are vertex-disjoint. Put the old sign on the first interlayer via; these vias form a matching.

For the present finite graph the spacings and length are

$$4P_x = 82\,606\,081\,152, \qquad 4P_y = 17\,686\,381\,696, \qquad L = 82\,606\,081\,252. \tag{58}$$

These numbers belong to the 23 573 144-spin circuit graph. At inverse temperature β' a unit-bond path transmits $(\tanh\beta')^L$, so it exactly reproduces the trivalent marginal at inverse temperature $\operatorname{atanh}((\tanh\beta')^L)$.

Replace this diluted bilayer by a fixed fully bonded graph. At each central vertex in layer zero partition ports into $\{W\}, \{E, S\}, \{N, R\}$; in layer one use $\{N\}, \{S\}, \{R\}$. The route just described uses at most one direction per group. Indeed horizontal motion occurs only in layer zero and motion in the y direction only in layer one, apart from the one-step north

ports of original vertices. In layer zero an original vertex uses W, E, N; a straight horizontal run uses W, E; the two vertices of a rightward turnaround use W, N and W, S; a via uses R with one of W, E, S, where S occurs only at the end of a north port. Thus E and S are never used together, nor are N and R. In layer one there are no x-directed links: a vertex uses N, S on a vertical run, or R with one of N, S at a via. Attach three forks to each central vertex by positive stems. The possible intercell links are east and north in layer zero, north in layer one, and a rung. Pair the four unused fork contacts by two extra neutral candidate links. This has eight vertices and twelve candidate edges per geometric cell. In-cell links may cross; the graph is nonplanar.

Keep the six stems. Replace each of the other six links $a - b$ by

$$a - h,\ k - b,\ h - p,\ h - q,\ p - k,\ q - k,\ p - q.$$

All-positive signs give positive transmission; negating $q - k$ makes it exactly zero. Indeed swapping p, q and flipping k, b preserves the conditional sum, making it independent of the terminal product. For a negative rung transmission negate $a - h$ instead. All candidate $q - k$ edges and the possible negative rung contacts form a fixed matching. No other bond is negative. The graph has $8 + 6 \cdot 4 = 32$ vertices and $6 + 6 \cdot 7 = 48$ edges per cell, all degrees exactly three, for every sign choice. The routing embeds, by internally disjoint paths in this same host, any periodic graph of maximum degree three whose edges join a cell to itself or to its east or north neighbor, in particular one copy of a finite cubic graph per cell. Taking $K_{3,3}$ proves nonplanarity.

If $t = \tanh \beta$, the active seven-bond switch transmits

$$\tau(t) = \frac{2t^4(1 + t)}{1 + 2t^3 + t^4}. \tag{59}$$

Its denominator counts even subgraphs and its numerator counts odd-boundary subgraphs. It increases strictly from zero to one: its derivative, after positive factors are removed, has numerator $4 + 5t + 2t^3 + 4t^4 + t^5 > 0$. The two stems multiply it by t^2. Thus the full physical-to-parent map for I is

$$\mathfrak{t}(\beta) = B\left(\operatorname{atanh}\left([t^2\tau(t)]^L\right)\right). \tag{60}$$

It is computable, analytic, strictly increasing and computably invertible on its image. Rational directed enclosures at the two window endpoints give

$$e^{-\mathfrak{t}(87/4)} \in (3/4, 19/25), \qquad e^{-\mathfrak{t}(45/2)} \in (13/20, 2/3).$$

The map stays in the parent's window and straddles $\mathfrak{t}_0$: the cut equation is below one at $z = 3/4$ and above one at $z = 2/3$, for every $\vartheta \in [1/2, 7/12]$. Consequently $\beta_\star = \mathfrak{t}^{-1}(\mathfrak{t}_0)$ is interior, right-c.e. and of halting degree.

The printer `crystal.py` takes a geometric cell and an edge ordinal $0, \dots, 47$, returning its two endpoint triples and sign. Its sign lookup uses only finitely many compressed route segments. The signed period is $(4P_x, 4P_y)$, containing $32(4P_x)(4P_y)$ spins. The chosen original site is $(552699460, 0, 0)$; the displayed vector is $(0, 4P_y, 0)$. Exact tracing gives

$$\mathcal{C}_\beta(n) = C^{\mathrm{bin}}_{\mathfrak{t}(\beta)}(4n). \tag{61}$$

Strict positive decrease survives, and $\cos(8\pi n/3) = \cos(2\pi n/3)$. The distinct simple conjugate poles remain distinct and simple under fourth-power subsampling. On the monotone side the positive stage coefficients on the subsequence still have rate $\mathfrak{b}^4$ and vanishing normalized prefactor. The preceding normal-convergence and Abel arguments apply to this subsequence, leaving precisely the real simple pole and the two rotated non-poles. This proves all correlation assertions of I.

14.7 Arbitrary microscopic exteriors

It remains to transfer uniqueness and mixing rather than merely identify one infinite-volume law. Given an arbitrary finite microscopic free set, first sum out every neutral switch whose four interior spins are free; by the swap-and-flip symmetry its sum is a constant independent of both terminals. A factor includes its routed wires, their active switches, and the forks and neutral switches attached to its vertices. Let Λ_{p} consist of parent sites whose visible spins and all three incident factors lie wholly in the free set. Condition on collar visible spins and on private spins in incomplete factors. All factors touching Λ_{p} can now be traced in full. Host spins outside all factors meet the rest only through neutral switches. At distance greater than a fixed collar width from the microscopic exterior, all neighboring neutral-switch interiors and the corresponding central–fork stars are free. After the neutral traces these stars are isolated and exterior-independent; observables in the traced switch interiors are recovered by fixed local kernels. The remaining background cases lie in the fixed-width collar, where the trivial total-variation bound suffices. The conditional law of decoded symbols in Λ_{p} is exactly the finite-volume parent Gibbs law with boundary symbols read from the collar. Factors involving only collar sites are constants, and rail gauge coordinates have uniform fibers.

The parent estimate (48), or (44) for H, is uniform in all those boundary symbols and in arbitrary finite shapes. A microscopic probe reduces, through its finite conditional decoration kernels, to a bounded function of nearby visible symbols. Averaging the uniform comparison over the random collar preserves the bound. The finite factor diameter and the fixed period convert parent distance to microscopic distance with fixed multiplicative and additive losses. This proves arbitrary-volume microscopic EWM and, by exhaustion, uniqueness among all DLR states. Global spin flip then gives zero means.

On periodic boxes the partition sum is the parent's sum times finite positive analytic normalizers and constant gauge multiplicities. The same fixed-collar surface comparison gives the infinite-volume pressure. Finite microscopic marginals are finite analytic combinations of the rational parent marginals and the conditional decoration kernels. This proves their real analyticity as well as that of the pressure. In particular uniqueness, EWM and analytic pressure all persist at the disorder line.

14.8 Finite data, margins and lag units

At rational $y = e^{-2\beta}$, every finite decoration transmission is rational, including the integer wire power. Thus $e^{-\mathfrak{t}(\beta)} = 1 - 3\mathcal{T}/8$ is rational. A fixed-lag parent correlation is a rational function of that coordinate and can be evaluated exactly. At a computable inverse temperature, the same formulas give uniform computability of each fixed-lag value.

Suppose a total computable $N(k)$ supplied a negative sampled lag by $N(k)$ for every $\beta \geq \beta_\star + 2^{-k}$. Choose a rational-$y$ grid with adjacent inverse temperatures less than 2^{-k} apart, covering a computable interior bracket for the cut. All signs up to $N(k)$ at these grid points are exactly decidable by (61). Every witnessed point is above the cut, and some point between $\beta_\star + 2^{-k}$ and $\beta_\star + 2 \cdot 2^{-k}$ must be witnessed for all large k. The least witnessed grid point approximates $\beta_\star$ with computably shrinking error, a contradiction. Uniform EWM supplies a fixed geometric envelope Mq^n, $q < 1$, for the sampled correlation. A computable positive magnitude guarantee $\varepsilon(k)$ would restrict every such witness to a computable finite lag range, giving the same contradiction. Fixed constants in this argument need not be efficiently small or uniformly recoverable from arbitrary model data; they belong to this one model.

There is also a window-uniform lower bound on the first parent witness. From (53), for

$13 \le n \le 10^{330}$,

$$\frac{\sum_j w_j \mathfrak{B}_j(n - F_j)}{\mathfrak{K}(n)} \le \frac{2^{-2680}}{1 - 2^{-15}} \exp(2n\, 16^{-272}).$$

Here $\omega \mathfrak{b}^{-13} > 2^{-1400}$ and $F_j \ge j + 272$ were used. The integer inequality $2 \cdot 10^{330} < 1024 \cdot 2^{1088}$ makes the exponent less than 1024. Since $e < 4$ and $(1 - 2^{-15})^{-1} < 2$, the ratio is less than 2^{-631}. Thus the factor $2\cos(2\pi n/3)$ cannot overturn the reference term. Earlier lags have no stage contribution. This proves positivity through 10^{330} parent lags everywhere in the parent window. The scale 10^{2536}, obtained by estimating the stage-zero excess and entrance weight at parent inverse temperature $1/2$, is an estimate, not a certified onset.

One sampled lag is four parent lags and $D_{\mathrm{geom}} = 17\,686\,381\,696$ geometric spacings. In I, the preceding bound therefore ensures positive sampled correlations whenever $4n \le 10^{330}$, at separation nD_{geom} spacings. The cosine period is three in units of d; no claim is made about a three-spacing oscillation. These observations prove Proposition 1.1 and explain its placement beside I.

15 Discussion

The diagrams in this paper separate three questions about finite local tables. Theorems A and C encode halting in temporal forgetting on the line. Theorem B transfers a programmed-window dichotomy to mixing and magnetic order for one square-lattice interaction. Theorems D–H ask which complete sets of inverse temperatures one can prescribe for mixing, point-field response and sampled decay shape; Theorem I places a halting-degree cut in one explicit crystal. Among magnets unique throughout the window, the sets with a negative point-field response somewhere are exactly the relatively effectively open sets, each realized by a single magnet. For the full EWM set, the same effective upper bound is known, with exact realization for explicit semialgebraic diagrams and sequential realization in general.

Theorem G realizes its response-sign set without leaving the mixing regime: it has one DLR measure, arbitrary-volume EWM and the real-analytic pressure $\log g_0(e^{-\beta})$ at every inverse temperature in the window. Off the prescribed set every response is nonnegative, and all sufficiently long column responses are positive with one uniform onset. On the set, all sufficiently long column responses are negative, with an onset that diverges on approach to its relative boundary. These are disorder lines in the correlation sense; the normalized-PCA reduction, or disorder variety, holds throughout the window. When the boundary is noncomputable, the information is carried by the family of local responses over all lags. The fixed linear probes in Section 11.6 have analytic response functions whose isolated real zeros are individually computable.

Theorem H prescribes the oscillatory set of one spin correlation sampled along one period vector on a fixed Ising geometry. Its converse uses uniqueness and the all-lag sign promise; G's converse uses uniqueness alone. The relation is through these sampled-correlation sets. In I, no assertion concerns the response at all displacements. The explicit crystal is constructed from its own binary-energy parent, and is not obtained by applying H to a cut.

In I, uniqueness among all DLR measures, arbitrary-volume EWM and analytic pressure hold across $\beta_\star$. What changes is the asymptotic shape of the specified correlation: positive decreasing decay through the cut, then a cosine of period three in units of d. On the monotone side the real pole shares its convergence circle with two non-pole singularities. The graph is nonplanar and has degree three; its 32-site geometric cell and its approximately 4.7×10^{22}-spin signed repeat describe different scales. The displayed vector spans approximately 1.8×10^{10} spacings. Every finite-lag correlation is computable, but neither a lag bound nor a positive signal bound can depend computably only on distance above the cut. The no-negative bound

through 10^{330} parent lags applies throughout the parent window. These are asymptotic correlations rather than a finite-data method of locating an uncomputable real.

Theorems D and E leave the following single-magnet realization problem in its full form: for rational $0 < a < b$, does every relatively effectively open $W \subset [a, b]$ equal $\mathcal{M}_\Phi$ for one finite rational magnet, with opposite translation-invariant ordered states on its complement? A sharp case asks for EWM on $[a, \beta_c)$ and order on $[\beta_c, b]$, where $\beta_c = a + (b - a)\Omega$ and $\Omega \in (0, 1)$ is the halting probability of a fixed universal prefix-free machine. The lower cut of Ω is effectively open. Theorem G realizes that cut for the response-sign observable, but its unique mixing state supplies no ordered complement. Purcell, Li and Cubitt realize a ground-state phase transition at a halting probability for a family of two-dimensional quantum Hamiltonians [43]; their varied parameter enters the couplings. Theorem B already gives one classical magnet with a halting-programmed diagram on disjoint inverse-temperature windows, and asserts nothing between those windows; it does not determine this cut as an EWM set.

Two obstructions explain why the present constructions stop short of that mixing-phase problem. Theorem F settles diagrams given explicitly as rational semialgebraic sets in $z = e^{-\beta/d}$. Its finite rational line-transfer mechanism has an EWM set semialgebraic in a Boltzmann coordinate, so its transition values are exp-algebraic: $e^{-\beta/d'}$ is algebraic for some positive integer d'. In particular those values are computable. The algebraic temperature substitutions sketched in Section 10.8 would retain this limitation. This ceiling belongs to the finite-table mechanism, with no restriction on general finite magnets or the programmed windows of Theorem B. The alphabets and index lists produced by its synthesis may be extremely large; all searches terminate, and no elementary size bound is claimed. In the response construction, Proposition 12.1 gives a different obstruction: strict positivity and a finite hidden Markov stationary row force full exponential forgetting and EWM. Keeping the iid hidden row under positive amplitude increases, or under strictly positive finite decorations freshly sampled from the hidden neighborhood and ignoring their own earlier values, therefore cannot create magnetic order. Decorations with memory or feedback require a separate analysis.

In Theorems A–C, the rational windows provide a computable map from programs to promised regimes. Their conclusions do not require a decision criterion at noise rates or temperatures outside those windows. In the dynamical problem, the observed periodic signal gives a cycle of row laws for the time-one kernel; its phase average is stationary. The witness supplies nonconvergence, and the number of stationary row laws on the nonhalting side remains undetermined by this argument. The corresponding Gibbs states of the space–time interaction distinguish the phases separately because the DLR specification permits time-translation symmetry breaking. Blocking by the clock period turns them into translation-invariant states, and doubling adds an exact symmetry that exchanges them.

Several restrictions on the models remain natural to investigate. One is how small the alphabet and range can be for a single rule whose convergence diagram encodes halting. Another is which arrangements of halting and nonhalting windows can be obtained with one fixed automaton away from the small-noise regime. For Theorem B, the finite-range square-lattice interaction is built from an anchored pattern of one child and several parents. It would be useful to determine whether the same all-DLR dichotomy can be realized with only pair interactions at nearest-neighbor range and finite energies.

How this work was done

The assistance named on the first page was used in developing and checking the mathematical arguments. The proofs distinguish deterministic implications, finite probability estimates

and explicitly stated external theorems. Independent mathematical checks informed the statements; exact finite calculations provide supplementary checks of the displayed examples. For the Ising reduction, integer identities certify every input of the finite Boolean circuit; the geometric and infinite-volume arguments are proved separately. The reproducibility appendix states the scope of each software check. The authors take responsibility for the proofs, their scope and the use of the cited results.

A Numerical choices

Fix the instruction set, finite track order, tape template and polynomial time allowance first. Choose w sufficiently large for every polynomial comparison with P_0, every constant-times-A comparison with v, and the collar inequalities. A lower bound of 10^{10} is compatible with all specified finite lower bounds, and w may be increased to satisfy the interpreter allowance and the fixed collar requirements. All native lengths are then fixed. In particular $d_1 \ll P_0$, and any fixed polynomial in w is $o(P_0)$. The estimates use fixed proof constants independent of word width. The following table lists the choices and the larger allowances used when estimates overlap.

Quantity	Value and role
q, A, B, v	$w^2, w^3, 128q, w^4$, respectively. They determine the propagation length, Shrink thresholds, provisional pass and Grow rounds.
P_0, P, T	$2^w, 5P_0$, and the duration of the fixed clock in Section 3.2.
H_0, d_0, d_1	$2q+2, (2A+4)H_0, qd_0$. These are the handle, Main cycle and active Shrink spacing.
$r_{\mathrm{last}}, s_0, s_1, s_c$	Start of final withdrawal, start of ordinary phase, end of Shrink(0), and the service date in early computation.
m, d, incoming pad	$m = \lfloor w/5 \rfloor$, $d = 4w+8$, and $64A$ child units. Common span $3w + 2m + 4 + 2(608) < d$: the $2 \cdot 608$ term is fixed slack; owner-level displacements are paid by the m-owner collar ($608 < m$) and do not enlarge K.
M, verifier threshold	$M = 128d$, count strictly greater than $4d$, epoch length $M+1$.
Vacant-band thresholds	$4B+4$ at all retired strobes, $8B+8$ before a partial pass. At most B positions can be entered from each side of a pass, so the central band stays vacant. Width $4B+4$ works throughout the ordinary and terminal estimates.
End gap and old material	Use $16B + 2d + 24$ in the early certificate and comparison with Δ. After removal of K and boundary positions it leaves a vacant band of width $4B+4$; also $\Delta < 16B + 2d + 24 < 3\Delta$.
Δ and the survivor budgets	$\Delta = 1024q$; the budgets are $3A + 3\Delta$ and $6A + 3\Delta$. The assignment allowance is $9d + 80$; surviving faces advance at least $B - 20$ in the partial pass.
Common full seed	Derived $100Ad_1$ bound per home and $N = 200Ad_1$ as uniform search allowance. For $b_{\mathrm{seed}} = 6N + 4w^6 + 20$, search $b_{\mathrm{seed}} + 1$ trials of $5d_0 + 1$ dates. The transport margin adds the relaxation and preparation delays to this bound.

Quantity	Value and role
Source density	$L = 60B + 3d + 40$ for the vacant interval comparison, at most five labels and $4d + 12$ components per label. The resulting deficit is at most 2,000,000A.
Terminal bands	Thin radius $100B$, near-source radius $300B$, wide half-width $v/5$. Permanent gain at most $35B$ across the event. Delete at most $(220A + 2w + 20)d_0$ terminal comparison dates.
Endpoint repair	Search length $3vd_0 + 2w^6 + 12d_0$, clear zone $100q$, retained guard at depth below $97q$. Use $4q + 4$ inspections for one search and $8q + 12$ for two. Strip length $17q + 25$ leaves a $q + 1$ backing after two-sided losses, with the nearest witness within $47q$. Final common exception has width at most $d + q + 42 < 2q$.
Directed advance	Startup in $6d_0$ clean ticks, allow $12d_0$ with the pulse. Cost $d + 4(q + 20) < 5q$ old sites and six cycle equivalents; overshoot $P + 10v$.
Nested tests	$M_0 = \lfloor 2w/5 \rfloor$; fixed multipliers with four rounding child units at each end; $M_0 \geq 2m \geq m + 3$ places W_0 inside the test. The centered cap test at rank j has half-extents $(4T/P + 4 + m + 2)P^{j+1} + 4P^j$ in space and $(6 + m + 2)T^{j+1} + 4T^j$ in time. The multipliers relative to V^{j+1} are independent of j and K.
D_0, H_n	Fix D_0 after P, T, w and test multipliers; set $H_n = D_0^{n+2}$. Fix the paint multiplier at this stage.
M_s, b_h, v_h, s_h	Choose the required short exponent and then its refresh, replica and reader counts. Require $b_h > 20$ and strict exponent slack in Lemma 4.5.
M_l, m_h	After b_h fix the long admission exponent and choose m_h for the long estimate. Use odd reader and channel counts.
$D_{\mathrm{lane}}, \ell_0, n_0$	Choose the stride and base word width after the helper counts. Increase n_0 last for the uniform probability and time bounds on all I_n.

A.1 Notation and local scopes

The following conventions distinguish global parameters from the explicitly local variables reused in individual calculations.

Notation	Meaning and scope
$\Delta, \Delta_{\mathrm{leaf}}$	$\Delta = 1024q$ is the ordinary depth margin; Δ_{leaf} is the enlarged-leaf overlap degree in Section 4.4.
H, H_0, H_n	H denotes a spatial home in organization proofs, the suffix length $\lceil p^{-6} \rceil$ throughout Section 4.2, and a cone height where explicitly stated. $H_0 = 2q + 2$ is always the handle; $H_n = D_0^{n+2}$ is a stripe-box radius. The decoding exception is I_{dec}.

Notation	Meaning and scope
L, P_0	L is a local gap bound, strip length or time horizon, defined afresh in its lemma, the segment length in Sections 4.2–4.3, and the leaf count $L = 2^{i-j}$ in Section 4.4; the Grow label L in the advance proof denotes the left donor. $P_0 = P/5$ is the packet length, including in the classification calculation.
D, D_{tr}	D denotes the dead symbol in state expressions. In numerical proofs D is explicitly a pair count, a $35B$ gain, a rear-wall interval, or the advance debit $d + 4(q + 20)$. D_{tr} is the ballot transport margin. $D_0, D_j, D_{\mathrm{lane}}$ have their displayed geometric, recursive and stride meanings.
$b_0, b_j, b_{\mathrm{dst}}$	$b_0 = 0, b_j$ are service lookahead. The destination coordinate in the ballot path equation is b_{dst}. The probability b_j in the iid sparsity proof is local to that proof. Other local b's denote a scale base, lock rank, interval end or epoch offset, as explicitly introduced.
W, W_0	W is $4B + 4$ in sterile-band proofs, $\lceil p^{-3/2} \rceil$ in Sections 4.2–4.3, and the explicitly displayed end-search duration in the end ledger. W_0 is the two-period region of Lemma 5.24, of extent $(2m + 8) \times (2m + 3)$ owner units.
K, S, S_v	$K \times S$ is the common exceptional column and time interval of the one-rank argument in Section 5 and in the table above. In Sections 6.2 and 6.3 K is the halted rank defined there. S_v is the support of an anchored term in Section 7.
M, M_0	$M = 128d$ is the verifier counter bound and $M_0 = \lfloor 2w/5 \rfloor$ the nested-test margin. In Section 7 M is the odd one-site observable of Theorem B.
N, p, Y	N is a nonfull-date allowance, a local Grow-cycle count, the length of a phase mean in Section 6.5, or the alphabet size in Section 7. The probe rate is p; a local endgap also called p is confined to ordinary geometry. The PCA probability table is $p(a \mid \xi)$. Y is the native carrier; Y_i are autonomous reference channels.
v, w, β_-, β_+	$v = w^4$ and w are native geometry parameters; v also denotes a space–time site in (3) and Section 7; $\beta_- < \beta_+$ are rational inverse-temperature endpoints.
$\Pi_T, \Phi^{[T]}, \Psi, \sigma$	In Section 7: the blocking map, the blocked interaction, the interaction of Theorem B and its involution. There $\mathcal{B}$ is the alphabet of Ψ, υ the uniform law of the free bits and θ_x the spatial shift.
Nonfull	At least one of a home's P positions lacks an actual matching served child at that fine date.

B Case tables for the numerical bounds

The rows below collect the timing reserves used in Lemmas 5.23, 4.3 and 4.4. The corresponding calculations are in the proofs of those lemmas. Each row uses the same native clock.

B.1 Directed interruption

State at interruption	Repair	Timing and vacancy bounds
Startup	Restart from the K-avoiding ordinary backing and use a complete clear pass.	At most $12d_0$ from the cut. Discard one cycle.
Free attack or free withdrawal	Compare to the last unmarked reference row. Rejoin its persistent guards. An uncommitted pair of provisional tails uses no reference positions.	At most $D = d + 4(q + 20)$ old positions and two touched cycles, with at most v additional retreat to the next complete cycle.
Established contact	Restore the interval connecting the two fixed guards. Each branch pays for its own short detached end.	At most D old positions. The later free displacement calculation is unnecessary if the Grow repair fits; otherwise use the final-withdrawal row.
Lone attained overshoot	Restore the backed free prefix or the genuinely joined one.	The restored prefix reaches at least $P + 9v - D > P + 8v$, before the next withdrawal.
Final withdrawal and normalization	Keep guards inside the target and switch to the in-home repair. Each old face was trimmed by ten sites.	At most two early Shrink inspections, included in the guard allowance.
Ordinary work	Use the old full target, its surviving guards, or the already observed service sample.	The sample has at least $P - D > P/2$ matching sites.

B.2 Reader and reference estimates

Quantity	Choice, after the data geometry has been fixed
Primitive pools	Thirteen wall probes, five mode probes, three groups of 59 replicas per reader. Thirty flags are needed to affect a decoded bundle. A reference group uses $b_h + 2$ probes and $2v_h - 1$ replicas.
Wall and mode rates	Wall refresh p^{12}, new wall choice Bernoulli p, CLEAN probability p^5, RUN probability $p^3(1 - p^2)$.
Reader radius and buffer	$W = \lceil p^{-3/2} \rceil$, $\ell'_0 = \lfloor p^{-1/2} \rfloor$, with the fixed extra boundary sites of Lemma 4.3.
Reader time	$B_0 = \lceil p^{-14} \rceil$ sets the prefix, $H = \lceil p^{-6} \rceil$ the suffix, and $S = \lceil p^{-4} \rceil$ the time budget. A fully blanked segment gets at least $S/2$ ticks to publish.
All-start mixing	Eight suffix slabs, enduring walls in the two exterior strips, and a CLEAN period of at least $5W + 10$ in each isolated segment.

Quantity	Choice, after the data geometry has been fixed
Anchor and cutoff	$A_n = 8^{n+n_0}$, $k_n = \lfloor A_n/8 \rfloor$, $R_n = 2^{\lfloor A_n/4 \rfloor}$, and the stripe width P^n.
Short and long supports	Short radius p^{-16}; long radius $p^{-(b_h+5)}$, including the reference refresh prefix of length $\lceil p^{-(b_h+3)} \rceil$. Leaf cutoff exponents 17 and $b_h + 6$ in the minimum side.
Tree and chain constants	Bounded child-pair count $8(P+2)^2(T+2)^2$, enlarged-leaf overlap degree fixed by geometry, and the geometric count $C_1 r_\lambda^4$ for each ordered ball. Their use precedes the final small-p choice.

C Dependencies of the lemmas

The table is the exact directed `\ref` graph on the labelled lemmas and Proposition 5.3, with self-references omitted, using each statement and its proof; the organization row includes its surrounding induction explanation and final assembly. It has 35 nodes and 90 edges. An arrow means that the row text references the result in its Uses column, including explanatory forward pointers. Used in is the exact reverse relation. References to definitions, equations, sections and theorems are outside this induced graph. Unlabelled native-rule and lower-rank premises remain as stated in each proof. The theorem-level arguments are also outside the graph: Section 6.3 uses Lemmas 4.5, 4.6, 6.1 and 6.3, and Section 7 uses Lemmas 6.4 and 6.5 with Lemmas 7.2 and 7.3.

At rank j, references to Proposition 5.3 within an induction step invoke a rank-$(j-1)$ clause or identify the current conclusion being proved; the assembly invokes the established rank-j lemmas. In the birth branch the computing test is proved before strong issue, incoming clearance and successor properness are used.

The supplied-interval Lemma 5.7 reproduces the first-clean-Purge count independently and uses only lower transitions and held-flank clearance. The held-flank proof fixed-backing pointer goes to the independent step 2 of Lemma 5.8, so this longer textual cycle also supplies no missing premise.

Every pair of reciprocal textual references has the following interpretation.
Proposition 5.3 and Lemma 5.10 are reciprocal only because the lemma invokes a lower-rank clause or names the current conclusion, while the assembly uses the completed lemma at rank j.
Proposition 5.3 and Lemma 5.21 are reciprocal only because the lemma invokes a lower-rank clause or names the current conclusion, while the assembly uses the completed lemma at rank j.
The transition/ascent pair: the transition proof only points to the separately centered cap tests of Lemma 6.3; the ascent uses the transition conventions.
The flanks/ordinary-response pair reuses only the fixed-backing induction in step 2, which is closed independently of the local response or repair.
The flanks/seed pair concerns the paired-seam application conditional on an already supplied full seed; obtaining that seed uses only the independently proved backing and ordinary geometry.
The flanks/attachment pair points from the held-wall mechanism to its issued-source application; attachment supplies the walls before invoking clearance.
The ordinary-response/seed pair points forward to the count of open-end episodes; the ordinary proof establishes the nonlatching geometry before that independent counting application.

The ordinary-response/attachment pair is an application pointer from ordinary work; the local calculation is proved earlier as Lemma 5.7 without an issued-source conclusion.
The attachment/components pair uses only the independent clean-evolution paragraph in the forward direction; the component count receives its initial source interval from attachment.
The attachment/remote pair points to an independently rederived sterile-band induction; source attachment does not use the remote-source conclusion.
The competitors/calendar pair uses only the independent old middle-wall selection, which is made without births or the terminal-exclusion conclusion.
The calendar/end-ledger pair uses the $97q$ guard only in the conditional filling remark; the end ledger selects that guard before invoking the clear calendar.

Longer textual cycles reduce to these pairs. Delete the pointer or independent-sub-argument direction of each pair (flanks to ordinary-response, seed and attachment; ordinary-response to seed and attachment; attachment to components and remote; competitors to calendar; calendar to end-ledger; transition to ascent) and every reference from a lemma of the one-rank induction to Proposition 5.3, which invokes a rank-$(j-1)$ clause or names the conclusion currently being established. The remaining graph is acyclic. In particular the cycle attachment–remote–thin-density–attachment is broken at the attachment/remote pointer, the cycles through the rank-$(j-1)$ classification clause cited by Lemma 5.5 are broken at that clause, and every cycle through Lemma 5.4 is broken at its pointer to Lemma 6.3.

Result	Uses	Used in
3.1	—	6.4
4.2	—	4.6
4.3	—	4.4, 4.5
4.4	4.3	4.5
4.5	4.3, 4.4	4.6, 6.4
4.6	4.2, 4.5	6.5
5.3	5.4, 5.10, 5.11, 5.17, 5.21, 5.23	5.5, 5.10, 5.12, 5.21, 6.3, 6.5
5.4	6.3	5.3, 5.5, 5.7, 5.22, 5.23, 5.24, 6.3
5.5	5.3, 5.4	5.8, 5.9, 5.12, 5.14
5.6	5.8, 5.9, 5.12	5.7, 5.8, 5.9, 5.10, 5.12, 5.17, 5.21
5.7	5.4, 5.6	5.8, 5.12
5.8	5.5, 5.6, 5.7, 5.9, 5.12	5.6, 5.9, 5.10, 5.12, 5.14
5.9	5.5, 5.6, 5.8	5.6, 5.8, 5.10, 5.17, 5.21, 5.24
5.10	5.3, 5.6, 5.8, 5.9	5.3, 5.12, 5.16, 5.17, 6.2
5.11	—	5.3, 5.17
5.12	5.3, 5.5, 5.6, 5.7, 5.8, 5.10, 5.13, 5.15	5.6, 5.8, 5.13, 5.14, 5.15, 5.20, 5.21
5.13	5.12	5.12, 5.14
5.14	5.5, 5.8, 5.12, 5.13	5.15, 5.16, 5.19
5.15	5.12, 5.14	5.12, 5.16, 5.19
5.16	5.10, 5.14, 5.15	5.17, 5.19, 5.21
5.17	5.6, 5.9, 5.10, 5.11, 5.16	5.3
5.18	—	5.21, 5.23, 5.24
5.19	5.14, 5.15, 5.16, 5.20	5.20, 5.21
5.20	5.12, 5.19, 5.21	5.19, 5.21
5.21	5.3, 5.6, 5.9, 5.12, 5.16, 5.18, 5.19, 5.20	5.3, 5.20
5.22	5.4	5.23
5.23	5.4, 5.18, 5.22, 5.24	5.3
5.24	5.4, 5.9, 5.18	5.23, 6.5

Result	Uses	Used in
6.1	—	6.2, 6.4, 6.5
6.2	5.10, 6.1	6.3
6.3	5.3, 5.4, 6.2	5.4
6.4	3.1, 4.5, 6.1	6.5
6.5	4.6, 5.3, 5.24, 6.1, 6.4	—
7.2	—	—
7.3	—	—

D Software checks: scope and results

The accompanying programs use Python 3.10 or later and only its standard library. They specify finite objects and replay exact identities. The mathematical arguments identify the Gibbs state and its infinite-lag asymptotics.

Object	Exact check and scope
Parent table	The printer gives 164 852 signed Walsh terms with total absolute coefficient $33\,293\,550\,733\,133\,931 < 2^{56}$. A separately coded assembly and scattering interpreter produces the same full multiset.
Boolean circuit	Regenerate the 85 041 ordered gates. Exact symbolic normalization, 315 decoder identities, the east quadratic, 68 arithmetic blocks, the full north polynomial and all 1024 selection rows prove its function on all 2^{323} inputs.
Parity core	Integer terminal polynomials and a nonnegative derivative polynomial verify the core's transmission and its strict monotonicity.
Crystal printer	The checks verify local degrees and the fixed matching, contract each switch type at rational transmissions, test the port cases and representative route bends, and certify the rational temperature endpoints. Its description gives the exact signed-cell size and sampled vector.
Supplementary identities	Exact rational arithmetic checks G's stage-start copy, the general selector coefficient bound, the return-series factorization and the integer estimate used in the 10^{330}-lag bound.

What is exact. Gate truth tables, coefficient dictionaries, tuple multisets, compressor identities, polynomial signs in the finite core certificate, and SHA-256 digests use integers. The named temperature endpoint checks use rational outward enclosures, repeated squaring and a Taylor remainder bound. The accompanying certificate proves a Boolean identity for the specified gate list; the thermal reduction is proved by the finite partition identities in Section 14.5.

What is approximate. The descriptions 4.7×10^{22} spins and 1.8×10^{10} spacings round displayed integers. The scale 10^{2536} is only an onset estimate. None of these approximations is used in a proof.

What these calculations do not decide. They do not approximate the halting set, evaluate a correlation to unbounded lag, or execute the general enumerator-to-Ising construction of H. Sampled route checks supplement the disjoint-lane proof; they do not enumerate the signed crystal's entire repeat. Uniqueness, arbitrary-volume EWM, the circle singularities and the finite-data impossibility theorem use the mathematical arguments. In particular the certificate alone makes no assertion about a thermodynamic limit.

E Reproducing the finite objects

From the `programs` directory, run

```
python -I -B export_circuit.py
python -I -B reference_scattering.py
python -I -B certify_circuit.py
python -I -B verify_results.py
python -I -B crystal.py check
```

The first command regenerates `herald-gates.json`, `binary-terms.tsv` and the untrusted component descriptions in `circuit-hints.json`. The second requires equality of the independently coded parent reference and the printed table, reporting 235 register instructions, 5731 payloads and 164 852 options. The third verifies every proposed component before using it and writes `CIRCUIT-CERTIFICATE.json`. It reports `PASS`, 323 input bits, 85 041 final gates, 63 011 raw gates, 68 arithmetic blocks, 96 population calls and 1024 selection rows. The fourth checks the delivered object digests and supplementary identities. The fifth prints a JSON record including

```
spins_per_geometric_cell: 32
bonds_per_geometric_cell: 48
signed_cell_spins: 46752085813280748601344
degree: 3
observed_site: [552699460, 0, 0]
lag: [0, 17686381696, 0]
```

It also checks the named window $[87/4, 45/2]$. A single geometric bond is obtained, for example, with

```
python -I -B crystal.py edge 0 0 0
```

No enumeration of the huge signed repeat is needed for that lookup. `binary_wave.py describe` prints the parent alphabet and coefficient count; its `horizontal`, `vertical` and `spin` commands query the finite tables by symbol ordinal. `herald_ising.py` and `square_ising.py` give the intermediate finite graph and routed realization. Imported modules define their register instructions, rail encoding, parity blocks and compressed route segments. `cubic_parity.py` regenerates the finite core's terminal polynomials and derivative checks. `core_polys.json` lists their coefficients.

Object identity. The full SHA-256 digests reproduced by the certificate are

```
herald-gates.json
a069b637730ec10e48d7402608e929f1e71
c8b558164388cdd07047a57eff2cc
binary-terms.tsv
d59dd512ce13ddbd6ab180b72634f277fdc
cc22b53861ff65c0060483aa84c2d
```

Concatenate the two lines of each digest. These digests identify the compared objects; exact equality decisions within the certifier use full integer or tuple data. The coefficient multiset comparison includes unreachable contexts. The reference is independently coded and checked equal to the printed table.

Building the manuscript. From the manuscript directory run the following command three times:

```
pdflatex -no-shell-escape -interaction=nonstopmode -halt-on-error main.tex
```

The bibliography is included directly. No bibliography program or shell escape is required. The formal results for the earlier higher-dimensional constructions belong to the version specified in the introduction; this distribution adds finite Python certificates rather than a new formalization.

References

[1] J. Bausch, T. S. Cubitt and J. D. Watson. Uncomputability of phase diagrams. *Nature Communications* 12 (2021), 452. doi:10.1038/s41467-020-20504-6.

[2] A. Bušić, J. Mairesse and I. Marcovici. Probabilistic cellular automata, invariant measures, and perfect sampling. *Advances in Applied Probability* 45 (2013), 960–980. doi:10.1239/aap/1386857853.

[3] T. S. Cubitt, D. Pérez-García and M. M. Wolf. Undecidability of the spectral gap. *Nature* 528 (2015), 207–211. doi:10.1038/nature16059.

[4] E. Domany and W. Kinzel. Equivalence of cellular automata to Ising models and directed percolation. *Physical Review Letters* 53 (1984), 311–314. doi:10.1103/PhysRevLett.53.311.

[5] S. Douglas and N. Mghirbi. Forgetting under noise is undecidable: ergodicity of positive-rate cellular automata at a fixed noise rate. Preprint, version 3.0.0 (2026). doi:10.5281/zenodo.24265098.

[6] S. Douglas and N. Mghirbi. One noisy cellular automaton with an uncomputable noise diagram. Preprint, version 1.4.0 (2026). doi:10.5281/zenodo.24204195.

[7] S. Douglas and N. Mghirbi. A perturbed cellular automaton with infinitely many ergodicity transitions. Preprint, version 1.0.0 (2026). doi:10.5281/zenodo.24122609.

[8] S. Douglas and N. Mghirbi. Exponential ergodicity and the space-time Gibbs states of one-dimensional probabilistic cellular automata. Preprint, version 1.2.0 (2026). doi:10.5281/zenodo.24122307.

[9] S. Douglas and N. Mghirbi. An uncomputable phase diagram at positive temperature for one classical lattice model. Preprint, version 1.2.0 (2026). doi:10.5281/zenodo.24144192.

[10] P. Gács. Reliable computation with cellular automata. *Journal of Computer and System Sciences* 32 (1986), 15–78. doi:10.1016/0022-0000(86)90002-4.

[11] P. Gács. Reliable cellular automata with self-organization. *Journal of Statistical Physics* 103 (2001), 45–267. Corrected and strengthened version arXiv:math/0003117v2.

[12] S. Goldstein, R. Kuik, J. L. Lebowitz and C. Maes. From PCA's to equilibrium systems and back. *Communications in Mathematical Physics* 125 (1989), 71–79.

[13] J. Kari. The nilpotency problem of one-dimensional cellular automata. *SIAM Journal on Computing* 21 (1992), 571–586. doi:10.1137/0221036.

[14] J. L. Lebowitz, C. Maes and E. R. Speer. Statistical mechanics of probabilistic cellular automata. *Journal of Statistical Physics* 59 (1990), 117–170. doi:10.1007/BF01015566.

[15] I. Marcovici and S. Taati. Positive-rate PCA and IPS with stationary Bernoulli measures are rapidly forgetful. Preprint (2026), arXiv:2605.16904.

[16] H. Marsan and M. Sablik. Strong stochastic stability of cellular automata. Preprint (2024), arXiv:2403.19932.

[17] H. Marsan and M. Sablik. Flexibility versus genericity of phase diagrams of perturbed continuous maps on the Cantor set. Preprint (2025), arXiv:2508.00461.

[18] H. Marsan, M. Sablik and I. Törmä. A perturbed cellular automaton with two phase transitions for the ergodicity. *Journal of Statistical Physics* 193 (2026), article 6. doi:10.1007/s10955-025-03564-0.

[19] J. Miękisz and S. Taati. Quasicrystalline Gibbs states in 4-dimensional lattice-gas models with finite-range interactions. Preprint (2026), arXiv:2512.24436v2.

[20] A. L. Toom. Algorithmical unsolvability of the ergodicity problem for binary cellular automata. *Markov Processes and Related Fields* 6 (2000), 569–577.

[21] A. L. Toom, N. B. Vasilyev, O. N. Stavskaya, L. G. Mityushin, G. L. Kurdyumov and S. A. Pirogov. Discrete local Markov systems. In R. L. Dobrushin, V. I. Kryukov and A. L. Toom (eds.), *Stochastic Cellular Systems: Ergodicity, Memory, Morphogenesis*, Manchester University Press, 1990, pp. 1–182.

[22] M. Gu, C. Weedbrook, A. Perales and M. A. Nielsen. More really is different. *Physica D* 238 (2009), 835–839. doi:10.1016/j.physd.2008.12.016.

[23] J.-R. Chazottes and M. Hochman. On the zero-temperature limit of Gibbs states. *Communications in Mathematical Physics* 297 (2010), 265–281. doi:10.1007/s00220-010-0997-8.

[24] J.-R. Chazottes and M. Shinoda. On the absence of zero-temperature limit of equilibrium states for finite-range interactions on the lattice $\mathbb{Z}^2$. Preprint (2020), arXiv:2010.08998v2.

[25] S. Barbieri, R. Bissacot, G. Dalle Vedove and P. Thieullen. Zero-temperature chaos in bidimensional models with finite-range potentials. *Advances in Mathematics* 457 (2024), 109906. doi:10.1016/j.aim.2024.109906.

[26] L. Gayral, M. Sablik and S. Taati. Characterisation of the set of ground states of uniformly chaotic finite-range lattice models. *Nonlinearity* 38 (2025), 075011. doi:10.1088/1361-6544/addbb8.

[27] S. Douglas and N. Mghirbi. Spontaneous magnetization and thermal decay are undecidable at positive temperature. Preprint, version 1.2.0 (2026). doi:10.5281/zenodo.23962706.

[28] H.-O. Georgii. *Gibbs Measures and Phase Transitions*, second extended edition. De Gruyter, 2011. doi:10.1515/9783110250329.

[29] F. Martinelli, E. Olivieri and R. H. Schonmann. For 2-D lattice spin systems weak mixing implies strong mixing. *Communications in Mathematical Physics* 165 (1994), 33–47. doi:10.1007/BF02099735.

[30] Y. Spinka. Finitary codings for spatial mixing Markov random fields. *Annals of Probability* 48 (2020), 1557–1591. doi:10.1214/19-AOP1405; arXiv:1803.10578.

[31] I. Syôzi. Statistics of Kagomé lattice. *Progress of Theoretical Physics* 6 (1951), 306–308. doi:10.1143/ptp/6.3.306.

[32] M. E. Fisher. Transformations of Ising models. *Physical Review* 113 (1959), 969–981. doi:10.1103/PhysRev.113.969.

[33] L. Čanová and J. Strečka. Reentrant phenomenon in the exactly solvable mixed spin-1/2 and spin-1 Ising–Heisenberg model on diamond-like decorated planar lattices. *Physica Status Solidi B* 247 (2010), 433–443. doi:10.1002/pssb.200945444; arXiv:0909.4191.

[34] H. Čenčariková, J. Strečka and M. L. Lyra. Reentrant phase transitions of a coupled spin-electron model on doubly decorated planar lattices with two or three consecutive critical points. *Journal of Magnetism and Magnetic Materials* 401 (2016), 1106–1122. doi:10.1016/j.jmmm.2015.11.018; arXiv:1706.09871.

[35] V. Powers and B. Reznick. A new bound for Pólya's theorem with applications to polynomials positive on polyhedra. *Journal of Pure and Applied Algebra* 164 (2001), 221–229. doi:10.1016/S0022-4049(00)00155-9.

[36] H. Duminil-Copin, A. Raoufi and V. Tassion. Sharp phase transition for the random-cluster and Potts models via decision trees. *Annals of Mathematics* 189 (2019), 75–99. doi:10.4007/annals.2019.189.1.2; arXiv:1705.03104v2.

[37] H. Duminil-Copin, M. Gagnebin, M. Harel, I. Manolescu and V. Tassion. Discontinuity of the phase transition for the planar random-cluster and Potts models with $q > 4$. *Annales Scientifiques de l'École Normale Supérieure* 54 (2021), 1363–1413. doi:10.24033/asens.2485; arXiv:1611.09877v2.

[38] V. Shapiro. Semi-analytic geometry with R-functions. *Acta Numerica* 16 (2007), 239–303. doi:10.1017/S096249290631001X.

[39] V. Yu. Protasov and A. S. Voynov. Matrix semigroups with constant spectral radius. *Linear Algebra and its Applications* 513 (2017), 376–408. doi:10.1016/j.laa.2016.10.013; arXiv:1407.6568.

[40] R. L. Dobrushin and S. B. Shlosman. Constructive criterion for the uniqueness of Gibbs field. In *Statistical Physics and Dynamical Systems*, Birkhäuser, Boston (1985), 347–370. doi:10.1007/978-1-4899-6653-7_20.

[41] D. Weitz. Combinatorial criteria for uniqueness of Gibbs measures. *Random Structures and Algorithms* 27 (2005), 445–475. doi:10.1002/rsa.20073.

[42] S. Shlosman. The critical temperature T_{cr}(Ising) is DS-computable. Preprint (2025). arXiv:2505.24750.

[43] J. Purcell, Z. Li and T. S. Cubitt. Chaitin phase transition. Preprint (2024). arXiv:2410.02600.

[44] A. V. Zarubin and F. A. Kassan-Ogly. The formation of magnetic reentrancy in the Ising model on a decorated square lattice. Preprint (2026). arXiv:2606.12572.

[45] V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields* 153 (2012), 511–542. doi:10.1007/s00440-011-0353-8.

[46] J. Stephenson. Ising model with antiferromagnetic next-nearest-neighbor coupling: spin correlations and disorder points. *Physical Review B* 1 (1970), 4405–4409. doi:10.1103/PhysRevB.1.4405.

[47] M. E. Fisher and B. Widom. Decay of correlations in linear systems. *Journal of Chemical Physics* 50 (1969), 3756–3772. doi:10.1063/1.1671624.

[48] I. G. Enting. Crystal growth models and Ising models: disorder points. *Journal of Physics C: Solid State Physics* 10 (1977), 1379–1388. doi:10.1088/0022-3719/10/9/011.

[49] I. Peschel and V. J. Emery. Calculation of spin correlations in two-dimensional Ising systems from one-dimensional kinetic models. *Zeitschrift für Physik B* 43 (1981), 241–249. doi:10.1007/BF01297524.

[50] P. Ruján. Order and disorder lines in systems with competing interactions: II. The IRF model. *Journal of Statistical Physics* 29 (1982), 247–262. doi:10.1007/BF01020785.

[51] J. Mairesse and I. Marcovici. Probabilistic cellular automata and random fields with i.i.d. directions. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 50 (2014), 455–475. doi:10.1214/12-AIHP530.

[52] J. Ouaknine and J. Worrell. Decision problems for linear recurrence sequences. In *Reachability Problems*, Lecture Notes in Computer Science 7550 (2012), 21–28. doi:10.1007/978-3-642-33512-9_3.

[53] P. Caputo and D. Parisi. Block factorization of the relative entropy via spatial mixing. *Communications in Mathematical Physics* 388 (2021), 793–818. doi:10.1007/s00220-021-04237-1.

[54] J. Casse and J.-F. Marckert. Markovianity of the invariant distribution of probabilistic cellular automata on the line. *Stochastic Processes and their Applications* 125 (2015), 3458–3483. doi:10.1016/j.spa.2015.05.001; arXiv:1401.5365v2.

[55] J. Casse and I. Marcovici. Probabilistic cellular automata with memory two: invariant laws and multidirectional reversibility. *Annales Henri Lebesgue* 3 (2020), 501–559. doi:10.5802/ahl.39; arXiv:1710.05490.

[56] A. Dudenhefner. Certified decision procedures for two-counter machines. In *FSCD 2022*, LIPIcs 228 (2022), 16:1–16:18. doi:10.4230/LIPIcs.FSCD.2022.16.

[57] P. Rebeschini and R. van Handel. Comparison theorems for Gibbs measures. *Journal of Statistical Physics* 157 (2014), 234–281. arXiv:1308.4117.

[58] P. Flajolet and R. Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009. doi:10.1017/CBO9780511801655.