

Polynomial critical-distance bounds for Gibbs preparation in the weak-field square-lattice Ising model

Yuxuan Zhang

Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL),
CH-1015 Lausanne, Switzerland

Department of Physics, Princeton University,
Princeton, New Jersey 08544, USA

26 October 2026 — Version 1

Abstract

We present a proof candidate for preparing the Gibbs state of the ferromagnetic transverse-field Ising model on a square torus below the classical critical inverse temperature. Write δ for the distance to that critical point. We obtain a sufficient transverse-field window $|g| \leq c\delta^p$ and a uniform circuit construction with gate count and classical bit cost bounded by $CN \max(1, \delta^{-q})[1 + \log(N/\varepsilon)]^a \text{poly}(b)$, for fixed finite exponents. The output is an actual mixed Gibbs state at global trace-distance error ε . The argument combines classical Ising mixing estimates, an exact doubled-space parent, quantitative spectral-flow stability, and explicit coherent access to a connected Pauli expansion. A relative bound for the global finite-volume remainder avoids an artificial volume crossover in the field exponent. The circuit analysis pays for growing spatial patches rather than treating them as free local oracles. Constants are deliberately conservative. This is a quantitative preparation result; no separation from general classical sampling algorithms is proved.

Status. Author-directed, AI-assisted proof candidate. The complete chain has received a sequential proof and source-interface audit, not independent review, external refereeing, or formal verification. The statement below is the claim to be reviewed; the problem catalog is not marked independently solved by publication.

1 Statement and contribution

On the $L \times L$ square torus, $L \geq 3$, let $N = L^2$ and

$$H(g) = -\frac{1}{4} \sum_{\langle i,j \rangle} Z_i Z_j - g \sum_i X_i, \quad \rho_{\beta,g} = \frac{e^{-\beta H(g)}}{\text{Tr } e^{-\beta H(g)}}. \quad (1)$$

Edges are unordered nearest-neighbor pairs. Fix $0 < \beta_{\min} < \beta_c = 2 \log(1 + \sqrt{2})$, take $\beta_{\min} \leq \beta < \beta_c$, and write $\delta = \beta_c - \beta$, $d_0 = \min\{1, \delta\}$. The classical coupling in conventional Ising units is $\beta/4$.

Theorem 1.1. *There are fixed positive constants c, C, p, q, a such that, whenever $|g| \leq c\delta^p$ and $0 < \varepsilon < 1/4$, a uniform quantum algorithm prepares $\tilde{\rho}$ satisfying*

$$\frac{1}{2} \|\tilde{\rho} - \rho_{\beta,g}\|_1 \leq \varepsilon. \quad (2)$$

Its elementary-gate count, classical bit-operation count, and ancillary space are each at most

$$CN \max\{1, \delta^{-q}\} [1 + \log(N/\varepsilon)]^a \text{poly}(b). \quad (3)$$

The input contains L, β, g, ε as explicit finite descriptions, with parameter bit length b . A fixed standard finite universal gate set with efficient synthesis and arbitrary gate connectivity is used. No Gibbs-state input, quantum random-access memory, dense local matrix oracle, or postselection is supplied.

One admissible parameter prescription is as follows. Let $C_* \geq 1$ be a fixed integer covering the classical polynomial finite-volume gap bounds in Lemma 3.1. Choose a fixed integer $a_0 > C_* + 14$ and set

$$q_0 = 4C_*, \quad A = 2a_0 + 29, \quad E = C_* + 8, \quad p = q_0 A + 2E + 1, \quad q = 28q_0 + 1. \quad (4)$$

The exponent a in (3) is fixed by reversible arithmetic and gate synthesis. These expressions are not optimized or useful numerical thresholds. The constant c also pays a fixed, very small convergence and implementation cutoff.

Rouzé and Stilck França [1] give weak-perturbation preparation for each fixed subcritical reference temperature. Their fixed-parameter implementation theorem does not by itself provide the polynomial critical-distance dependence in (3); constants involving block dimensions and boundary couplings must be paid when those blocks grow. Our contribution is the quantitative parent-to-circuit chain below. The construction uses established classical mixing [2, 3], stability [4, 5], and simulation methods [6, 7]. Auxiliary parents and local spectral transport are also used by Xun, Cao, and Watanabe [8]; their local-channel framework is not an elementary-gate complexity theorem for this near-critical regime.

The quantum task in (2) is more than sampling a basis string. It still does not imply a computational quantum advantage: no lower bound against general randomized classical algorithms is established. If the output is measured in the computational basis, the resulting distribution is close in total variation by contractivity, but no classical sampling separation follows.

2 The exact parent and local ground spaces

Double every physical site with an ancilla, and use the unnormalized maximally entangled vector $|I\rangle = \sum_x |x, x\rangle$. The thermofield vector

$$|\Omega_{\beta, g}\rangle = (\text{Tr } e^{-\beta H(g)})^{-1/2} (e^{-\beta H(g)/2} \otimes I) |I\rangle$$

purifies $\rho_{\beta, g}$. Operators below that contain H act only on the physical copy.

Let $P_i = (I - Z_i Z'_i)/2$ be the mismatch projection. For fixed neighboring physical spins τ , define the single-site conditional probabilities

$$p_i(s \mid \tau) = \frac{\exp[(\beta/4)s \sum_{j \sim i} \tau_j]}{2 \cosh[(\beta/4) \sum_{j \sim i} \tau_j]}, \quad s = \pm 1.$$

On the matched pair, let h_i project onto $\sqrt{p_i(- \mid \tau)}|00\rangle - \sqrt{p_i(+ \mid \tau)}|11\rangle$, and let it annihilate the mismatched pair. Then $P_i h_i = 0$ and

$$Q_0 = \sum_i (P_i + h_i)$$

has unique zero vector $\Omega_{\beta,0}$. Indeed, all mismatch flags commute with all terms, every interior mismatch costs at least one, and the matched sector is the symmetric ground-state transform of classical heat-bath dynamics with one update clock per site.

Set

$$K_g = e^{-\beta H(g)/2} e^{\beta H(0)/2}, \quad Q_g = \sum_i \sum_{D_i \in \{P_i, h_i\}} (K_g D_i K_g^{-1})^\dagger (K_g D_i K_g^{-1}). \quad (5)$$

Invertibility of K_g identifies the common kernel exactly as the span of $\Omega_{\beta,g}$. The two annihilators in (5) must remain separate: conjugating $P_i + h_i$ and squaring would introduce cross terms.

For a square Λ , the canonical restriction includes only anchors whose whole star lies in Λ . Write I_Λ for these centers and $S_\Lambda = \Lambda \setminus I_\Lambda$ for the shell. Both physical shell labels τ and ancilla shell labels η are free. Its ground projector is

$$\mathcal{P}_\Lambda = \sum_{\tau, \eta} |\tau, \eta\rangle \langle \tau, \eta|_{S_\Lambda} \otimes |\psi_\tau\rangle \langle \psi_\tau|_{I_\Lambda}, \quad |\psi_\tau\rangle = \sum_x \sqrt{\pi^\tau(x)} |x, x\rangle. \quad (6)$$

The nonzero gap is at least the smaller of one and the worst fixed-boundary heat-bath gap on I_Λ .

Lemma 2.1 (All-observable boundary reduction). *For a doubled inner set B separated from the shell, let $C = \partial^+ B$ be its physical exterior neighbor set. Its conditional reduced density is*

$$\rho_B^\tau = \sum_z \pi^\tau(C = z) |\phi_z\rangle \langle \phi_z|, \quad |\phi_z\rangle = \sum_{x_B} \sqrt{\pi(x_B | z)} |x_B, x_B\rangle. \quad (7)$$

The vectors ϕ_z do not depend on τ , and $\|\rho_B^\tau - \rho_B^{\tau'}\|_1 \leq \|\pi^\tau|_C - \pi^{\tau'}|_C\|_1$.

Proof. Tracing the doubled complement forces its bra and ket configurations to coincide. Group them according to $C = z$ and use the nearest-neighbor Markov property. Equation (7) is a common classical-to-quantum channel, so trace-norm contraction proves the inequality. Inner observables have no off-diagonal shell-label matrix elements in (6); hence the same bound controls compression of every inner quantum observable. No exponential factor in $|B|$ is introduced. $\square$

3 Quantitative classical inputs

Lemma 3.1. *There are universal $a_* > 0, C_0 \geq 1$, a fixed integer $C_* \geq 1$, and constants $c_\gamma, C_\xi > 0$ such that, uniformly for $0 \leq \beta < \beta_c$:*

- (i) *canonical radius- r square restrictions have gap at least $a_*(1+r)^{-C_*}$;*
- (ii) *the full L -torus reference gap is at least $a_*(1+L)^{-C_*}$;*
- (iii) *for an observable O_B buffered inside Λ ,*

$$\|\mathcal{P}_\Lambda O_B \mathcal{P}_\Lambda - \langle O_B \rangle_{\Omega_{\beta,0}} \mathcal{P}_\Lambda\| \leq C_0 |B| \|O_B\| e^{-\text{dist}(B, S_\Lambda)/(C_\xi)}, \quad \xi \leq C_\xi d_0^{-2}; \quad (8)$$

- (iv) *the reference gap also obeys the volume-uniform bound*

$$\text{gap}(Q_0) \geq \gamma = c_\gamma d_0^{q_0}, \quad q_0 = 4C_*, \quad (9)$$

with $\gamma \leq 1$ after decreasing c_γ .

Proof. We identify the classical input and its quantitative use. The FK parameter is $u = 1 - e^{-\beta/2}$, and $u_c - u$ is comparable to $\beta_c - \beta$ near criticality. Lubetzky–Sly [2, Theorem 3.1 and Section 4.2] give polynomial finite-volume gaps for arbitrary fixed and periodic boundaries at criticality. Their

upper mixing argument extends uniformly to $u \leq u_c$: in the boundary-disagreement exploration it uses the wired, bounded-aspect-ratio crossing *upper* bound strictly below one. This bound persists under decreasing u . The conditional FK exploration for mixed spin boundaries is retained, including its no-plus/minus-connection conditioning; no inverse probability of that conditioning is inserted. The critical lower-gap argument is not imported. This is an extension of the upper-bound proof, not a monotonicity assertion about Glauber gaps. Enlarging C_* and decreasing a_* covers the fixed and periodic estimates. Equation (6) transfers them to (i)–(ii).

Duminil-Copin–Garban–Pete [3, Lemma 2.13] give exponential wired crossing decay beyond a length

$$\ell \leq C d_0^{-1} \log(e/d_0) \leq C d_0^{-2}.$$

Their dyadic rectangle estimate extends to intermediate scales by enlargement and to squares by a fixed rectangle cover. To separate an interior point from a changing boundary, cover a square annulus by a fixed number of overlapping bounded-aspect rectangles, and use domain Markov and wired domination. Physical crossings, rather than identification wires on the boundary, are counted. The boundary-disagreement exploration thus gives a joint marginal bound $C|C|e^{-u/(C\ell)}$ at buffer u . Only rectangles of bounded aspect ratio are needed, including inward-shifted rectangles near a side and corner pieces. Lemma 2.1 gives (iii), retaining its support factor. The comparison with the global expectation follows by conditioning the torus law on the same shell.

For (iv), consider a rectangle of long side R and bounded aspect ratio. Split it into two overlapping blocks with overlap $w = \lceil K\ell \log(R+2) \rceil$ and average cut positions over the middle third. Strong mixing makes the failure to couple the complement of the other block at most $\eta_R = R^{-4}$, for fixed sufficiently large K . The two-block continuous chain, with rate one per block, has gap at least $1 - \eta_R$. One direct check uses bad/good noncoalesced states: the random-scan transition is bounded by

$$\begin{pmatrix} \eta_R & 1 - \eta_R \\ \eta_R/2 & (1 - \eta_R)/2 \end{pmatrix},$$

whose spectral radius is $(1 + \eta_R)/2$. Good means the complement of the next block agrees; updating that block coalesces. Passing from the rate-two random scan to continuous time gives the bound.

The averaged site coverage is at most $1 + Cw/R$. The block Dirichlet-form comparison therefore incurs inverse-gap loss at most $(1 + Cw/R)/(1 - R^{-4})$ per scale. Splitting the longer side contracts the maximal side by a fixed factor within two steps, preserving bounded aspect ratios. Stop at

$$R_0 = C\ell[1 + \log(\ell + 2)]^2 \leq C d_0^{-4}.$$

The accumulated logarithmic loss is bounded by a constant times $\sum_{\text{scales}} [\ell \log(R+2)/R + R^{-4}]$, which is bounded uniformly in the initial size. Leaf gaps from (i) yield $c d_0^{4C_*}$. The two constant-overlap reductions in [2, Section 4.2] remove the periodic directions, costing only two constants; their two-side crossing upper estimate is uniform below criticality as above. Small tori use (ii). This proves (9). $\square$

Both bounds $\text{gap}(Q_0) \geq \gamma$ and $\text{gap}(Q_0) \geq a_*(1 + L)^{-C_*}$ are retained. They have different roles in the perturbation argument.

4 A volume-uniform parent perturbation

Put $t = \beta/2$, $A_{\text{cl}} = -H(0)$, $V_X = \sum_i X_i$, and $\tau_s(O) = e^{sA_{\text{cl}}} O e^{-sA_{\text{cl}}}$. With $U'_I = g\tau_{-s}(V_X)U_I$ and $U_I(0) = I$, one has $K_g = \tau_t(U_I(t))$. Thus each $D_i \in \{P_i, h_i\}$ has expansion

$$K_g D_i K_g^{-1} = \sum_{m \geq 0} g^m \int_{t \geq u_1 \geq \dots \geq u_m \geq 0} \tau_t \text{ad}_{\tau_{-u_1}(V_X)} \cdots \text{ad}_{\tau_{-u_m}(V_X)}(\tau_{-t}(D_i)) d\mathbf{u}. \quad (10)$$

The latest time is the outermost commutator, as follows from the conjugation ODE.

All Ising edges commute. For physical support S , $\|\tau_s(O)\| \leq e^{2|s||S|}\|O\|$, and support expands by at most one neighbor shell. Initially $\tau_{-t}(D_i)$ has at most thirteen physical sites and norm at most e^{10t} . At a subsequent commutator at most $5(13 + 4r)$ star centers can meet current support, and each adds at most four sites. A star has norm at most e^{2t} . The time simplex supplies $t^m/m!$. The sum of norms at order m is consequently bounded by

$$e^{36t} z^m \frac{(13/4)_m}{m!}, \quad z = 40te^{10t}|g|. \quad (11)$$

Here $(a)_m$ is the rising factorial. The support radius is at most $3 + 2m$. At $m = 0$ the term is exactly D_i .

For $z \leq 1/4$, the series converges absolutely and its tail is bounded by $C2^{-M}$. It agrees with the finite-dimensional ODE solution, uniformly in volume. Grouping the separate Gram terms by total order k gives a Hermitian interaction with norm majorant $Cz^k(13/2)_k/k!$ and radius $3 + 2k$. Its difference from Q_0 starts at $k = 1$. For a fixed site at most $C(1 + k)^2$ anchors cover it, and each support has at most $C(1 + k)^2$ sites. Therefore, with a fixed $\mu > 0$,

$$\sup_x \sum_{X \ni x} |X| e^{\mu \text{diam } X} \|(Q_g - Q_0)(X)\| \leq C \sum_{k \geq 1} (1 + k)^4 2^{k/2} z^k \frac{(13/2)_k}{k!} \leq C'|g| =: J. \quad (12)$$

One may take $\mu = (\log 2)/8$. Taking one factor z out shows the last bound uniformly as $g \rightarrow 0$. Wraparound reduces support sizes and does not change the bound.

5 Quantitative stability and the global remainder

For stability only, interpolate $\mathcal{H}_s = Q_0 + s(Q_g - Q_0)$, $0 \leq s \leq 1$. This is not the physical preparation path. Intermediate Hamiltonians need not be frustration free. On a bootstrap interval with unique ground eigenvalue and gap at least $\gamma/2$, exact spectral transport gives the following estimate.

Lemma 5.1. *Take $a_0 > C_* + 14$, $M = a_0 + 14$, $A = 2M + 1$. If $J \leq c\gamma^{M+1}$, there is a unitary U_s transporting the ground projector to that of Q_0 such that*

$$U_s^\dagger \mathcal{H}_s U_s - Q_0 - E_s I = \sum_x Z_x, \quad Z_x = Z_x^\dagger, \quad Z_x \mathcal{P}_0 = 0, \quad (13)$$

and, with $\Pi_{x,r}$ the normalized partial-trace localization onto a radius- r square,

$$\|Z_x\| \leq M_0, \quad \|Z_x - \Pi_{x,r} Z_x\| \leq M_0(1 + r)^{-a_0}, \quad M_0 = CJ\gamma^{-A}. \quad (14)$$

The constants are independent of volume, δ and g .

Proof. The fixed-moment spectral-flow construction is that of [4, Section 4.2]. Its quantitative adaptation, retaining the small factor J and preventing an exponential in γ^{-1} , is proved in Appendix A. $\square$

Proposition 5.2 (All-volume relative bound). *Under Lemmas 3.1 and 5.1,*

$$\left| \left\langle \psi, \sum_x Z_x \psi \right\rangle \right| \leq b_{\text{rel}} \langle \psi, Q_0 \psi \rangle, \quad b_{\text{rel}} \leq CM_0 a_*^{-1} [1 + \sqrt{C_0} (1 + \xi)^E], \quad E = C_* + 8. \quad (15)$$

Proof. For $L \geq 128$, use nonwrapping square balls and $r_* = \lfloor L/8 \rfloor$. Let $\mathcal{P}_r$ be the canonical local ground projector at x , and for $r \geq 4$ define

$$A_r = \Pi_{x, \lfloor r/2 \rfloor} Z_x, \quad W_x(r) = (I - \mathcal{P}_r) A_r (I - \mathcal{P}_r).$$

Since $Z_x \Omega_{\beta,0} = 0$, $\|A_r \Omega_{\beta,0}\| \leq CM_0 (1+r)^{-a_0}$. Apply (8) to A_r^2 , not just A_r , to get

$$\|\mathcal{P}_r A_r^2 \mathcal{P}_r\| \leq CM_0^2 [(1+r)^{-2a_0} + C_0 (1+r)^2 e^{-r/(C\xi)}].$$

Using $\|A_r \mathcal{P}_r\|^2 = \|\mathcal{P}_r A_r^2 \mathcal{P}_r\|$ gives

$$\|Z_x - W_x(r)\| \leq CM_0 f(r), \quad f(r) = (1+r)^{-a_0} + \sqrt{C_0} (1+r) e^{-r/(C\xi)}. \quad (16)$$

Frustration freeness implies $\mathcal{P}_r \leq \mathcal{P}_{r-1}$. Thus the self-adjoint increment $B_{x,r} = W_x(r) - W_x(r-1)$ annihilates $\mathcal{P}_r$. For the first fixed radius set the preceding W to zero; its cost is absorbed into constants. The local gap yields

$$|\langle \psi, B_{x,r} \psi \rangle| \leq a_*^{-1} (1+r)^{C_*} \|B_{x,r}\| \langle \psi, Q_{0,x,r} \psi \rangle.$$

Summing anchors counts each reference interaction at most $C(1+r)^2$ times. Equation (16) therefore bounds $W = \sum_x W_x(r_*)$ relatively by

$$CM_0 a_*^{-1} \sum_{r \geq 1} (1+r)^{C_*+2} [f(r) + f(r-1)] \leq CM_0 a_*^{-1} [1 + \sqrt{C_0} (1 + \xi)^{C_*+4}].$$

The polynomial tail is summable; the exponential tail is bounded by its elementary integral.

The remaining operator $\mathcal{R} = \sum_x [Z_x - W_x(r_*)]$ is self-adjoint and *exactly* annihilates $\mathcal{P}_0$. Its norm is at most

$$CM_0 [L^{2-a_0} + \sqrt{C_0} L^3 e^{-L/(C\xi)}].$$

The finite-torus gap in Lemma 3.1(ii), rather than an absolute Weyl estimate, now gives

$$\begin{aligned} |\langle \psi, \mathcal{R} \psi \rangle| &\leq a_*^{-1} (1+L)^{C_*} \|\mathcal{R}\| \langle \psi, Q_0 \psi \rangle, \\ (1+L)^{C_*} \|\mathcal{R}\| &\leq CM_0 [1 + \sqrt{C_0} (1 + \xi)^{C_*+3}]. \end{aligned}$$

Here $\sup_{u \geq 0} u^v e^{-u/(C\xi)} = (C\xi v/e)^v$ for fixed $v > 0$, while $C_* + 2 - a_0 < 0$. No commutation of $\mathcal{R}$ with Q_0 or positivity of $\mathcal{R}$ is required.

For $3 \leq L < 128$, use global annihilation directly: $|\langle \psi, \sum_x Z_x \psi \rangle| \leq CM_0 a_*^{-1} L^2 (1+L)^{C_*} \langle \psi, Q_0 \psi \rangle$. The finitely many sizes cost a fixed constant. The larger exponent $E = C_* + 8$ covers all bounds, proving (15). $\square$

Choose J so that $J \leq c\gamma^{M+1}$ and $b_{\text{rel}} \leq 1/4$. The rotated Hamiltonian minus its ground energy is bounded below by $(3/4)Q_0$, so its gap is at least $3\gamma/4$. Starting at $s = 0$, continuity excludes a first point where the gap falls to $\gamma/2$, closing the bootstrap uniformly in all finite tori. Since $J \leq C|g|$, $\gamma = c_\gamma d_0^{q_0}$ and $\xi \leq C d_0^{-2}$, it suffices that

$$p > \max\{q_0(M+1), q_0 A + 2E\}.$$

The choice (4) satisfies this. The endpoint (5) identifies the exact Gibbs purification. The global relative remainder is the useful quantitative refinement: it removes an unnecessary size crossover and its additional field restriction. It does not change the underlying stability method.

6 Constructive finite-gate preparation

The gap alone does not establish (3). In particular, a dense matrix on a ball of polylogarithmic radius would have exponential dimension. We give coherent access to short symbolic Pauli paths, localize an approximate transport, and pay for every patch propagator.

6.1 Two physical paths

First vary $\beta_s = s\beta$ at $g = 0$, then vary $g_s = sg$ at fixed β . At $\beta = g = 0$, $P_i + h_i = I - |\Phi^+\rangle\langle\Phi^+|$ on the doubled site. Thus $O(N)$ gates prepare the initial Bell product. Lemma 3.1 gives a uniform gap on the first path. The endpoint gap theorem of Section 5, applied at every intermediate field, supplies the second. The paths join exactly. Their derivative jump is handled by separate evolutions. Neither path is the linear perturbation interpolation used to prove stability.

6.2 Pauli-path access and its normalization

The coefficients of $D_i = P_i, h_i$ in the Pauli basis are elementary functions of β on at most six qubits; they and two derivatives are bounded on $[0, \beta_c]$. They can be evaluated to k bits in $\text{poly}(b, k)$ time. For an Ising edge Pauli $B = Z_u Z_v$, conjugation of a Pauli P is unchanged if they commute and otherwise is

$$e^{aB} P e^{-aB} = \cosh(2a)P + \sinh(2a)BP. \quad (17)$$

The phase of BP is retained. Since all Ising edges commute, only edges touching the original support need be processed; the added Z factors cause no further recursive expansion.

At commutator step r of (10), reversibly construct and deduplicate the neighboring star-center list. It has at most $b_r = 5(13 + 4r)$ entries, padded with zeros. A left/right multiplication bit implements the commutator sign. All operations are on Pauli labels and lattice coordinates, with polynomial cost in $m, b, \log N$. One does not enumerate the exponentially many complete paths.

For a deliberately loose linear-combination-of-unitaries (LCU) normalization, each binary choice in (17) uses a uniform register and a scalar-amplitude flag with normalization $2e^{t/2}$. A commuting or missing edge contributes zero on the unused branch. Let C_D be a fixed normalization for D_i . Initial conjugation costs at most $C_D 2^{20} e^{10t}$; each star costs at most $2^4 e^{2t}$; a commutator costs another factor two; final conjugation costs at most $2^{52+16m} e^{2t(13+4m)}$. Including centers and simplex volume gives

$$\alpha_m \leq C_D 2^{72} e^{36t} z_{\text{LCU}}^m (13/4)_m / m!, \quad z_{\text{LCU}} = 40t e^{10t} 2^{20} |g|. \quad (18)$$

Shrink the fixed field prefactor so that $|g| \leq 2^{-48}$, hence $z_{\text{LCU}} \leq 1/8$ for $t \leq t_c = \log(1 + \sqrt{2})$. The sum of (18) and each fixed polynomial-in- m weighted sum is a constant independent of N, δ, ε . The constant is large, but not an exponential in the truncation order.

The order register through order M_{tr} can be prepared with $O(M_{\text{tr}})$ computed rotations. Uniform ranks among b_r possibilities use binary prefix counts and rotations with polynomial bit cost; no uncharged factor two per branching step is assumed. Coefficient flags give top-left amplitudes rather than a successful measurement event. Sorting m independent uniform times on $[0, t]$ realizes the ordered simplex average, with factor $t^m / m!$ already included in (18). On a grid of Q_{time} points, the error is bounded by $C(m+1)^2 \alpha_m / Q_{\text{time}}$ using the integrand derivative and Lipschitz continuity of sorted coordinates. Ties have a fixed convention. After summation, only $O(\log(1/\eta) + \log(M_{\text{tr}} + 1))$ time bits are needed for operator error η ; there is no loop over grid points.

On the field path, derivatives introduce polynomials $mg^m s^{m-1}$ and $m(m-1)g^m s^{m-2}$, evaluated at $s = 0$ without dividing by g_s . The first path has $g = 0$ and only constant-size terms. Product

block encodings with fresh work registers implement the two *separate* truncated Gram terms and their first two derivatives.

6.3 Parent truncation

Let $Q^{(M_{\text{tr}})}(s)$ be this finite-range Gram sum. It is positive but need not be frustration free. The parent and its first two path derivatives have total truncation error at most

$$CN(M_{\text{tr}} + 1)^2 2^{-M_{\text{tr}}}.$$

Choose $M_{\text{tr}} = C \lceil \log(CN/(\varepsilon\gamma)) \rceil$ with a sufficiently large fixed C . Weyl's inequality gives a uniform gap of at least half the exact lower bound, and isolated-eigenvector perturbation gives endpoint vector error $O(\varepsilon)$. Below, γ denotes a conservative fixed fraction of the lower gap *after* truncation.

Write $Q(s) = \sum_x q_x(s)$. Its anchored terms have declared radius $r_0 = O(M_{\text{tr}})$, and

$$\|q_x\|, \|q'_x\|, \|q''_x\| \leq C.$$

A patch restriction includes only declared supports wholly contained in that patch. Its block-encoding normalization is at most C times the patch volume. No patch gap is assumed, and no patch is diagonalized.

6.4 An explicit approximate transport

Define the real odd filter

$$F_\sigma(u) = \frac{1}{2} \text{sgn}(u) \text{erfc}(\sigma|u|/\sqrt{2}), \quad \widehat{F}_\sigma(\omega) = \frac{i(1 - e^{-\omega^2/(2\sigma^2)})}{\omega}, \quad (19)$$

using the Fourier convention $\widehat{F}(\omega) = \int F(u) e^{i\omega u} du$. The value at zero is zero. Integration by parts in the sine transform proves (19). Put

$$\mathcal{A}_\sigma(s) = \int_{\mathbb{R}} F_\sigma(u) e^{iuQ(s)} Q'(s) e^{-iuQ(s)} du, \quad U'_\sigma = i\mathcal{A}_\sigma U_\sigma.$$

For the normalized ground vector in parallel-transport gauge,

$$\|\Omega'(s) - i\mathcal{A}_\sigma(s)\Omega(s)\| \leq CN\gamma^{-1} e^{-\gamma^2/(2\sigma^2)}. \quad (20)$$

Indeed the excited-ground matrix element of $i\mathcal{A}_\sigma$ is $-(1 - e^{-\Delta^2/(2\sigma^2)})Q'_{m0}/\Delta$, compared with $-Q'_{m0}/\Delta$ for the exact derivative. The odd filter has zero diagonal part. This uses energy differences and does not require the truncated ground energy.

Let $\ell \geq 1$ be an integer of order $1 + \log(N/(\varepsilon\gamma))$, choose a dyadic σ between $\gamma/(8\sqrt{\ell})$ and $\gamma/(4\sqrt{\ell})$ by rational comparisons, and set $T = 32\ell/\gamma$. The frequency error in (20) and the tail

$$\int_{|u|>T} |F_\sigma(u)| du \leq \frac{C}{\sigma^2 T} e^{-\sigma^2 T^2/2}$$

give total vector error $O(\varepsilon)$. Increasing the fixed coefficient in ℓ assigns any required fraction of the error budget. Unitary variation of constants integrates these errors without an exponential in N .

6.5 Spatial localization

The finite-range real-time Lieb–Robinson iteration for $Q(s)$, at spatial weight $1/r_0$, has rate at most $\kappa_0 = Cr_0^4$. Both the number of terms covering a site and their support size cost $O(r_0^2)$. Replace the evolution of q'_x by that under $Q_{B_R(x)}$, the restricted parent in a radius- R square. Duhamel and the commutator bound give, for $|u| \leq T$,

$$\|\tau_u^Q(q'_x) - \tau_u^{Q_{B_R(x)}}(q'_x)\| \leq CNTr_0^4 \exp[\kappa_0 T - (R - 4r_0)/r_0]. \quad (21)$$

The margin accounts for initial and omitted support radii. The crude count N for omitted anchors is sufficient; it will enter a logarithm, not a gate count.

Define

$$d_x(s) = \int_{-T}^T F_\sigma(u) \tau_u^{Q_{B_R(x)}(s)}(q'_x(s)) du, \quad \mathcal{A}_R = \sum_x d_x.$$

Choosing

$$R = \left\lceil 4r_0 + r_0 \left[\kappa_0 T + \log \frac{CN^2 Tr_0^4}{\varepsilon \sigma} \right] \right\rceil \leq C\gamma^{-1} \ell^6$$

makes the global flow error $O(\varepsilon)$. The N^2 pays both anchor sums. If the ball covers the torus the restriction is exact. Let $V_R \leq CR^2 \leq C\gamma^{-2} \ell^{12}$. Then

$$\|d_x\| \leq CT, \quad \|d'_x\| \leq C(T + V_R T^2). \quad (22)$$

The second bound follows by differentiating the unitary conjugation; no exponential derivative factor occurs.

A block encoding uses

$$d_x = \frac{1}{2} \int_0^T \operatorname{erfc}(\sigma u / \sqrt{2}) [\tau_u^{Q_{B_R}}(q'_x) - \tau_{-u}^{Q_{B_R}}(q'_x)] du.$$

Uniform time and sign registers, a reversibly evaluated erfc flag, and coherent restricted-parent simulations give normalization CT . The integrand derivative is bounded by $C(\sigma + V_R)$, so a grid size $CT^2(\sigma + V_R)/\eta$ suffices for error η . Again only its logarithm affects register size. Binary controlled times have total simulated duration $O(T)$. Coefficients are computed coherently in both the outer parameter s and time u ; measuring them is unnecessary.

6.6 Outer spatial decomposition and charged patch circuits

A global block encoding nested inside a second global simulation would introduce an avoidable N^2 cost. Instead apply the overlapping-region unitary decomposition of Haah–Hastings–Kothari–Low [6, Lemma 6 and time-dependent extension] to the localized generator $\mathcal{A}_R$. Its interaction rate, at weight $1/R$, is at most $CR^4 T$. An overlap

$$W = \left\lceil CR \left[R^4 T + \log \frac{CN^2(1 + R + T)}{\varepsilon} \right] \right\rceil \leq C\gamma^{-6} \ell^{31} \quad (23)$$

makes total decomposition error $O(\varepsilon)$. First split into overlapping strips, then split those in the other coordinate. This produces at most CN patch propagators on at most $V_W \leq CW^2 \leq C\gamma^{-12} \ell^{62}$ sites. Inverses are implemented by reversing the corresponding circuit. Time-dependent restrictions use the same Duhamel commutator estimate with suprema of local norms. The underlying decomposition is used, not a constant-range theorem that grants growing patches free implementation.

For a patch P , its generator is the literal restriction $\sum_{x: B_R(x) \subset P} d_x(s)$, with normalization at most $CV_W T$. To apply a selected Pauli, scan its $O(V_W)$ data qubits, compare their coordinates to the short symbolic support list, and apply the controlled Pauli letter. Coherent center labels do not add a second scan over all centers. Thus, for some fixed arithmetic exponent a_{bit} and $B = b + \lceil \ell \rceil$,

$$G_{\text{parent}} \leq CV_W B^{a_{\text{bit}}}.$$

The time-independent Taylor/Dyson simulation of an inner restricted parent costs at most $CV_R T B^2$ calls, using a fixed polynomial logarithmic overhead. It is uniform in coherent parameters and binary controlled times. Therefore a d_x oracle costs at most $CV_W V_R T B^{a_{\text{bit}}+2}$.

Time-dependent Dyson simulation [7, Theorem 3 and Corollary 4] of an outer patch for path time one needs at most $CV_W T B^2$ such calls. Equation (22) bounds the derivative required for discretization. The grid may be large but its number of bits is logarithmic; coefficients are evaluated on the quantum time register. Consequently

$$\begin{aligned} G_{\text{patch}} &\leq CV_W^2 V_R T^2 B^{a_{\text{bit}}+4}, \\ G_{\text{total}} &\leq CN V_W^2 V_R T^2 B^{a_{\text{bit}}+4} \leq CN \gamma^{-28} \ell^{138} B^{a_{\text{bit}}+4}. \end{aligned} \quad (24)$$

The use of Taylor/Dyson constructions also pays for preparing their elementary coefficients; no uncharged phase-optimization algorithm is assumed. Classical circuit generation and workspace satisfy the same type of bound. If a prescribed patch is larger than the torus, the whole system already has size at most that patch bound.

6.7 Precision and the final mixed state

Use fixed fractions of ε for parent truncation, approximate transport, time cutoff, spatial restriction, outer decomposition, simulation, and finite arithmetic. Every per-call numerical error can be chosen below ε divided by a fixed polynomial in N, γ^{-1}, ℓ, b . All precision logarithms are $O(b + \ell)$. Special-function arguments are bounded by constants or $O(\sqrt{\ell})$; series evaluation, square roots, reversible arithmetic and finite-gate synthesis have polynomial bit cost. If a top-left simulation error η^2 is needed for clean-ancilla isometry error $O(\eta)$, precision bits only double. Work registers are fresh or uncomputed; no accepted measurement branch is treated as free.

The input β, g defines the target exactly. Rounding coefficients is charged as an implementation error. The constant β_c is computable; evaluate it until δ is bracketed within a factor of two, giving a conservative rational lower gap with cost polynomial in $b + \log(1/\delta)$. The promise $\beta < \beta_c$ ensures termination.

Discard the second copy and work registers. Trace-distance contractivity gives (2) from the vector-norm budget. This is an actual quantum mixed state, including transverse-field coherence. Finally use $\gamma = c_\gamma d_0^{q_0}$ and $[1 + \log(1/d_0)]^A \leq C_A d_0^{-1}$ for each fixed A . Equation (24) becomes (3) with $q = 28q_0 + 1$ and $a = a_{\text{bit}} + 142$, up to a polynomial in b . Replace the field prefactor by $c/\max(1, \beta_c - \beta_{\min})^p$ to convert the d_0^p condition to δ^p also when $\delta > 1$. This completes the candidate proof of Theorem 1.1.

7 Scope, verification and authorship

This is a sufficient window near the *classical* critical point on its high-temperature side. It is not a claim at criticality, throughout the transverse-field phase diagram, or for a physically useful value of c . General connectivity is permitted. The exponents are fixed but large and unoptimized. A

gate-efficient algorithm here does not prove that a natural thermal bath mixes quickly, and does not establish a lower bound for classical basis sampling.

The manuscript combines the G02 parent, stability and compiler records from rounds 83, 87, 88 and 90 with the improved relative remainder of round 99. The source package retains those original records and hashes. This version corrects an old bibliographic attribution: arXiv:2610.27939 is by Xun, Cao and Watanabe; the Sengoku–Cao–Watanabe item is a distinct manuscript cited there. The audit checked canonical shell labels, all-observable compression, classical rate conventions, support-weighted perturbation norms, retention of the perturbative factor in spectral flow, all torus sizes, coherent coefficient access and the two nested simulation costs. These are analytic checks, not an execution of a large-system preparation circuit. Independent proof and novelty review remains outstanding.

Human and AI contributions. The author set the research directions, designed the harness coordinating model tasks, tools, memory and checking, supplied physical and mathematical intuitions and proof skeletons, and guided development and revision. AI systems assisted detailed derivations, primary-source comparison, computational checks, and manuscript drafting and revision. The assembling agent’s audit is not independent review.

A Polynomial constants in exact spectral transport

This appendix proves Lemma 5.1, making explicit where inverse-gap powers arise in the standard spectral-flow method [4]. We use dimension $d = 2$ but retain d in intermediate counts.

Choose a real even Schwartz function w with integral one and smooth Fourier transform supported in $(-1/8, 1/8)$. Set $w_\gamma(t) = \gamma w(\gamma t)$. An odd integrated kernel F_γ , with the spectral-flow sign convention, has inverse-frequency multiplier outside that support and moments

$$\int |t|^k |w_\gamma(t)| dt \leq C_k \gamma^{-k}, \quad \int |t|^k |F_\gamma(t)| dt \leq C_k \gamma^{-k-1}. \quad (25)$$

All moment orders are fixed, independently of δ and volume. Positivity of w is not required.

Write $V = Q_g - Q_0$ and $\mathcal{H}_s = Q_0 + sV$. On the bootstrap interval define

$$\mathcal{T}_s(O) = \int w_\gamma(t) e^{it\mathcal{H}_s} O e^{-it\mathcal{H}_s} dt, \quad D_s = \int F_\gamma(t) e^{it\mathcal{H}_s} V e^{-it\mathcal{H}_s} dt.$$

The first map removes all ground/excited matrix elements, and the second generates exact transport of the isolated ground projection. Let $\alpha_s(O) = U_s^\dagger O U_s$ denote transport back to the reference projection.

The exponential interaction bound (12) gives a real-time Lieb–Robinson bound with fixed velocity and decay constants for $J \leq 1$. For a local observable at x , split filter times at a fixed multiple of a radius r . Short times give an exponentially small spatial tail; long times use (25). With

$$b_0 = a_0 + 2d + 4, \quad M = b_0 + d + 4 = a_0 + 3d + 8,$$

one obtains

$$\|(I - \Pi_{x,r})\mathcal{T}_s(O)\| \leq C \|O\| \gamma^{-M} (1+r)^{-M}, \quad (26)$$

$$\|(I - \Pi_{x,r})D_{s,x}\| \leq C J \gamma^{-M-1} (1+r)^{-M}. \quad (27)$$

For exponentially decaying anchored V_x , decompose into support shells before applying these inequalities; fixed polynomial support factors are summable.

Decompose each $D_{s,x}$ into consecutive-radius shells. There are at most $C(1+r)^d$ possible anchors of a radius- r support containing two specified sites. For $F_{b_0}(u) = (1+u)^{-b_0}$, the interaction norm is therefore bounded by

$$\|D_s\|_{F_{b_0}} \leq CJ\gamma^{-M-1} \sum_{r \geq 0} (1+r)^{b_0+d-M} \leq CJ\gamma^{-M-1} =: \kappa. \quad (28)$$

The summand exponent is -4 . Since $b_0 > d$, summability and convolution constants are volume uniform, also on finite tori. Enforce $J \leq c\gamma^{M+1}$ *before* applying the parameter-time locality bound, so $\kappa \leq 1$. Integrating the commutator with D_s , and then localizing by normalized partial trace, gives for fixed local support X

$$\begin{aligned} \|(\alpha_s - I)O\| &\leq C\kappa|X|\|O\|, \\ \|(I - \Pi_{X(r)})(\alpha_s - I)O\| &\leq C\kappa|X|\|O\|(1+r)^{d-b_0}. \end{aligned} \quad (29)$$

The parameter-time exponential is $e^{C\kappa} - 1 \leq C'\kappa$, retaining the factor J . It is not $e^{\text{poly}(\gamma^{-1})}$.

The difference $\mathcal{T}_s - \mathcal{T}_0$ must likewise retain J . Differentiating real-time evolution gives a Duhamel integral of commutators with V . Inside its real-time cone there are $O((1+|u|)^d)$ anchors; outside, the summed exponential tail has the same order bound. For fixed-support O ,

$$\|\partial_s \tau_t^{\mathcal{H}_s}(O)\| \leq CJ\|O\| |t|(1+|t|)^d.$$

For $|t| \leq cr$, split anchors at distance $r/4$, truncate their exponential support tails at $r/8$, and truncate the evolved observable at $r/2$. The near terms remain local; their final evolution and the far terms cost CJ times a fixed polynomial in r times $e^{v'|t|-\mu'r}$. Choose c small enough. For $|t| > cr$, the displayed norm bound and the M th filter moment give

$$\|(I - \Pi_{x,r})(\mathcal{T}_s - \mathcal{T}_0)(O)\| \leq CJ\|O\| \gamma^{-M} (1+r)^{-(M-d-1)}. \quad (30)$$

The same coarse prefactor controls the norm at $r = 0$. Subtracting two separate bounds (26) would lose J and would be insufficient.

Write $Q_0 = \sum_x h_x^{(0)}$, $V = \sum_x V_x$ with the preceding anchored shells, and set

$$Y_x = \alpha_s \mathcal{T}_s(h_x^{(0)} + sV_x) - \mathcal{T}_0(h_x^{(0)}).$$

Both terms commute with the reference ground projection. Since that projection has rank one, $c_x = \langle \Omega_{\beta,0}, Y_x \Omega_{\beta,0} \rangle$ and $Z_x = Y_x - c_x I$ satisfy $Z_x \mathcal{P}_0 = 0$ exactly. Summing uses $\mathcal{T}_s(\mathcal{H}_s) = \mathcal{H}_s$ and $\sum_x c_x = E_s$, proving (13); one does not assume $\mathcal{T}_0(h_x^{(0)}) = h_x^{(0)}$ term by term.

Expand

$$Y_x = (\alpha_s - I)\mathcal{T}_s(h_x^{(0)}) + (\mathcal{T}_s - \mathcal{T}_0)(h_x^{(0)}) + s\alpha_s \mathcal{T}_s(V_x).$$

For the first term, decompose the filtered local term into radius- R shells of norm at most $C\gamma^{-M}(1+R)^{-M}$. Each support has $O((1+R)^d)$ sites. For $R \leq r/2$ use (29); for larger R use its norm bound. This yields

$$C\kappa\gamma^{-M} [(1+r)^{d-b_0} + (1+r)^{d-M+1}] \leq CJ\gamma^{-(2M+1)}(1+r)^{-a_0}.$$

Equation (30) controls the second term with a stronger radius exponent. Filtering and transporting the exponentially decaying shells of V_x gives the same bound for the third. Norm bounds at radius zero obey the same larger prefactor. Scalar subtraction at most doubles the norm and leaves localization tails unchanged. Thus (14) holds with $A = 2M + 1$, completing Lemma 5.1.

References

- [1] C. Rouzé and D. Stilck França, *Rapid mixing of Gibbs samplers via quantum Dobrushin–Shlosman conditions*, [arXiv:2610.03578v1](#) (2026), Sections 3.3 and 6.
- [2] E. Lubetzky and A. Sly, *Critical Ising on the square lattice mixes in polynomial time*, [arXiv:1001.1613v4](#) (2010), Theorem 3.1, Section 3.2, and Section 4.2.
- [3] H. Duminil-Copin, C. Garban, and G. Pete, *The near-critical planar FK-Ising model*, [arXiv:1111.0144v4](#), Theorem 1.2 and Lemma 2.13.
- [4] B. Nachtergaele, R. Sims, and A. Young, *Quasi-locality bounds for quantum lattice systems. Part II. Perturbations of frustration-free spin models with gapped ground states*, [arXiv:2010.15337v3](#), especially Sections 2, 4 and 5.
- [5] S. Michalakis and J. P. Zwolak, *Stability of frustration-free Hamiltonians*, [arXiv:1109.1588](#) (2011).
- [6] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, *Quantum algorithm for simulating real time evolution of lattice Hamiltonians*, [arXiv:1801.03922v4](#) (2020), Lemma 6 and the time-dependent extension.
- [7] G. H. Low and N. Wiebe, *Hamiltonian simulation in the interaction picture*, [arXiv:1805.00675v2](#) (2019), Theorem 3 and Corollary 4.
- [8] G. Xun, W. Cao, and H. Watanabe, *Local-channel stability of classical $\mathbb{Z}_2$ lattice gauge models*, [arXiv:2610.27939v1](#) (2026), Section II.4 and the auxiliary-parent construction in the supplement.