

Single-exponential temperature cost for one-dimensional Gibbs preparation

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

Version 1 — 22 October 2026

Abstract

We give a uniform implementation of quantum-belief-propagation Gibbs preparation on a one-dimensional chain. For fixed local dimension and interaction range, put $x = 1 + \beta J_{\text{int}}$ and $Q = 1 + \log(N/\varepsilon)$. The algorithm prepares the Gibbs state to trace distance ε , using at most $\exp(Cx^3)N(1+b)^cQ^{Cx+c}$ elementary gates and classical bit operations, including the expanded circuit description. Here b records local and scalar input lengths, and C, c depend only on fixed local parameters. Arbitrary onsite energy magnitudes enter through their binary description lengths. The logarithmic exponent is $O(x)$, rather than uniform in temperature. The construction completes the implementation route proposed by Rouzé and Stilck França: recursive block purification and collective reset encoding are combined with coherent quantile arithmetic, a logarithmically normalized correction oracle, and bounded finite-bit sampling of local updates. All child preparations, inverse calls, arithmetic precision and classical assembly are charged. The result concerns state preparation; it does not approximate the entire global semigroup on arbitrary inputs.

1 Question, input model and main result

Rapid mixing does not by itself give an efficient quantum circuit. In the one-dimensional smoothed heat-bath construction of Rouzé and Stilck França [1], the reset block has length exponential in inverse temperature. A dense representation of its Gibbs state, or a separate encoding of every reset jump, therefore produces a doubly exponential temperature prefactor. Their Remark 3.12 proposes recursive quantum belief propagation (QBP) and collective jump encoding to remove this cost. The remaining issue is to implement the proposed primitives and compose their errors without an uncharged purification or simulation oracle.

We supply that implementation analysis. The recursion is inherited; the contribution is a bound that includes its success amplification, the coherent correction, ordinary binary input, and classical circuit construction. The distinction matters because a polynomial-cost local primitive need not remain polynomial after it is invoked recursively at increasingly fine accuracy.

Fix a local Hilbert-space dimension and a finite interaction range. Let

$$H_N = \sum_X h_X, \quad \rho_\beta(H_N) = \frac{e^{-\beta H_N}}{\text{Tr } e^{-\beta H_N}}$$

be a chain Hamiltonian. The interaction strength J_{int} bounds the sum of intersite interaction norms incident on a site; onsite terms are excluded. Changing this convention by a fixed factor

changes only the constants below. Grouping a fixed number of neighboring sites reduces the problem to nearest-neighbor form.

Local Hermitian matrices, β , and supplied bounds have ordinary binary descriptions. They may be exact rationals or supplied to the absolute precision specified in Section 7. Let b bound their per-neighborhood raw description length and scalar bit lengths, including integer magnitude bits. Repeated supports are aggregated with their full input-reading cost charged to b . Site-address bits are charged separately. A succinct symbol for an arbitrarily large real number is not an exact-real oracle.

Theorem 1.1. *There are constants $C, c > 0$, depending only on the fixed local dimension, range and gate/arithmetic conventions, with the following property. For $\beta \geq 0$, $0 < \varepsilon \leq 1/8$, and the input above, a uniform randomized quantum algorithm outputs a density operator $\hat{\rho}$ satisfying*

$$\frac{1}{2} \|\hat{\rho} - \rho_\beta(H_N)\|_1 \leq \varepsilon.$$

Writing $x = 1 + \beta J_{\text{int}}$ and $Q = 1 + \log(N/\varepsilon)$, both its elementary-gate count and classical bit-operation count are bounded by

$$\exp(Cx^3) N(1+b)^c Q^{Cx+c}. \quad (1.1)$$

The classical cost includes input processing and expansion of the complete gate list. Every random seed has bounded work. The output density operator includes the algorithm's discarded classical randomness.

For larger requested error, run the theorem with $\min(\varepsilon, 1/8)$. In the bounded-local-input, requested-precision model of [1, Theorem 3.11 and Appendix B.1], b is polynomial in x, Q , so (1.1) has the requested form $\exp(\text{poly}(x))N \text{polylog}(N/\varepsilon)$ at each fixed temperature. For unrestricted input lengths the factor $(1+b)^c$ remains explicit. Neither a temperature-independent power of Q nor the stronger proposed prefactor $\exp(O(x^2))$ is proved.

Earlier recursive Gibbs preparation [6] establishes the usefulness of joining thermal regions, but its stated temperature-dependent power of N does not yield (1.1). The matrix-product-operator construction of [7, Theorem 3] has a quasi-linear size bound with an overhead exceeding a fixed polylogarithm at inverse-polynomial error. Kim's block-Petz implementation [2] is another nearly linear preparation method at fixed parameters; we use its ordered-Dyson circuit construction, while tracking the different QBP temperature and recursion costs explicitly.

2 Imported analytic and simulation interfaces

We distinguish the imported theorems from the implementation proved here. The analytic input from [1] consists of QBP transport, Gibbs-preserving dressed resets, one-dimensional locality estimates, and the discrete Dobrushin–Shlosman contraction theorem. The precise uses are Section 3.1 and Proposition 3.4 for transport and correction, Theorem 4.4 for discrete contraction, Propositions 5.7–5.8 for locality and influence, and Proposition 5.11 with Eq. (5.31) for block parameters.

For an interval B , let

$$\mathcal{R}_B(X) = \rho_\beta(H_B) \otimes \text{Tr}_B X.$$

QBP dressing gives an invertible O so that the completely positive part of the local generator is

$$\mathcal{T}_B = \text{Ad}_O \mathcal{R}_B \text{Ad}_{O^{-1}}.$$

If V is a Stinespring stack for this map and $S = V^\dagger V$, the exact source generator is

$$\mathcal{K}_B(X) = \text{Tr}_{\text{env}}(V X V^\dagger) - \frac{1}{2} \{S, X\} - i[G_H(S), X], \quad G_H(S) = \frac{i}{2} \tanh\left(\frac{\beta}{4} \text{ad}_H\right)(S). \quad (2.1)$$

It preserves the target Gibbs state. The corresponding update is $\Phi_B = e^{T\mathcal{K}_B}$. A localized construction replaces the ambient Hamiltonian by H_R on a neighborhood $R \supset B$, while retaining the Gibbs reset on B . This last distinction keeps the purification recursion independent of the eventual localization radius.

The algorithmic inputs are the interaction-picture Hamiltonian simulator of Low and Wiebe [3, Theorem 3, Corollary 4 and Lemma 6], the ordered-Dyson construction described in [2, Appendix C.2], Grover's $\pi/3$ fixed-point amplification [4], and the stacked-jump Lindblad simulator of Chen, Gao, Wang and Zhou [5, Definition 1.1 and Theorem 1.2]. We construct their actual oracles and account for their classical coefficient preparation below.

A block encoding with normalization α means a unitary whose selected ancilla block, multiplied by α , approximates the intended operator in operator norm. Rectangular stacks use specified input and output projections. These projections are components of a circuit construction, not physical postselection granted at no cost.

3 A coherent QBP coefficient from local matrices

Write a regional Hamiltonian as $H = A + D$, where $A = \sum_i a_i$ is onsite and D contains the interactions. On L sites an explicit local-term encoding has normalization $\alpha_D = O(LJ_{\text{int}})$. Let $\alpha_A \geq \sum_i \|a_i\|$, $\alpha = \alpha_A + \alpha_D$, and let a constant-size boundary interaction W have a local-term encoding with normalization $J \geq \|W\|$.

Lemma 3.1. *For $\theta > 0$ and $0 < \eta < 1/4$, the QBP filter*

$$\mathcal{F}_{\theta,H}(W) = \int_{\mathbb{R}} p_{\theta}(t) e^{itH} W e^{-itH} dt, \quad p_{\theta}(t) = \frac{4}{\pi\theta} \log \coth \frac{\pi|t|}{\theta} \quad (3.1)$$

has a block encoding of normalization J and error $J\eta$. Both its gate and classical construction costs are bounded by a fixed-degree polynomial in

$$L, b, 1 + \theta\alpha_D, \log(1/\eta).$$

The construction is coherent in an additional binary parameter s in $H(s) = H_0 + sW$, with its bit length charged.

Proof. For $t > 0$, expanding the logarithm expresses p_{θ} as a mixture of symmetric Laplace laws with rates $2\pi(2j+1)/\theta$ and weights $8/[\pi^2(2j+1)^2]$. The weights sum to one. Consequently the absolute-time survival function is

$$S_{\theta}(t) = \frac{8}{\pi^2} \left[\text{Li}_2(q) - \frac{1}{4} \text{Li}_2(q^2) \right], \quad q = e^{-2\pi t/\theta}. \quad (3.2)$$

In particular $S_{\theta}(t) \leq e^{-2\pi t/\theta}$. Formula (3.2) is used as a scalar arithmetic expression, without enumerating the mixture. The dilogarithm on $[0, 1/2]$ has a geometrically convergent power series; on $[1/2, 1]$, use

$$\text{Li}_2(y) = \frac{\pi^2}{6} - \log y \log(1-y) - \text{Li}_2(1-y).$$

Endpoint limits and guarded arithmetic give P -bit values in $\text{poly}(P)$ bit operations.

Choose $K = O(\log(1/\eta))$ with $K \geq \log(32/\eta)$, and $T = \theta K$. Cap $|t|$ at T , retaining its sign. The resulting normalized filter error is at most $\eta/16$. Let

$$q(u) = \min\{(1 - S_{\theta})^{-1}(u), T\}.$$

For $M = 2^n$ midpoint indices $u_j = (j + 1/2)/M$, monotonicity gives

$$\int_0^1 |q(u) - q(u_{j(u)})| du \leq T/M. \quad (3.3)$$

Indeed each cell contributes at most its width times the oscillation of q , and those oscillations telescope. Since $t \mapsto e^{itH}(W/J)e^{-itH}$ is 2α -Lipschitz, $M \geq 64(1 + \alpha)T/\eta$ bounds the grid error by $\eta/32$. Only $\log M$ bits and uniform Hadamards are needed.

Compute each capped quantile reversibly by bisection to time error $\eta/[64(1 + \alpha)]$. On $[0, T]$,

$$(1 - S_\theta)'(t) = 2p_\theta(t) \geq \frac{16}{\pi\theta} e^{-2\pi K}.$$

Thus an ambiguous CDF comparison can be resolved within the required time error using $O(K + |\log \theta| + \log(1/\eta) + \log(1 + \alpha))$ guard bits. The cap is treated by the same comparison rule. There are polynomially many bit operations; the numerical size of M never appears as a loop length.

It remains to implement the coherent time evolution. The interaction-picture theorem [3] has cost polynomial in $1 + \alpha_D T$ and precision bits when supplied with encodings of $e^{itA} D e^{-itA}/\alpha_D$ and onsite evolution. Fixed-dimensional onsite matrices can be diagonalized to the required accuracy in polynomial bit cost. Near-degenerate eigenvalues may be grouped at the requested accuracy, so inverse spectral gaps need not be paid. Their eigenvalue phases, including reduction modulo 2π , are evaluated by ordinary binary arithmetic; their cost depends on the integer bit lengths, not the onsite energy values. Explicit local-term sums encode D/α_D with $O(L)$ fixed-size controlled operations. Conjugation by the charged onsite circuits supplies the interaction-picture oracle.

For a coherent requested time $t = xT$, $x \in [-1, 1]$, simulate xH for the common duration T . A bounded scalar rotation encodes x , and binary multiplication supplies xs in the interaction-picture oracle. Uniform norm and derivative bounds allow common segment counts for every x . Negative times and the interpolation parameter require no node table. If $\alpha_D = 0$, onsite evolution suffices.

Let the resulting contraction block be P_t , with $\|P_t - e^{itH}\| \leq \xi$. Fresh product ancillas encode $P_t(W/J)P_t^\dagger$, with error at most 2ξ beyond the W -encoding error. Prepare uniform index/sign registers, compute the time, apply this controlled product, uncompute the time, and unprepare the uniform registers. The projected block is an average with positive weights summing to one. Normalization is therefore J , independent of M and Hilbert-space dimension.

Taking $\xi \leq \eta/64$, normalized boundary error at most $\eta/32$, and allocating the remaining slack to arithmetic and gate synthesis proves the claim. Regional Hamiltonian rounding by Δ_H changes the normalized filter by at most $C\theta \|\Delta_H\|$, by Duhamel and $\mathbb{E}|t| \leq \theta$. Thus required fractional input bits are polynomial in the listed quantities. Temperature rounding costs additional $\log(1 + \alpha)$ bits, already charged through $b + \log L$. $\square$

4 Dressing and recursively amplified purification

4.1 Ordered products and the inverse

For a merge of two adjacent intervals, set $H(s) = H_{\text{left}} + H_{\text{right}} + sW$. The exact transport identity [1, Section 3.1] is

$$A(s) = -\frac{\beta}{2} \mathcal{F}_{2\beta, H(s)}(W), \quad O'(s) = A(s)O(s), \quad O(0) = I, \quad O(1)e^{-\beta H(0)}O(1)^\dagger = e^{-\beta H(1)}. \quad (4.1)$$

The parameter 2β in this coefficient is required for full-temperature purification. With $\gamma = \beta J/2$, positivity of the filter and Duhamel give

$$\|A(s)\| \leq \gamma, \quad \|A(s) - A(s')\| \leq 8\gamma^2|s - s'|, \quad \|O(1)\|, \|O(1)^{-1}\| \leq e^\gamma. \quad (4.2)$$

Fix an integer $1 \leq \kappa \leq Cx$ dominating all merge coefficients and take $D = O(\kappa^2)$ as a common Lipschitz bound.

Lemma 4.1. *The operators $O(1)$ and $O(1)^{-1}$ have block encodings with normalization at most e^κ , absolute error η , and fixed-degree polynomial gate and classical cost in $L, b, x, \log(1/\eta)$.*

Proof. Put $B(s) = A(s)/\kappa$. The ordered-Dyson expansion can be written as

$$O(1) = \sum_{j \geq 0} \frac{\kappa^j}{j!} \mathbb{E}[B(t_{(j)}) \cdots B(t_{(1)})], \quad (4.3)$$

where the t_i are independent uniform variables and parentheses denote sorted order. This is the cube form of the ordered simplex integral. Choose $K+1 \geq 2e\kappa$ and $K+1 \geq \log_2(16/\eta)$. The tail after K is at most $2\kappa^{K+1}/(K+1)! \leq 2^{-K} \leq \eta/8$. Replace each time by a midpoint on a grid $M = 2^n \geq 8De^\kappa/\eta$. Sorting is nonexpansive in maximum coordinate distance. Telescoping products and summing their exponential weights bounds the grid error by $De^\kappa/(2M) \leq \eta/16$.

Prepare the length register with probabilities $(\kappa^j/j!)/\alpha$, where $\alpha = \sum_{j=0}^K \kappa^j/j! \leq e^\kappa$, together with K uniform n -bit time registers. Reversibly sort the selected times, retaining comparison bits; apply the ordered coefficient encodings on separate work registers; then undo sorting and state preparation. Repeated grid indices remain in the sum. This is the construction of [2, Appendix C.2, Eqs. (203)–(207)], instantiated with Lemma 3.1. It uses polynomially many operations in K, n , without a list of M^K tuples.

Normalized coefficient error ρ contributes at most $\kappa e^\kappa \rho$. Choose $\rho \leq \eta/(16\kappa e^\kappa)$. The $K+1$ rational factorial weights and their square roots can be prepared with polynomial bit cost. Allocate unitary synthesis error at most $\eta/(16\alpha)$ to this preparation and the remaining controls. All extra precision is $O(\kappa + \log(1/\eta))$.

Finally, the same construction with coefficient $-A(1-s)$ solves

$$Y(s) = O(1-s)O(1)^{-1}, \quad Y(0) = I, \quad Y(1) = O(1)^{-1}.$$

It therefore encodes the inverse with identical bounds. Reversing the first encoding alone would encode the adjoint and would not supply this inverse. $\square$

4.2 Amplification and the full recurrence

Proposition 4.2. *Let $h = \lceil \log_2 L \rceil$, $z = \log(1/\delta)$, and $0 < \delta < 1/4$. An L -site Gibbs purification can be prepared to vector error δ , up to phase, at gate and classical construction cost*

$$C_{\text{pur}}(L, \delta) \leq L^c e^{C\kappa h} [1 + b + \kappa + h + z]^c [1 + z + (\kappa + 1)h]^h. \quad (4.4)$$

All inverse calls, work-register tests and expanded circuit output are included.

Proof. Suppose the two compiled children prepare a normalized vector $\tilde{\phi}$ within error e of an exact product purification ϕ . Let K_O be the scaled projected block from Lemma 4.1, with $\|K_O - O\| \leq \eta$. If $\eta \leq e^{-\kappa}/2$, then for every normalized input

$$\|K_O \tilde{\phi}\| \geq e^{-\kappa}/2, \quad p_{\text{good}} \geq p_0 = \frac{1}{4}e^{-4\kappa}.$$

There is no assumed partition-function ratio.

Use the $\pi/3$ scheme [4]. If $W|0\rangle = \sqrt{p}g + \sqrt{1-p}b$, let R_0, R_{good} multiply the input-zero line and good subspace by $\omega = e^{i\pi/3}$. Define

$$W_0 = W, \quad W_{j+1} = W_j R_0 W_j^\dagger R_{\text{good}} W_j.$$

Writing $q = 1 - p$, the bad amplitude is multiplied by $1 + (\omega - 1)(\omega p + q) = -(\omega - 1)^2 q$, whose modulus is q . Failure becomes q^3 , while the good direction is preserved. After j levels it is q^{3^j} . Choosing the first

$$3^j \geq p_0^{-1} \log(128/\delta^2)$$

uses at most $3p_0^{-1} \log(128/\delta^2)$ calls to W or its inverse and makes the failure-amplitude contribution at most $\delta/4$. Every call includes both compiled child preparations. Multi-controlled phase tests use reversible conjunctions of the relevant work bits. Synthesis error $\delta/(100 \cdot 3^j)$ per phase leaves a total error below $\delta/4$.

These are inverses of the actual finite gate sequences. The amplification identity consequently applies to their actual output, and its good direction is exactly the normalized $K_O \tilde{\phi}$. The normalization inequality gives

$$\text{dist} \left(\frac{K_O \tilde{\phi}}{\|K_O \tilde{\phi}\|}, \frac{O\phi}{\|O\phi\|} \right) \leq 2e^\kappa \eta + 2e^{2\kappa} e. \quad (4.5)$$

Preparing each child to $\delta/(32e^{2\kappa})$ and taking $\eta \leq \delta e^{-\kappa}/32$ makes the complete node error less than δ . Old child work is retained as purification environment. Child error is therefore controlled through (4.5), rather than charged as an unexplained ideal-oracle error at every amplification call.

At a leaf, fixed-dimensional diagonalization and shifted Boltzmann weights give a purification by a fixed-size basis change and rotations. Weights below the assigned accuracy can be omitted. Large energies, very small weights and near-degeneracies require additional bits, not operations proportional to the energy or inverse weight.

Along a balanced tree path, logarithmic accuracy increases by at most $\log 32 + 2\kappa$ at each edge. Hence

$$z_{\text{leaf}} \leq z + (\log 32 + 2\kappa)h.$$

For a fixed-degree polynomial P , a safe recurrence is

$$C_h(\delta) \leq e^{C\kappa} (1 + \log(1/\delta)) [2C_{h-1}(\delta/(32e^{2\kappa})) + P(2^h, b, \kappa, \log(1/\delta))]. \quad (4.6)$$

The polynomial includes all nonrecursive merge work and phase-test registers. Amplification reuses registers; at most $2L$ tree nodes contribute. Replacing each amplification factor by its largest value along the path and summing these nodes proves (4.4). Expanding every repeated child gate list costs no more than its charged gate count, so the same bound holds classically. $\square$

For a reset interval of length $\ell \leq e^{Cx}$, its tree height is $O(x)$. Equation (4.4) therefore implies

$$C_{\text{pur}}(\ell, \delta) \leq e^{Cx^3} (1 + b)^c [1 + \log(1/\delta)]^{Cx+c}. \quad (4.7)$$

Polynomial factors in x and $\exp(O(x^2 \log(2+x)))$ are absorbed conservatively in e^{Cx^3} . The larger localization region will have a precision-dependent size. We prepare a Gibbs state only on the reset interval, not recursively on that larger region.

4.3 Collective reset and loss

On a region $R \supset B$, the boundary dressing uses $-(\beta/4)\mathcal{F}_{\beta, H_R(s)}(W_{\partial B})$. There are at most two interval ends, so its coefficient bound remains $O(x)$, independently of $|R|$. Lemma 4.1 implements this dressing and its separate inverse at fixed-degree polynomial cost in $|R|, b, x, \log(1/\eta)$.

Swap the old block input into an environment and prepare a Gibbs purification on B and another environment. This is an isometry V_0 for the reset, with operator error equal to the purification vector error. Define

$$V = (I_{\text{env}} \otimes O_R)V_0 O_R^{-1}. \quad (4.8)$$

It is the dressed reset stack in (2.1). Independent-ancilla block products give normalization $\alpha_V \leq e^{Cx}$. No Kraus list is enumerated. If Π selects the successful output of the stack encoding U_V , then $U_V^\dagger \Pi U_V$ encodes S/α_V^2 . The projection is implemented by a success-bit check, using $\Pi = (I + (2\Pi - I))/2$. If the stack error is e_V , its loss error is at most

$$(2\alpha_V + e_V)e_V. \quad (4.9)$$

All requested component errors are chosen smaller by e^{-Cx} factors to account for the product norms; this changes only their logarithmic precision budgets.

5 The exact coherent correction with logarithmic normalization

The straightforward commutator encoding of $G_H(S)$ can carry an onsite energy normalization. We instead implement its odd integral directly. Let $E \geq \max(1, \|H\|)$, $B_0 \geq \max(1, \|S\|)$, and let an actual encoding of S/B_0 be available. Its cost will be supplied by (4.8), not treated as an oracle of unit cost.

Lemma 5.1. *For $\beta > 0$, $0 < \eta < 1/4$, and Hermitian S , $G_H(S)$ in (2.1) has an encoding of error η and normalization*

$$\alpha_G = O(B_0[1 + \log(1 + \beta E B_0/\eta)]).$$

Its unitary completion can be chosen exactly Hermitian and involutive. Beyond a constant number of calls to the S encoding and its inverse, the gate and classical cost is polynomial of fixed degree in regional size, input bits, $1 + \beta\alpha_D$, $\log B_0$, and $\log(1/\eta)$. In particular $\log E$, not E , is charged.

Proof. The source correction identity [1, Eq. (3.11)], with $\tau_t(S) = e^{itH} S e^{-itH}$, is

$$G_H(S) = \frac{1}{\beta} \int_0^\infty \frac{\tau_t(S) - \tau_{-t}(S)}{\sinh(2\pi t/\beta)} dt. \quad (5.1)$$

It has the Hermitian sign convention of (2.1); this can also be checked entrywise using $\int_0^\infty \sin(at)/\sinh(bt) dt = (\pi/2b) \tanh(\pi a/2b)$.

Set

$$\begin{aligned} \nu &= 1 + \beta E B_0/\eta, & a &= \frac{1}{64\nu}, & U &= 2 + \lceil \log_2(128 B_0/\eta) \rceil, \\ r &= \lceil \log_2(1/a) \rceil + 2, & \alpha_G &= 2B_0(r+1). \end{aligned}$$

After $u = 2\pi t/\beta$, the small- u omission is bounded using $\|\tau_t(S) - \tau_{-t}(S)\| \leq 4EB_0|t|$ and $\text{csch } u \leq 1/u$. Its norm is at most $\beta E B_0 a/\pi^2 \leq \eta/(64\pi^2)$. The upper tail has norm at most $4B_0 e^{-U}/\pi < \eta/64$.

Let

$$I = \log \tanh(U/2) - \log \tanh(a/2), \quad w = B_0 I/\pi.$$

Since $I \leq \log(4/a) \leq r \log 2$, $w \leq \alpha_G$. The truncated integral is w times the average of $\sigma \tau_{\sigma t}(S/B_0)$, with uniform sign σ and u -density $\text{csch } u/I$ on $[a, U]$. A scalar rotation with projected value w/α_G pads the normalization.

The inverse CDF has an explicit formula. Put $L_a = \log \tanh(a/2)$, $L_U = \log \tanh(U/2)$, and

$$\ell(v) = (1-v)L_a + vL_U, \quad u(v) = \log \frac{1+e^{\ell(v)}}{1-e^{\ell(v)}}, \quad t(v) = \frac{\beta u(v)}{2\pi}. \quad (5.2)$$

Evaluate it reversibly on uniform midpoint indices. Stable arithmetic uses $\log(a/2) + \log(\tanh(a/2)/(a/2))$ at the small endpoint and $\log 1p(2e^\ell/[-\text{expm1}(\ell)])$ for the inverse. Throughout the interval,

$$1 - e^\ell \geq e^{-U}, \quad \left| \frac{du}{d\ell} \right| \leq Ce^U.$$

Thus $O(U + |\log \beta| + \log(1/\zeta) + \log(1/a))$ bits suffice for absolute time error ζ . Range-reduced logarithm and exponential series have polynomial bit cost. There is no operation count proportional to $1/a$ or e^U .

The normalized conjugation is $2E$ -Lipschitz. Apply (3.3) with a power-of-two grid

$$M \geq \max(1, 128\alpha_G E \beta U / \eta).$$

Its operator error is at most $2\alpha_G E(\beta U/(2\pi))/M$. Choose time error $\zeta \leq \eta/(128\alpha_G E)$, propagator error $\xi \leq \eta/(128\alpha_G)$, and normalized loss error $\rho \leq \eta/(128\alpha_G)$. Fresh product ancillas and the actual compiled propagator adjoint give conjugation error at most $2\xi + \rho$. Give the scalar rotation the same normalized error budget. Together with the two tails these bounds sum to less than η , leaving slack for synthesis.

Let W be the assembled finite unitary and K its projected block. Introduce one flag and define

$$U_G = (H_{\text{flag}} \otimes I) \begin{pmatrix} 0 & W \\ W^\dagger & 0 \end{pmatrix} (H_{\text{flag}} \otimes I). \quad (5.3)$$

This unitary is exactly Hermitian and squares to I ; its zero block is $(K + K^\dagger)/2$. Hermitization cannot increase error to the Hermitian target. The inverse in (5.3) is the exact reversal of the compiled sequence.

The maximum real time is $O(\beta[1 + \log B_0 + \log(1/\eta)])$. Lemma 3.1's charged interaction-picture construction therefore applies with polynomial bit budgets and no enumerated grid. This proves the stated cost. $\square$

For the reset loss take $B_0 = \alpha_V^2 \leq e^{C_x}$, padding to a known bound if needed. Equation (4.9) shows that a loss encoding at normalized error ρ is obtained by compiling the stack with $e_V \leq B_0 \rho / [4(1 + \alpha_V)]$. Hence the loss encoding and its inverse cost a constant number of actual stack calls. For a regional size m ,

$$\alpha_G \leq e^{C_x} [1 + b + \log m + x + \log(1/\eta)]. \quad (5.4)$$

If precision substitution produces powers of $1 + \log(1 + b)$ with exponent $O(x)$, the inequality

$$y^k \leq (k/e)^k e^y \quad (y, k \geq 1)$$

absorbs them into $e^{O(x \log(2+x))}$ times a fixed-degree polynomial in $1 + b$. The exponent of the input-bit factor need not grow with temperature.

6 Local channel simulation and its classical cost

The stacked-jump simulator [5] takes an exactly Hermitian involution for the Hamiltonian and a projected rectangular jump stack whose successful output labels are nonzero. Equation (5.3) supplies the first. Add an output flag equal to one on every successful stack output to supply the second. This adds no enumeration of environment labels.

Let $\tilde{V}, \tilde{G}$ be the exact projected matrices of the compiled encodings after scaling by their normalizations. They define a genuine Lindbladian

$$\tilde{\mathcal{K}}(X) = \text{Tr}_{\text{env}}(\tilde{V}X\tilde{V}^\dagger) - \frac{1}{2}\{\tilde{V}^\dagger\tilde{V}, X\} - i[\tilde{G}, X].$$

No exact Gibbs fixed point is required of this nearby generator. Stinespring multiplication bounds, valid with an arbitrary reference, give

$$\|\mathcal{K} - \tilde{\mathcal{K}}\|_\diamond \leq 2e_G + 2(2\alpha_V + e_V)e_V, \quad (6.1)$$

where $e_V = \|V - \tilde{V}\|$, $e_G = \|G - \tilde{G}\|$. Both semigroups are completely positive and trace preserving. Duhamel therefore yields

$$\|e^{T\mathcal{K}} - e^{T\tilde{\mathcal{K}}}\|_\diamond \leq T \|\mathcal{K} - \tilde{\mathcal{K}}\|_\diamond. \quad (6.2)$$

The stability factor is linear in T . Taking

$$e_G \leq \frac{\delta}{32(1+T)}, \quad e_V \leq \frac{\delta}{64(1+T)(1+\alpha_V)}$$

puts the perturbation error below $\delta/4$. Compile the separate loss stack used inside G to Lemma 5.1's required accuracy and allocate $\delta/2$ to the exact-oracle simulator and its remaining synthesis.

The source normalization is $\alpha_{\text{sim}} = \alpha_G + \alpha_V^2$, with $\tau = \alpha_{\text{sim}}T$. In particular it is not a normalization of the physical Hamiltonian H . The query bound is

$$O\left(\tau + \frac{\log(1/\delta)}{\log(e + \log(1/\delta)/\tau)}\right).$$

A fixed polynomial in $1 + \tau + \log(1/\delta)$ suffices for our bound. Classical work requires more than this query statement. In [5, Sections 6.2, 7.2 and 9.2–9.4], the coefficient support has size $O(q)$, and its binomial/filter coefficients are computable in polynomially many bit operations. We may choose

$$q = \lceil C(1 + \tau + \log(1/\delta)) \rceil, \quad J_{\text{grid}} = 2^j \geq \max(2\tau, 8q, 20\tau^2/\delta).$$

For sufficiently large universal C , this satisfies the source truncation criterion $\sqrt{\tau}(C_0\tau/q)^q \leq \delta/20$. The time position is binary. The source's two-by-two rotations and Algorithm 2 use $O(\log J_{\text{grid}})$ binary decisions with $O(q)$ retained labels, not a table of J_{grid} entries. Preparing coefficients, synthesizing rotations, controlling all oracle calls and inverses, and writing the resulting circuit have fixed-degree polynomial cost in $q, \log J_{\text{grid}}, \log(1/\delta)$, and the charged oracle registers.

Suppose now the reset length, update duration and localization region obey

$$\ell, T \leq e^{Cx}, \quad m \leq e^{Cx}[1 + \log(1/\delta)].$$

Equations (4.7) and (5.4), with the preceding oracle and simulator costs, give

$$C_{\text{upd}} \leq e^{Cx^3}(1+b)^c[1 + \log(1/\delta)]^{Cx+c} \quad (6.3)$$

for both gates and classical bit operations. This includes the reset purification, regional dressing and inverse, collective loss, correction, and actual simulator coefficients. Powers of $\log(1+b)$ are absorbed as above.

7 Global parameters and a bounded preparation schedule

7.1 One rounded Hamiltonian

After fixed-range grouping, let the chain have $n \leq N$ sites and $s = O(N)$ aggregated terms. Work at inverse temperature one with $K = \beta H_N$, rounding each term once to a Hermitian dyadic matrix with error at most $\varepsilon/(128s)$. Then $\|\widetilde{K} - K\| \leq \varepsilon/128$. Differentiating the normalized exponential gives

$$\frac{d}{dt}\rho(K + tV) = - \int_0^1 \rho^{1-u}[V - \text{Tr}(\rho V)I]\rho^u du.$$

Schatten Hölder bounds its trace norm by $2\|V\|$. Thus

$$\|\rho(\widetilde{K}) - \rho(K)\|_1 \leq \varepsilon/64. \quad (7.1)$$

All updates subsequently use the same $\widetilde{K}$.

For precision-supplied inputs, it suffices to allocate local errors $\varepsilon/[Cs \max(1, \beta)]$ and temperature error $\varepsilon/[C(1 + \sum_X \|\widehat{h}_X\|)]$, with C enlarged for matrix dimension and product rounding. The needed fractional bits are

$$O(\log(N/\varepsilon) + \log(1 + \beta) + \log(1 + \sum_X \|h_X\|)).$$

Input integer bits and redundant raw descriptions remain charged through b .

From constant-dimensional entry sums compute a conservative interaction bound $\overline{J} \leq C\beta J_{\text{int}} + 1$ for $\widetilde{K}$, excluding onsite terms. Put $X = 1 + \overline{J} \leq Cx$. In the remainder set

$$Q_* = 1 + \lceil \log_2(64N/\varepsilon) \rceil = O(Q).$$

If $\beta = 0$, prepare the maximally mixed product state. If all interactions vanish, use the charged single-site purification circuits.

7.2 Explicit temperature-dependent blocks

For the dimensionless chain, the analytic source bounds have the form

$$\begin{aligned} M_{\text{dr}}^4 &\leq e^{2(X-1)}, & \lambda &= c_0/X, \quad K_{\text{loc}} \leq C_0 e^{2(X-1)} X^7, \\ A_u &\leq C_1[(u+1) + X^2 e^{8(X-1)}(u+X)^2]. \end{aligned} \quad (7.2)$$

Here M_{dr} bounds dressing norms and A_u is the boundary-influence sum in [1, Propositions 5.7–5.11]. The constants depend only on fixed local parameters and can be hardwired.

Choose fixed sufficiently large integers A, C_L and define

$$u_0 = \lceil AX^2 \rceil, \quad \ell = \lceil C_L X^2 (u_0 + 2)^3 \rceil 2^{\lceil 12X \rceil}, \quad T = (u_0 + 2)2^{\lceil 4X \rceil}. \quad (7.3)$$

These are computed with polynomial bit arithmetic; $\log \ell, \log T = O(X)$. We verify that they satisfy the source's influence inequality rather than assuming monotonicity in block duration. Choose $C_L \geq \max(256, 160C_1)$. Equation (7.2), $u_0 \geq X$, and $2^{12X} \geq e^{8X}$ imply

$$\ell \geq 32(u_0 + 1)A_{u_0}, \quad \ell \geq 256(u_0 + 1)^2.$$

Also

$$M_{\text{dr}}^{-4} T \geq e^2 e^{(4 \log 2 - 2)X} (u_0 + 2) \geq \log(8[1 + 16(u_0 + 1)]).$$

In Eq. (5.31) of [1], these estimates bound each of the first two positive contributions by $\ell/8$. For its remaining locality contribution use

$$(1 - e^{-\lambda/2})^{-1} \leq CX, \quad \frac{e^{-\lambda/2}}{(1 - e^{-\lambda/2})^2} \leq CX^2.$$

After division by ℓ , that contribution is bounded by

$$\begin{aligned} CK_{\text{loc}}T(u_0 + 1)e^{-c_0u_0/(2X)} &\leq C'(A + 2)^2 X^{11} e^{(2+4\log 2 - c_0A/2)X} \\ &\leq C'(A + 2)^2 e^{-(c_0A/2 - 17)X}. \end{aligned}$$

Choose A so that $c_0A/2 > 17$ and $C'(A + 2)^2 e^{-(c_0A/2 - 17)} \leq 1/4$. It then holds for every $X \geq 1$. The source inequality yields $\sum_B [D_B(j) - 1] \leq -\ell/2$ for every site j , i.e. the required Dobrushin–Shlosman contraction parameter $1/2$.

7.3 Small chains and discrete contraction

If $n < \ell$, prepare the entire rounded chain by Proposition 4.2 at vector error $\varepsilon/16$. Its height is at most $O(X)$. Hence (1.1) applies, and (7.1) completes the error estimate. No exponential family of repeated full-chain labels is generated.

For $n \geq \ell$, retain all labels

$$a = 2 - \ell, \dots, n, \quad B_a = [a, a + \ell - 1] \cap [n].$$

There are $m_{\text{lab}} = n + \ell - 1 \leq 2n$ labels, and each site belongs to exactly ℓ labelled intervals, including edge multiplicities. Set

$$P = \frac{1}{m_{\text{lab}}} \sum_a \Phi_{B_a}, \quad a_{\text{tot}} = m_{\text{lab}}/\ell, \quad k = \lceil 4a_{\text{tot}}Q_* \rceil.$$

The discrete contraction theorem [1, Theorem 4.4, Eqs. (4.16)–(4.17)] gives, from any initial state,

$$\left\| P^k(\rho_0) - \rho(\widetilde{K}) \right\|_1 \leq 2n(1 - 1/(2a_{\text{tot}}))^k \leq 2Ne^{-2Q_*} < \varepsilon/32. \quad (7.4)$$

We choose a computational-basis product state. The update count is deterministically $O(NQ)$; no Poisson tail is needed for the preparation task.

7.4 Localization, finite random bits and the error sum

Put $\delta = \varepsilon/(64k)$. The localization statement in [1, Proposition 5.7 and Appendix B.1] bounds channel error by $TC_{\text{loc}}e^{-\mu v}$, with $\mu = c_0/(2X)$, $C_{\text{loc}} \leq e^{CX}$, for $v \geq \max(\text{range}, \ell)$. It is sufficient to choose

$$v = \max\{\text{range}, \ell, \lceil CX[X + 1 + \log(1/\delta)] \rceil\}$$

with a fixed sufficiently large C . The localized region has size $m \leq e^{CX}[1 + \log(1/\delta)]$, and localization error is at most $\delta/2$. Section 6 supplies the remaining $\delta/2$, so each compiled update satisfies

$$\|\Lambda_a - \Phi_{B_a}\|_{\diamond} \leq \delta.$$

Its gate and classical cost is (6.3), with x replaced by X and $\log(1/\delta) = O(Q)$.

For each step draw a uniform p -bit integer j and choose label index

$$f(j) = \left\lfloor \frac{m_{\text{lab}}j}{2^p} \right\rfloor, \quad p = \lceil \log_2(64km_{\text{lab}}/\varepsilon) \rceil.$$

Each index has either $\lfloor 2^p/m_{\text{lab}} \rfloor$ or $\lceil 2^p/m_{\text{lab}} \rceil$ preimages. Its law q thus obeys $\text{TV}(q, \text{unif}) \leq m_{\text{lab}}/2^p \leq \varepsilon/(64k)$. This is a bounded fair-bit procedure. Writing $P_q = \sum_a q_a \Phi_{B_a}$ and $\Lambda_q = \sum_a q_a \Lambda_a$, channel contractivity gives

$$\|P_q - P\|_{\diamond} \leq \varepsilon/(32k), \quad \|\Lambda_q - P_q\|_{\diamond} \leq \delta.$$

Telescoping the k steps yields

$$\|\Lambda_q^k - P^k\|_{\diamond} \leq \varepsilon/32 + k\delta \leq 3\varepsilon/64. \quad (7.5)$$

Combining (7.1), (7.4) and (7.5) gives trace-norm error at most $6\varepsilon/64$, within the theorem's 2ε budget with slack for dyadic error rounding. The randomized algorithm's output is precisely $\Lambda_q^k(\rho_0)$ after discarding its random bits. Individual trajectories need not themselves be Gibbs states.

8 Classical assembly and completion of the cost bound

Compile the at most $2N$ local templates to the common accuracy δ . Extracting their overlapping input neighborhoods costs at most $O(Nm \text{ poly}(b, Q))$, covered by N times (6.3). Pre-sample the k labels classically, then execute the corresponding ordinary circuits. This does not require a coherent selector over N templates or quantum random-access memory.

Even classical random access is unnecessary for the stated cost class. Form request records (label, step), sort by label, merge with the sorted template file, attach the expanded gate list to each request, and sort the resulting records by step. Merge sorting variable-length records uses $O(\log(N+k))$ passes, each linear in the total written size, with polynomial-length comparisons. The total template and expanded circuit size is

$$O((N+k)C_{\text{upd}} \text{ poly}(Q)).$$

Rebase local wire addresses by their region offsets, at polynomial cost in $\log N$ per gate. Fixed-size grouping is unpacked with constant overhead. Work registers, initialization, discard, and the bounded random-bit count are included. Every seed produces k updates within the same size bound.

Since $k \leq CNQ$, $X \leq Cx$, and $\log(N+k) = O(Q)$, substitution of (6.3) proves (1.1), after increasing the fixed constants. The input-reading and rounding work $N \text{ poly}(b, Q)$ is smaller than this bound. The small-chain branch was covered separately. This completes the proof of Theorem 1.1.

9 Scope and reproducibility

The result fills the implementation gap identified in [1, Remark 3.12], with a conservative cubic temperature exponent and an explicit temperature-dependent logarithmic exponent. QBP, recursive merging, collective jump encoding, Dobrushin–Shlosman contraction, ordered Dyson sums, fixed-point amplification, interaction-picture evolution and stacked Lindblad simulation are inherited methods. The new contribution is their fully charged implementation and composition in the stated input model.

The source archive accompanying this paper contains the exact problem statement, the frozen five-stage derivation, dependency metadata, scalar and finite-matrix controls, and the subsequent independent review records. The finite controls check filter arithmetic, ordered products, amplification, correction signs and global finite-bit accounting; they do not execute an end-to-end large-chain preparation or replace the analytic proof.

This work was developed in an author-directed, AI-assisted research workflow. AI systems assisted derivation, literature comparison, computation and drafting. The release underwent separate agent review of the proof and primary-source interfaces, followed by manuscript correspondence and compilation checks. That is internal review, not external peer review or formal proof certification. The comparison uses the cited versions and the sources checked through 22 October 2026.

References

- [1] C. Rouzé and D. Stilck França, *Rapid mixing of Gibbs samplers via quantum Dobrushin–Shlosman conditions*, [arXiv:2610.03578v1](#) (2026).
- [2] L. Kim, *Rapid mixing of quantum spin chains at any finite temperature*, [arXiv:2610.01190v1](#) (2026). Theorem 6.1, Corollary 6.2 and Appendix C.2.
- [3] G. H. Low and N. Wiebe, *Hamiltonian Simulation in the Interaction Picture*, [arXiv:1805.00675v2](#) (2019).
- [4] L. K. Grover, *A different kind of quantum search*, [arXiv:quant-ph/0503205v1](#) (2005).
- [5] B. Chen, M. Gao, X. Wang and S. Zhou, *Query-Optimal and Gate-Efficient Lindbladian Simulation*, [arXiv:2609.18757v1](#) (2026).
- [6] E. Bilgin and S. Boixo, *Preparing Thermal States of Quantum Systems by Dimension Reduction*, [arXiv:1008.4162v2](#) (2010).
- [7] T. Kuwahara, Á. M. Alhambra and A. Anshu, *Improved thermal area law and quasi-linear time algorithm for quantum Gibbs states*, [arXiv:2007.11174v2](#) (2021).