Quantum simulation of open systems with bounded jump activity

Jacob Kitchen

Supplementary Notes

These notes give the full formal statements, construction and proofs supporting the main article. The input consists of coherent factor oracles; access to a readable classical formula is not assumed. The notes use the same notation as the article and are self-contained. For $\lambda \geq 1$, define

$$\Phi_\lambda(x) = x + \frac{\lambda}{\log(e + \lambda/x)} \quad (x > 0), \qquad \Phi_\lambda(0) = 0. \tag{1}$$

The no-event propagator is

$$K(T) = \mathcal{T} \exp\left[-\int_0^T \left(\tfrac{1}{2} L(t)^\dagger L(t) + iH(t) \right) dt \right]. \tag{2}$$

1. Access assumptions and channel statement.
2. Weighted transfer and the compression certificate.
3. Adjustable Cayley cells.
4. Joint event network and parameter choice.
5. Circuit implementation and resource bounds.
6. Stopping, time discretization and data errors.
7. Accretive propagators.
8. Lower bounds and cap saturation.
9. Activity certificates and physical examples.

Supplementary Note 1. Access assumptions and channel statement

Let $\mathcal{H} = \mathbb{C}^d$, with a fixed jump alphabet of size m. Write

$$\begin{aligned} L(t) &= \sum_{\mu=1}^{m} |\mu\rangle \otimes L_\mu(t), \\ \mathcal{J}_t(X) &= \sum_{\mu} L_\mu(t) X L_\mu(t)^\dagger, \\ G_t &= L(t)^\dagger L(t) = \mathcal{J}_t^*(I), \\ \mathcal{L}_t(X) &= -i[H(t), X] + \mathcal{J}_t(X) - \tfrac{1}{2}\{G_t, X\}. \end{aligned} \tag{3}$$

We work in finite dimension and denote the solution channel of (3) by $\mathcal{E}_{T:0}$. Our diamond norm has no prefactor; for states we retain the usual $1/2$ in the trace distance.

Supplementary Definition 1 (Raw factor access)**.** The input oracles are unitary encodings of $F(t) = L(t)/\alpha$ and $F_H(t) = H(t)/\beta$, with $\|F(t)\|, \|F_H(t)\| \leq 1$ and known input and output embeddings. On any chosen uniform dyadic grid, we are given controlled coherent selectors for these encodings and their adjoints. Each selector preserves its time-address register. A call to either selector or its adjoint counts as one raw query. The query model gives black-box access to these unitaries, without a readable classical description of the unknown instance. In the gate model we charge the complete circuits implementing the selectors.

Garbage blocks are unrestricted and are retained in the transducer. Besides the oracles, the input consists of finite descriptions of T, α, β, the regularity bounds and the desired accuracy. A declared zero normalization means that the corresponding branch is absent. Define

$$A_{\mathrm{D}} = \alpha^2 T/2, \qquad A_{\mathrm{H}} = \beta T. \tag{4}$$

These are normalization budgets; the actual norms and their time integrals need not be known. Building a stack from separate jump encodings incurs the label-preparation and multiplexing costs, and may increase the normalization. A classical list of encodings does not, by itself, provide the coherent SELECT of Definition 1.

Assume the physical maps satisfy

$$\begin{aligned} \|H(t) - H(s)\| &\le K_H |t-s|, \\ \|\mathcal{J}_t - \mathcal{J}_s\|_\diamond &\le K_J |t-s|. \end{aligned} \tag{5}$$

Discontinuities in the Kraus phases or basis are allowed by (5). One may replace the Lipschitz constants by an effective modulus of continuity, provided it is supplied as input; the modulus determines the mesh and hence the address cost.

Fix the unraveling specified by the jump map and let N_T count its jumps. A hard cap r means $\Pr_\rho(N_T > r) = 0$ for every initial state ρ. More generally, suppose

$$\sup_\rho \Pr_\rho(N_T > r) \le p_{\mathrm{tail}}. \tag{6}$$

For a promise restricted to an input subspace, take the supremum only over that subspace. All conclusions then hold on it, with arbitrary references. The coefficient, jump-map and regularity bounds remain bounds on the whole domain. The input state is supplied: neither finding the promised subspace nor preparing a state in it is included in simulation.

Supplementary Theorem 2 (Event-sensitive channel simulation)**.** *Under Definition 1 and* (5)*, let $r \ge 0$ be an integer for which* (6) *holds. For $0 < \varepsilon < 1/8$ and $p_{\mathrm{tail}} \le \varepsilon/8$, there is a trace-preserving quantum algorithm approximating $\mathcal{E}_{T:0}$ to diamond error ε with*

$$\begin{aligned} Q &= O(1 + \Phi_\lambda(B)), \\ B &= A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}}(r+\lambda)}, \quad \lambda = \log(C/\varepsilon). \end{aligned} \tag{7}$$

Here C is a universal constant. Zero total action is identity and uses no queries. The algorithm initializes one system input, retains finite compiler registers coherently, and discards them at the end. It uses no postselection.

For hard caps, $A_{\mathrm{D}} \ge r \ge 1$ and $\varepsilon \le 1/100$, the worst-case query complexity is $\Theta(\Phi_\lambda(B))$, for supplied upper normalization budgets with system dimension and oracle registers allowed to grow. Gates, memory and source precision obey the separate bounds in Supplementary Note 5.

Supplementary Note 5 charges loading and finite-precision synthesis of the selected unitaries. The lower bound allows growing dimension and jump-label registers and counts the maximum number of queries. It says nothing about optimal gate count, fixed dimension, or an expected-query bound without a worst-case stopping limit.

Supplementary Note 2. Weighted transfer and compression

Let a unitary transducer on a known public/private direct sum be

$$S = \begin{pmatrix} A & C_0 \\ B_0 & D \end{pmatrix}, \qquad S \begin{pmatrix} I \\ \Gamma \end{pmatrix} = \begin{pmatrix} V \\ \Gamma \end{pmatrix}. \tag{8}$$

Here V must be unitary on the whole public space. To compile the transducer we need controlled $S, S^\dagger$, public routing, and a reflection about the public subspace with all compiler registers zero. No preparation oracle for Γ is used.

Supplementary Lemma 3 (Averaged disk filter)**.** *Suppose* $\Omega = \Omega^\dagger$, $\Omega^2 = I$, *and* $\|R - bI\| \le 1 - b$, *where* $0 < b \le 1$. *Then*

$$\left\| \frac{I + (iR\Omega)^2}{2} \right\| \le \vartheta, \qquad \vartheta = \frac{2(1-b)}{2-b}. \tag{9}$$

Proof. For $b < 1$, the disk condition is equivalent to

$$\|Rz\|^2 + \frac{b}{1-b} \|(I-R)z\|^2 \le \|z\|^2 . \tag{10}$$

It also holds for $R' = \Omega R \Omega$. Apply it first to R' and then to R. The sum of the two squared displacement norms is at least half the squared norm of their sum, namely $(I - RR')z$. Thus

$$\left\|RR'z\right\|^2 + \frac{b}{2(1-b)} \left\|(I-RR')z\right\|^2 \le \|z\|^2 . \tag{11}$$

Equivalently $\|RR' - (1-\vartheta)I\| \le \vartheta$, and hence $\|I - RR'\| \le 2\vartheta$. Since $(iR\Omega)^2 = -RR'$, this proves the claim. At $b = 1$, $R = I$ and the filter is zero. This is the linear-operator form of the averaged-composition estimate of[1] Proposition 2.4. □

Supplementary Theorem 4 (CQC with a weighted-transfer certificate)**.** *Assume* (8), $\|\Gamma\|^2 \le W$, $W, s \ge 0$, $0 < b \le 1$, *and an invertible analysis map* Z *for which*

$$\begin{aligned} \kappa(Z) \le e^s, \qquad & ZDZ^{-1} = iR\Omega, \\ \Omega^2 = I, \qquad & \Omega = \Omega^\dagger, \qquad \|R - bI\| \le 1 - b. \end{aligned} \tag{12}$$

For each integer $k \ge 1$, *a bounded-coefficient combination of reuse circuits defines a linear combination* $\widetilde{V}$ *with a block encoding of normalization* $3/2$, *satisfying*

$$\left\| V - \widetilde{V} \right\| \le \frac{W}{2k} e^s \vartheta^k. \tag{13}$$

It uses $O(k)$ *controlled calls to* $S, S^\dagger$. *If the right side is at most* $1/4$, *constant-overhead oblivious amplification gives a clean initialized isometry with error at most five times that bound. Coefficient preparation and reuse-register operations use* $O(k \,\mathrm{polylog}(k/\varepsilon))$ *additional gates, apart from supplied public-subspace routing and reflection costs, at the requested circuit accuracy. Neither* Z *nor* Γ *is implemented.*

Proof. Write

$$f(z) = (1+z^2)/2, \qquad g_n(z) = \frac{1}{n} \sum_{j=0}^{n-1} z^j. \tag{14}$$

The inherited reuse construction[2] has public compression $M_n = V - C_0 g_n(D)\Gamma$. Supplementary Note 5 describes the actual controlled loop. From unitarity and (8),

$$V^\dagger C_0 = \Gamma^\dagger (I - D). \tag{15}$$

For $c_j = 2^{-k}\binom{k}{j}$, the identity

$$z^{2j} g_{4k}(z) = \frac{(4k+2j) g_{4k+2j}(z) - 2j g_{2j}(z)}{4k} \tag{16}$$

expresses $f^k g_{4k}$ as a signed combination of reuse averages. The zero-length term has coefficient zero. The coefficients sum to one and their absolute sum is $3/2$, because $\mathbb{E}j = k/2$. All lengths are at most $6k$. Their combination gives

$$\widetilde{V} = V - C_0 f(D)^k g_{4k}(D)\Gamma. \tag{17}$$

Equation (15) and the polynomial identity $(I - D)g_{4k}(D) = (I - D^{4k})/(4k)$ yield

$$V^\dagger(V - \widetilde{V}) = \frac{\Gamma^\dagger f(D)^k (I - D^{4k})\Gamma}{4k}. \tag{18}$$

Only polynomials of the same D have been commuted. Since $\|D\| \le 1$, Lemma 3 and (12) prove (13).

The signed SELECT block is $\widetilde{V}/(3/2)$. Attenuation by $3/4$ makes it $\widetilde{V}/2$. One standard oblivious-amplification step has good block

$$\tfrac{3}{2}\widetilde{V} - \tfrac{1}{2}\widetilde{V}\widetilde{V}^\dagger\widetilde{V}. \tag{19}$$

If $\left\|\widetilde{V} - V\right\| \le \eta \le 1/4$, each singular value is $1 + e$, $|e| \le \eta$, and the amplified value is $1 - 3e^2/2 - e^3/2$. The bad-register amplitude is at most 2η; the polar unitary differs from V by at most 2η. The complete clean error is therefore at most 5η. The reflection includes every public input column. This argument cannot be applied directly to a nonunitary subblock of V. □

Composition need not preserve the disk hypothesis. For example, two real rotation transducers with public cosine $0 < a < 1$ can compose to give

$$R = \begin{pmatrix} a & 0 \\ -(1 - a^2) & a \end{pmatrix}. \tag{20}$$

On $(1, -a)$, this matrix preserves norm without fixing the vector. Such a vector violates the displacement inequality in Lemma 3 for every $b > 0$, ruling out $\|R - bI\| \le 1 - b$. We must therefore prove a disk bound for the composed network; local cell bounds alone will not suffice.

Supplementary Note 3. Adjustable Cayley cells

In a cell of width δ, let F be the normalized rectangular factor and F_H the normalized Hermitian Hamiltonian. From the supplied encoding U_F, form

$$O_F = \begin{pmatrix} 0 & U_F^\dagger \\ U_F & 0 \end{pmatrix}. \tag{21}$$

Known good embeddings E_c, E_b in its opposite flag sectors obey $E_b^\dagger O_F E_c = F$. For the Hamiltonian use the analogous involution O_H and the plus-flag embedding E_H, so $E_H^\dagger O_H E_H = F_H$. Put $O = O_F \oplus O_H$. The three embeddings are mutually orthogonal. These operations cost a constant number of controlled raw queries, retain all encoding garbage, and require no polar-factor or positive-square-root oracle.

Choose $q > 0$ and define

$$p = \frac{\delta\alpha^2}{4q}, \qquad x = p + \frac{\delta\beta}{2},$$
$$B_c = \sqrt{p}\, E_c^\dagger + \sqrt{\delta\beta/2}\, E_H^\dagger, \qquad B_b = i\sqrt{q}\, E_b^\dagger. \tag{22}$$

Let B be the block column with rows B_c, B_b. Its public carry space has dimension d and its bath has dimension md. Directly,

$$BB^\dagger = \mathrm{diag}(xI_d, qI_{md}),$$
$$2BOB^\dagger = \begin{pmatrix} \delta\beta F_H & -i\sqrt{\delta}\,\alpha F^\dagger \\ i\sqrt{\delta}\,\alpha F & 0 \end{pmatrix} = G_{\mathrm{cell}}. \tag{23}$$

Although B depends on q, the physical matrix G_{cell} does not.

Supplementary Lemma 5 (Weighted cell)**.** *Let $R_0 = (I + BB^\dagger)^{-1}$. Define the known free unitary*

$$S_0 = \begin{pmatrix} (I - BB^\dagger)R_0 & -2iR_0B \\ 2B^\dagger R_0 & i(I - 2B^\dagger R_0 B) \end{pmatrix} \tag{24}$$

Then $S = S_0(I_{\mathrm{public}} \oplus O)$ transduces $U = (I - iG_{\mathrm{cell}}/2)(I + iG_{\mathrm{cell}}/2)^{-1}$. For $y = Uu$, $v = (u + y)/2$, its private witness is

$$\gamma = (I + iO)B^\dagger v, \qquad \|\gamma\|^2 = 2x\,\|v_c\|^2 + 2q\,\|v_b\|^2. \tag{25}$$

With zero incoming bath, the cell's carry and jump operators are

$$A = L^\dagger L/2 + iH, \qquad R_\delta = (I + \delta A/2)^{-1},$$
$$M_\delta = 2R_\delta - I, \qquad J_\delta = \sqrt{\delta}\,LR_\delta. \tag{26}$$

They satisfy $M_\delta^\dagger M_\delta + J_\delta^\dagger J_\delta = I$.

Proof. The rows of B are explicitly orthogonal, so the unitarity of (24) reduces to scalar two-dimensional rotations and the phase i on their orthogonal complement. No singular-vector computation is required. Substitution of γ in $S_0(u \oplus O\gamma) = y \oplus \gamma$ gives $u - v = iBOB^\dagger v$, which is the physical Cayley relation. The identity $(I - iO)(I + iO) = 2I$ proves the norm formula. Eliminating the zero-input bath in (23) yields (26). Its last identity follows from $A + A^\dagger = L^\dagger L$. □

The circuit consists of the raw involution and the known coin. Resolvents appear in the proof to identify its target, so no inverse oracle is needed.

Supplementary Note 4. Joint event network

Fix J uniform cells and let $R = r + 1$. A history is an increasing list $h = ((j_1, \mu_1), \ldots, (j_n, \mu_n))$, $n \le R$. Associate one public copy of $\mathcal{H}$ to each valid history, including the empty history. Index cells by $j = 0, \ldots, J - 1$, with $\mathrm{last}(\emptyset) = -1$. At vertex (j, h), $|h| \le r$, $j > \mathrm{last}(h)$, the physical Cayley cell couples the carry wire h to the bath wires $h + (j, \mu)$. For each fixed j, these spaces are disjoint. Fix any order for the disjoint vertices at a given time. Their chronological product is therefore a unitary V_J on the whole public direct sum.

On the initialized empty-history column, a bath wire has zero amplitude before its birth. The column is exactly the chronological rational instrument with at most one jump per cell. Histories of length r continue their preferred evolution; the $(r + 1)$-st event is retained with its time and label, after which the history is frozen. Stopping immediately at the r-th jump would remove valid subsequent no-jump evolution.

Energy on the entire input boundary

Compose the local transducers before converting them to ordinary circuits. Their private spaces are disjoint mathematical summands, represented by an address and one shared payload. Private oracle $O_{j,h}$ commutes with the free coin at every other vertex. Thus

$$S_J = S_{\mathrm{free},J}(I_{\mathrm{public}} \oplus \Omega), \qquad \Omega = \bigoplus_{j,h} O_{j,h}. \tag{27}$$

A coherent vertex selector realizes Ω in a constant number of raw calls. History-independent Lindblad data use only the time part of that address. For a public vector u, define Γu by stacking the local witnesses of Lemma 5 evaluated on the physical prefix midpoints. Cancellation of internal public wires then gives

$$S_J \begin{pmatrix} I \\ \Gamma \end{pmatrix} = \begin{pmatrix} V_J \\ \Gamma \end{pmatrix}. \tag{28}$$

Supplementary Lemma 6 (Full-boundary energy)**.** *For the global witness of* (27)*,*

$$\|\Gamma\|^2 \le A_{\mathrm{D}}/q + A_{\mathrm{H}} + q(r+2). \tag{29}$$

The estimate holds for arbitrary inputs on every public history wire, including incoming bath superpositions and arbitrary references.

Proof. At a fixed time, the incoming carries and the outgoing carries are subsets of a unit-norm public vector. Midpoint convexity gives $\sum_h \left\|v_{c,j,h}\right\|^2 \le 1$. Their total contribution is at most $2Jx = A_{\mathrm{D}}/q + A_{\mathrm{H}}$.

Let U_b, Y_b be the total squared norms of bath inputs and outputs over all vertices. A bath input is touched for the first time at its unique parent vertex, so $U_b \le 1$. Let N be history length on the public boundary. Each local unitary couples depths $n, n+1$, and changes its expectation by $\|y_b\|^2 - \|u_b\|^2$. Consequently

$$Y_b - U_b = \langle N \rangle_{\mathrm{out}} - \langle N \rangle_{\mathrm{in}}. \tag{30}$$

Since $\langle N \rangle_{\mathrm{in}} \ge U_b$ and $\langle N \rangle_{\mathrm{out}} \le R$, we have $Y_b \le R$. Equation (25) and midpoint convexity give the remaining contribution at most $q(U_b + Y_b) \le q(r+2)$. Reverse scattering from an occupied child input is included in this identity; monotonicity of every full-boundary path is neither asserted nor needed. □

An exact private-transfer calculation

Use the bright coordinates of B_c, B_b. In these coordinates the carry and bath cosines are $a = (1-x)/(1+x)$ and $b_0 = (1-q)/(1+q)$; private inputs also receive phase i. There is no public carry–bath mixing. Consequently each history gives one scalar path: carry rotations for the root, a bath head followed by carry rotations for a nonempty history, and just a bath head at overflow.

For a path with port weights $w_0, \ldots, w_n$, write its real private compression as R_{path}, before the phase i. Its Cayley admittance is

$$\begin{aligned} K_{\mathrm{path}} &= (I - R_{\mathrm{path}})(I + R_{\mathrm{path}})^{-1}, \\ (K_{\mathrm{path}})_{jj} &= w_j, \qquad (K_{\mathrm{path}})_{jk} = 2\sqrt{w_j w_k} \quad (j > k). \end{aligned} \tag{31}$$

The entries above the diagonal are zero. To see the formula, write the public amplitude change at port j as $-2\sqrt{w_j}\, v_j$, with private midpoint v_j, and sum the changes at earlier ports. More explicitly, let ξ_j, η_j be the private input and output at port j, $v_j = (\xi_j + \eta_j)/2$, and

$d_j = (\xi_j - \eta_j)/2$. For the real rotation with cosine $(1 - w_j)/(1 + w_j)$ and sine $2\sqrt{w_j}/(1 + w_j)$, the public amplitude u_j satisfies

$$u_{j+1} = u_j - 2\sqrt{w_j}\, v_j, \qquad d_j = w_j v_j - \sqrt{w_j}\, u_j. \tag{32}$$

The private compression sets the public input to zero, so $u_0 = 0$ and

$$d_j = w_j v_j + 2 \sum_{k<j} \sqrt{w_j w_k}\, v_k. \tag{33}$$

Since $d = (I - R_{\mathrm{path}})\xi/2$ and $v = (I + R_{\mathrm{path}})\xi/2$, this is exactly the lower-triangular admittance in (31). For finite nonnegative weights, $I + R_{\mathrm{path}}$ is invertible. As in the cell proof, this inverse is used only to analyze the circuit.

Choose $s, \nu > 0$ and weight all private rails at vertex (j, h) by

$$Z_{j,h} = \exp(-sj/J - \nu|h|). \tag{34}$$

Then $\kappa(Z) \le e^{s+\nu r}$. The weight is scalar on the entire vertex payload, including encoding garbage, so it commutes with Ω and with the known local bright-coordinate changes. On a nonroot path the first shift weight is $\gamma = e^{-s/J-\nu}$, and the remaining weights are $\rho = e^{-s/J}$. The head belongs to its parent's depth, which accounts for the extra factor $e^{-\nu}$.

Supplementary Lemma 7 (Weighted path disk)**.** *For the joint private block there are $\widehat{R}$ and the same involution Ω such that $ZDZ^{-1} = i\widehat{R}\Omega$ and*

$$\left\| \widehat{R} - bI \right\| \le 1 - b, \qquad b = \frac{1}{1 + c}, \tag{35}$$

where it suffices to take

$$c = \max\left\{ \frac{2(q + x\gamma^2)}{1 - \gamma^2}, \frac{2x(1 + \rho^2)}{1 - \rho^2}, 2x \right\}. \tag{36}$$

Root and head-only paths obey the same upper bound.

Proof. Let T_s be the strict lower shift with first weight γ and subsequent weights ρ, and $W_p = \mathrm{diag}(q, x, \ldots, x)$. Equation (31) gives

$$\widehat{K} = W_p^{1/2} F(T_s) W_p^{1/2}, \qquad F(T_s) = (I + T_s)(I - T_s)^{-1}. \tag{37}$$

Direct multiplication gives

$$\mathrm{Re}\, F(T_s) = (I - T_s^\dagger)^{-1}(I - T_s^\dagger T_s)(I - T_s)^{-1}. \tag{38}$$

Thus $\widehat{K}^\dagger \widehat{K} \le c\, \mathrm{Re}\, \widehat{K}$ follows by congruence from

$$(I + T_s^\dagger) W_p (I + T_s) \le c(I - T_s^\dagger T_s). \tag{39}$$

The left quadratic form is at most $2\sum_j (w_j + w_{j+1} t_j^2)|z_j|^2$, where t_j is the next shift weight; the terminal $w_{j+1} t_j^2$ term is zero. The right form is diagonal. Equation (36) therefore suffices. Expanding the Cayley transform $\widehat{R} = (I - \widehat{K})(I + \widehat{K})^{-1}$ shows that its disk inequality is equivalent to $b\widehat{K}^\dagger \widehat{K} \le (1 - b)\, \mathrm{Re}\, \widehat{K}$. The private complement has $\widehat{R} = I$. Taking the direct sum proves the result without replacing a nonnormal operator norm by a spectral radius. □

Balancing action, events and precision

Let $\lambda = \log(C/\varepsilon) \ge 1$, assume $A_{\mathrm{D}} > 0$, and choose

$$q_* = \sqrt{A_{\mathrm{D}}/(r + \lambda)}, \qquad q \in [q_*, 2q_*],$$
$$s = \lambda, \qquad \nu = \lambda/(r + \lambda),$$

$$v = \sqrt{A_{\mathrm{D}}(r+\lambda)}, \qquad B = v + A_{\mathrm{H}}. \tag{40}$$

The implemented q can be dyadic. Take a certified $B_+ \in [B, 2B]$ and a dyadic grid satisfying

$$J \geq \max\{1, r, \lambda, \lceil 2v + A_{\mathrm{H}} \rceil\}. \tag{41}$$

Then $x = (A_{\mathrm{D}}/q + A_{\mathrm{H}})/(2J) \leq 1/2$, and $\kappa(Z) \leq e^{2\lambda}$. Since $1 - e^{-2\nu} \geq (1 - e^{-2})\nu$ for $0 < \nu \leq 1$, (36) gives

$$\begin{aligned} c &\leq 3q(r+\lambda)/\lambda + 5x + 5xJ/\lambda \\ &\leq 11B/\lambda. \end{aligned} \tag{42}$$

The first inequality uses $r \leq J$; the second also uses $J \geq \lambda$. Meanwhile (29) gives $W \leq 5v + A_{\mathrm{H}} \leq 5B$.

Retain the exact filter rate, rather than replacing it by a linear exponential estimate. Since $\vartheta = 2c/(1+2c)$, choose

$$k = \left\lceil \frac{4\lambda}{\log(1 + \lambda/(22B_+))} \right\rceil. \tag{43}$$

The inequality $\log(1+u) \leq u$ gives $k \geq 88B_+$, so $W/(2k) \leq 1/4$. Equations (13) and (42) give block error at most $e^{-2\lambda}/4$. Theorem 4, with a constant fraction of ε reserved for synthesis, proves

$$Q = O\left(1 + \frac{\lambda}{\log(1 + \lambda/B)}\right) = O(1 + \Phi_\lambda(B)). \tag{44}$$

The equality of orders follows by splitting $\lambda/B \leq 1$ and $\lambda/B \geq 1$. The additive one covers integer rounding.

If $A_{\mathrm{D}} = 0$, omit the bath and factor ports. A single Hamiltonian carry path has $W \leq A_{\mathrm{H}}$, $\kappa(Z) \leq e^{\lambda}$ and $c \leq 2A_{\mathrm{H}}/\lambda$ on $J \geq \max\{1, \lambda, A_{\mathrm{H}}\}$. The same argument gives (44) with $B = A_{\mathrm{H}}$. This is the usual Hamiltonian endpoint. With both actions zero, the target is identity.

For a prescribed finite instrument, (41) is a hypothesis: subdividing a cell would change that instrument. For continuum simulation we are free to choose the grid, and choose it fine enough for both (41) and the discretization estimate below.

Supplementary Note 5. Circuit implementation and resources

A history address and one system payload suffice to represent the exponentially many network wires. We implement the free unitary and coefficient preparation on these addresses and charge their routing and inverse circuits below.

Head-and-suffix paths

Use a canonical array of at most $R = r + 1$ time/label pairs, its length, and public/private and head/carry flags. Bright-coordinate addresses are given by

$$\begin{aligned} (j, h, \mathrm{carry}) &\longleftrightarrow (h, j), \\ (j, h, \mu, \mathrm{bath}) &\longleftrightarrow (h + (j, \mu), \mathrm{head}). \end{aligned}$$

Deleting the last event inverts the head map. Since the map is injective on valid addresses, compute the new address, swap registers, and uncompute the old address using the inverse map. Each such call scans one canonical record at a cost of $O(R\,\mathrm{poly}(w))$ gates, where $w = O(\log(Jm(R+1)))$.

The free unitary is a bath-head rotation followed by a uniform carry suffix, with phase i on private inputs. Its head has cosine b_0, negative when $q > 1$, and sine $2\sqrt{q}/(1+q)$. Reversibly

compute from the bits of t an ordered decomposition of $[t, J)$ into $O(\log J)$ aligned dyadic intervals. For $J = 2^n$, an interval of length 2^ℓ uses the scalar cascade identity of[3]:

$$C_{2^\ell}(a) = W_{\text{out}}\, \mathrm{Rot}(a^{2^\ell})\, W_{\text{in}}^\dagger. \tag{45}$$

The central rotation couples the path's public coordinate to its first private leaf. The private transforms use

$$\begin{aligned} M_{\text{in}}(z) &= \frac{1}{\sqrt{1+z^2}} \begin{pmatrix} z & 1 \\ 1 & -z \end{pmatrix}, \\ M_{\text{out}}(z) &= \frac{1}{\sqrt{1+z^2}} \begin{pmatrix} 1 & z \\ z & -1 \end{pmatrix}. \end{aligned} \tag{46}$$

Apply these in decreasing bit order: on bit j, set $z = a^{2^j}$ and condition on all lower bits being zero. Address offsets and high-bit tests select the interval. The identity holds on all incoming columns, including nonvacuum private inputs.

This gives $O(\log^2 J)$ structural rotations for all paths in superposition, plus arithmetic and record routing. Overflow paths omit the suffix; the root omits the head. Unused private coordinates retain phase i. Reverse the address changes, interval order and gates for the inverse. Neither direction queries the factor oracle.

The private oracle in (27) selects a direct-sum block by vertex address, preserves that address and clears its address scratch. It leaves block-encoding garbage in the private payload. These properties ensure the commutations in (27) and preserve valid logical addresses with initialized reversible scratch. Extend the operation by the identity on extra public coordinates, give them zero witness, and assign phase i to private spectators.

Reuse and coefficient circuits

For reuse, take one payload and a zero-initialized register of dimension at least $n+1$. On the public sector, prepare it uniformly over $0, \ldots, n-1$. For $\ell = 0, \ldots, n-1$, apply S when the register equals ℓ, then increment it on the private sector. Each call consists of the known free unitary and the selected raw involution. Unprepare the public uniform state after the loop. Summing the recurrence $z_{\ell+1} = B_0 u/\sqrt{n} + D z_\ell$, with $B_0 = (I - D)\Gamma$, gives $M_n = V - C_0 g_n(D)\Gamma$. A coherently selected n controls the stopping of every operation. The inverse reverses the complete loop.

For the signed coefficients of Theorem 4, prepare three branches with probabilities $2/3, 1/6, 1/6$. The first uses $j \sim \mathrm{Bin}(k, 1/2)$ and positive length $4k + 2j$. The other two use $j = 1 + \mathrm{Bin}(k-1, 1/2)$: one has length $4k + 2j$ with positive sign; the other has length $2j$ with negative sign. Before rotation synthesis these coefficients are exact. Retain the branch and count registers until unpreparation.

Prepare the binomial state with a count register and one reusable coin, using logarithmic space and $O(k)$ arithmetic rotations. Given the state for n trials, append a Hadamard coin and increment the count if it is one. At new count j, the coin amplitudes are proportional to $\sqrt{n+1-j}$ and $\sqrt{j}$. A rotation controlled by j, n returns the coin to zero and leaves the state for $n+1$ trials. Use the identity outside counts $0, \ldots, n+1$. Reversing these operations gives the inverse without storing or erasing a list of Bernoulli outcomes.

For coherent reuse length n, prepare the uniform interval by a binary prefix tree. Compute each child's and parent's intersection sizes with $[0, n)$, rotate by the square root of their ratio, and uncompute the sizes; choose a unitary extension on empty branches. A loop of maximum length $6k$ handles all selected lengths, at their maximum cost. Coefficient preparation is called a constant number of times.

Amplification reflects about zero coefficient, attenuation, reuse and scratch registers tensored with the entire public history space. It requires no reflection about the unknown system state

and is not restricted to the root history. The full-public proof then applies to the initialized-root column; the circuit may initialize the root at the start. Afterwards, select an output block or trace out the records.

A finite gate and memory bound

Let $G_F, G_H, G_{\mathrm{addr}}, G_{\mathrm{emb}}$ denote the actual costs of controlled source calls, address computations and known good-embedding circuits. Write a_F, a_H for the source workspace widths and b_{in} for the bit length of the scalar inputs. Full-unitary synthesis accuracy $O(\varepsilon/Q)$ gives the sufficient gate bound

$$\begin{aligned} G_{\mathrm{total}} = O\big(Q[&G_F + G_H + G_{\mathrm{addr}} \\ &+ G_{\mathrm{emb}} + a_F + a_H \\ &+ R\,\mathrm{poly}(w) + b_{\mathrm{work}}^8]\big). \end{aligned} \tag{47}$$

$$\begin{aligned} b_{\mathrm{work}} &= O(b_{\mathrm{in}} + \log \Xi), \\ \Xi &= \frac{(J+1)(Q+1)(R+1)(m+1)}{\varepsilon}. \end{aligned} \tag{48}$$

The unoptimized exponent eight covers reversible arithmetic. Space consists of one system/oracle payload, $O(Rw)$ record and address qubits, $O(\log Q)$ reuse/count registers and $\mathrm{poly}(b_{\mathrm{work}})$ arithmetic scratch.

Repeated squaring for a^{2^j} amplifies error by at most $2J$, so $O(\log J + \log(Q/\varepsilon))$ guard bits suffice. Cosine error u^2 causes rotation error $O(u)$, including at the endpoints. Compute the bath sine from its formula to avoid cancellation, and use reciprocals for large q. Division, square roots, bounded-angle trigonometric evaluation and bisection take polynomially many bit operations; retain and then uncompute their scratch. A certified lower approximation to the logarithm in (43), within a factor two, avoids exact real comparisons and increases the query count by at most a constant factor. The bounds include uniform coefficient preparation, inverses and all loop controls.

Charge matrix loading and encoding construction to the source costs. Arbitrary explicit tables need not have cheap SELECT access under (47) or the query theorem. Classical record readout costs its length; observable estimation requires repetitions or an estimation algorithm. Full tomography is not part of the output guarantee.

Supplementary Note 6. Stopping and approximation errors

We prove the continuum assertion in Theorem 2 by stopping the physical process before approximating its coefficients. This order matters: rounding a jump operator can destroy an exact cap.

Stopping as a completely positive operation

Set $R = r + 1$, and introduce a counter with basis $0, \dots, R$. Let

$$S_c = \sum_{n=0}^{r} |n+1\rangle\langle n|, \qquad P_c = S_c^\dagger S_c. \tag{49}$$

The stopped Lindblad generator has Hamiltonian $P_c \otimes H(t)$, jumps $S_c \otimes L_\mu(t)$, and loss $P_c \otimes G_t$. Count R is absorbing. Its jump map is $\mathrm{Ad}_{S_c \otimes I} \circ (\mathrm{id} \otimes \mathcal{J}_t)$. Its rational no-event operator is $P_c \otimes M_\delta + (I - P_c) \otimes I$, and its event operators are $S_c \otimes J_{\delta,\mu}$. Starting at counter zero gives the joint finite instrument with only the record length retained; the event times and labels are traced out. Equality on all system matrix units also gives equality in the presence of an arbitrary reference.

Supplementary Lemma 8 (Linear channel-tail error)**.** *If* (6) *holds, the original and stopped system channels differ by at most* $2p_{\mathrm{tail}}$ *in unhalved diamond norm on the promised input space.*

Proof. Decompose both channels at the first overflow. Their no-overflow completely positive map is common. Their two overflow maps have the same effect operator F_R, because the continuations after first overflow are trace preserving. For an all-input promise, $0 \le F_R \le p_{\mathrm{tail}} I$; on a promised input subspace, $P_{\mathrm{in}} F_R P_{\mathrm{in}} \le p_{\mathrm{tail}} P_{\mathrm{in}}$. The diamond norm of a completely positive map is the operator norm of its effect on the identity. Apply this identity after restricting the input embedding when necessary. The triangle inequality therefore bounds the difference by $2p_{\mathrm{tail}}$. The linear dependence on the tail probability comes from comparing the reduced channels. Comparing coherent trajectory amplitudes generally incurs a square-root probability loss. □

Grid and data errors

Put $\gamma = \beta + \alpha^2$. Both the physical and stopped generators have diamond norm at most 2γ. Since $\|G_t - G_s\| \le \|\mathcal{J}_t - \mathcal{J}_s\|_\diamond$, their variation is bounded by

$$\left\| \widehat{\mathcal{L}}_t - \widehat{\mathcal{L}}_s \right\|_\diamond \le 2(K_H + K_J)|t - s|. \tag{50}$$

No factor depending on counter dimension appears.

Supplementary Lemma 9 (One-cell channel approximation)**.** *For constant coefficients of norm at most* α, β*, let*

$$\mathcal{C}_\delta(X) = M_\delta X M_\delta^\dagger + \delta \mathcal{J}(R_\delta X R_\delta^\dagger). \tag{51}$$

Then $\mathcal{C}_\delta$ *is CPTP and*

$$\left\| \mathcal{C}_\delta - e^{\delta \mathcal{L}} \right\|_\diamond \le 4\gamma^2 \delta^2. \tag{52}$$

Proof. Trace preservation is Lemma 5, and positivity is explicit. Let $g = \beta + \alpha^2/2$. Accretivity gives $\|R_\delta\| \le 1$, $\|R_\delta - I\| \le \delta g/2$, and $M_\delta = I - \delta A R_\delta$. Expanding around $I + \delta\mathcal{L}$, the no-event quadratic and resolvent remainders are bounded by $2g^2\delta^2$, and the jump remainder by $\alpha^2 g \delta^2$. Their sum is at most $2\gamma^2\delta^2$. The integral remainder of the exact semigroup is at most $\delta^2 \|\mathcal{L}\|_\diamond^2 / 2 \le 2\gamma^2\delta^2$, since the semigroup is CPTP and has diamond norm one. Add these estimates. □

Freeze the coefficients at the left endpoint of each cell. Apply Duhamel's formula with (50), then telescope the CPTP cell products to obtain

$$\delta_{\mathrm{grid}} \le [K_H + K_J + 4\gamma^2] T^2 / J. \tag{53}$$

All intervening channels have diamond norm one, so telescoping introduces no additional stability factor.

At the sampled points, suppose the actual signal blocks are contractions $\widetilde{F}, \widetilde{F}_H$, with $\widetilde{F}_H$ Hermitian, and

$$\left\| \widetilde{F} - F \right\| \le \zeta_F, \qquad \left\| \widetilde{F}_H - F_H \right\| \le \zeta_H. \tag{54}$$

Hermitianizing an approximate Hamiltonian block through the plus-flag involution does not increase its error or norm. The Stinespring difference for the jump maps is at most $2\alpha^2\zeta_F$ in diamond norm. Thus

$$\delta_{\mathrm{data}} \le 2A_{\mathrm{H}}\zeta_H + 8A_{\mathrm{D}}\zeta_F. \tag{55}$$

Here the approximate samples define a piecewise-constant stopped generator. Duhamel's formula compares it with the exact stopped generator; the cell estimate compares it with its

rational instrument. No approximate coefficient is needed between samples, nor must the samples preserve the original grading or cap: the counter enforces the stop.

Let η be the clean error of the compiled actual finite isometry, including full-unitary synthesis. Tracing its records and work registers gives channel error at most 2η. The complete bound is

$$\left\|\mathcal{E}_{T:0} - \widetilde{\mathcal{E}}\right\|_{\diamond} \le 2p_{\text{tail}} + [K_H + K_J + 4\gamma^2]T^2/J + 2A_{\text{H}}\zeta_H + 8A_{\text{D}}\zeta_F + 2\eta. \tag{56}$$

For example, choose

$$\begin{aligned} p_{\text{tail}} &\le \varepsilon/8, & \eta &\le \varepsilon/8, \\ J &\ge 4[K_H + K_J + 4\gamma^2]T^2/\varepsilon, & & \\ \zeta_H &\le \frac{\varepsilon}{16(1+A_{\text{H}})}, & \zeta_F &\le \frac{\varepsilon}{64(1+A_{\text{D}})}. \end{aligned} \tag{57}$$

Also impose (41) and round J upward to a power of two. Together with (44), these choices prove the upper statement of Theorem 2. For $T = 0$ the channel is identity.

Equation (57) controls errors in the signal blocks. The compiler takes the actual unitary completions as its exact input, so errors in implementing their controlled calls and inverses require a separate full-operator hybrid budget, such as $O(\varepsilon/Q)$ per call in (47). This additional budget also controls the garbage blocks, whose accuracy cannot be inferred from the good blocks.

Which regularity is used. Time-dependent unitary mixing and phases of the Kraus labels preserve $\mathcal{J}_t$, so they preserve the channel comparison. A prescribed coherent environment wavefunction can still change. For example, choosing $L(t) = A$ on dyadic rationals and $L(t) = -A$ elsewhere gives the same constant jump map as $L(t) = A$, while reversing the one-event continuum amplitudes almost everywhere. A selector sampled at dyadic times cannot recover that amplitude gauge. Thus (5) suffices for the channel theorem; a fixed coherent trajectory output needs stronger regularity of the chosen factors.

Counting also depends on the specified jump map. Regularity of the uncounted generator alone is insufficient: $L(t) = c(t)I$ has zero dissipator but a variable counting process. Any label-dependent feedback must therefore be part of the declared payload and jump map.

Supplementary Note 7. Accretive propagators

In the first-event network the root carry block is $M_{J-1} \cdots M_0$. After compiling on the whole public space, we can select this block to obtain the no-event propagator. No bound on the number of physical jumps is needed for this block. The complete physical channel also contains jump branches.

Supplementary Theorem 10 (Factor-access accretive propagation)**.** *Suppose Definition 1 holds and $A(t) = L(t)^\dagger L(t)/2 + iH(t)$ is Lipschitz with known constant K_A. There is a unitary circuit whose designated good block approximates $K(T)$ in* (2) *to operator error ε, at normalization one, using*

$$O\Big(1 + \Phi_\lambda(A_{\text{H}} + \sqrt{A_{\text{D}}\lambda})\Big) \tag{58}$$

raw queries. For $A_{\text{D}} \ge 1$ and $\varepsilon \le 1/100$, this is optimal up to constants in the worst case over supplied normalization budgets. No event promise, gap, mixing-time bound or initial-state overlap is required.

Proof. Use the network with $r = 0$ and amplify its unitary full-public target before selecting the carry block. That block contains exactly the no-event factors, regardless of the probability of jumps in the physical process. Equation (44) gives (58).

Both e^{-tA} and the Cayley factor M_δ are contractions. Writing $g = \alpha^2/2 + \beta$, an exact resolvent expansion and the integral Taylor remainder give $\|M_\delta - e^{-\delta A}\| \le 2g^3\delta^3$. Freezing and telescoping therefore give

$$\delta^K_{\mathrm{grid}} \le K_A T^2/(2J) + 2g^3T^3/J^2. \tag{59}$$

For actual sampled contractions, the resolvent identity gives $M_X - M_Y = -\delta R_X(A_X - A_Y)R_Y$, so

$$\delta^K_{\mathrm{data}} \le A_{\mathrm{H}}\zeta_H + 2A_{\mathrm{D}}\zeta_F. \tag{60}$$

Choose the grid, data and compiled clean-isometry errors to sum to ε, and also impose (41). The mesh size contributes through the charged address registers and scalar arithmetic. The scalar and Hamiltonian families in Supplementary Note 8 already give the matching lower bound; the event family is unnecessary for this output problem. □

Using the designated block for state preparation succeeds on $|\psi\rangle$ with probability approximately $\|K(T)|\psi\rangle\|^2$; the corresponding success-amplification cost must be included in that application. The theorem also requires access to the factors themselves. Starting from an encoding of A or of $L^\dagger L$ requires a factor realization and its cost.

Supplementary Note 8. Lower bounds and cap saturation

Set $a = A_{\mathrm{D}}$, $h = A_{\mathrm{H}}$, and $\lambda = \log(1/\varepsilon)$ in this section. Constant rescalings of ε do not affect the estimates for $\varepsilon \le 1/100$. All lower bounds permit controls, inverses and coherent time selection, and use a worst-case bounded number of raw calls.

An event lower bound for the reduced channel

Supplementary Lemma 11. *For every real $a \ge r \ge 1$, with r integral, there is a stationary Lindblad family with dissipative action a, zero Hamiltonian and at most r jumps on every input, whose reduced-channel simulation requires at least $\sqrt{ar}/8$ raw queries at error $\varepsilon \le 1/100$.*

Proof. Set $L = \lfloor 2a/r \rfloor$, $M = rL$, and $x \in \{0, 1, \dots, L\}^r$. Group b contains either no marked address or the one marked address x_b. On r live/sink qubits, put $A_b = |\mathrm{sink}\rangle\langle\mathrm{live}|_b$ and supply the stationary rectangular factor

$$F_x = \frac{1}{\sqrt{M}} \sum_{b=1}^{r} \sum_{j=1}^{L} |b, j\rangle \otimes \mathbf{1}[x_b = j] A_b. \tag{61}$$

Take the jump stack $\sqrt{2}F_x$, $H_x = 0$, supplied normalization $\alpha = \sqrt{2}$, and horizon $T = a$. These give the stated action. All jumps lower the excitation number and $G_x = (2/M)\sum_b \mathbf{1}[x_b \ne 0] A_b^\dagger A_b$ commutes with it. Thus no-event evolution preserves grades and the cap holds on every input. Both physical regularity constants are zero. Since $a \le M \le 2a$, a marked qubit survives with probability $q = e^{-2a/M} \le e^{-1}$. On the known all-live input, the reduced output is a product with factors

$$\begin{aligned} \sigma_0 &= |\mathrm{live}\rangle\langle\mathrm{live}|, \\ \sigma_j &= q|\mathrm{live}\rangle\langle\mathrm{live}| \\ &\quad + (1-q)|\mathrm{sink}\rangle\langle\mathrm{sink}|. \end{aligned} \tag{62}$$

Different nonzero positions have the same marginal output; only marked versus unmarked will be used.

A legal raw completion prepares the uniform superposition of all M addresses, computes $\mathbf{1}[x_b = j]$ in one membership query, and applies a known normalization-one encoding of A_b, controlled by b. The good output retains the address, selects membership response one and

the lowering encoding's good flag, and is exactly F_x. The input and output embeddings are known. Each raw call, inverse and control uses one membership query; time selection adds no information. No list of marked addresses or sampler for their actual rates is supplied.

Purify any simulator, retaining all of its arbitrary garbage, and call the final states ϕ_x. On an edge with $x_b = 0$, $y_b = j > 0$, the output live projector has probabilities $p_x \geq 1 - \varepsilon$ and $p_y \leq e^{-1} + \varepsilon$. Cauchy–Schwarz in its two measurement sectors gives

$$\begin{aligned} |\langle \phi_x, \phi_y \rangle| &\leq \sqrt{p_x p_y} + \sqrt{(1 - p_x)(1 - p_y)} \\ &\leq \sqrt{e^{-1} + \varepsilon} + \sqrt{\varepsilon} < 3/4. \end{aligned} \tag{63}$$

This is valid for every simulator purification; no ideal environment alignment has been assumed.

Let S_L be the star adjacency with edges $0 \leftrightarrow j$, and $w = (1/\sqrt{2}, 1/\sqrt{2L}, \ldots, 1/\sqrt{2L})$. The entrywise nonnegative adversary

$$\Gamma_{\mathrm{adv}} = \sum_{b=1}^{r} I^{\otimes(b-1)} \otimes S_L \otimes I^{\otimes(r-b)} \tag{64}$$

has top eigenvector $w^{\otimes r}$ and norm $r\sqrt{L}$. The usual progress on the mathematical input-label lift starts at that value. For a membership address (b, j), the difference-masked adversary is a single edge on factor b with norm one. A phase query changes progress by at most two, also in superposition over addresses. At the end, positivity and (63) bound progress by $3r\sqrt{L}/4$. This proves the claimed $r\sqrt{L}/8$ bound. This is the standard adversary method[4], applied to a reduced-output guarantee.

Finally $L \geq a/r$, so $r\sqrt{L}/8 \geq \sqrt{ar}/8$. The supplied normalization is a budget; the factor may have smaller actual norm. In the all-marked instance, the actual loss action is $T \, \|G_x\| \, /2 = ar/M$, of order r, whereas the supplied budget is a. The lower bound therefore concerns this raw-factor oracle contract, not a claim about an independently supplied optimal normalization of the physical loss operator. The declared Hamiltonian upper budget can likewise be positive while its signal in this family is zero. □

Precision and Hamiltonian contributions

For the stationary cap-one family

$$L_\theta = \alpha \cos\theta |\mathrm{sink}\rangle\langle \mathrm{live}|, \qquad H_\theta = 0, \tag{65}$$

live survival probability is $e^{-2a\cos^2\theta}$. A raw encoding is one unknown two-dimensional rotation R_θ tensored with a known encoding of the lowering operator. All query entries have Laurent degree one in $e^{i\theta}$. A probability after Q queries has trigonometric degree at most $2Q$. Averaging it at $\theta, -\theta, \theta + \pi$ and $-\theta + \pi$ leaves a polynomial of degree at most Q in $y = \cos^2\theta$, with no larger approximation error. The established exponential-approximation theorem[5] Theorem 2 therefore with interval length $2a$ and uniform error $O(\varepsilon)$ gives, for $a \geq 1$,

$$Q = \Omega(d_\lambda(a)), \quad d_\lambda(a) = \sqrt{a\lambda} + \frac{\lambda}{\log(e + \lambda/a)}. \tag{66}$$

For an operator block use its live matrix element, or its unconditioned good/live probability. The same argument applies. This is a transfer of approximation degree through a legal factor oracle, not a new approximation theorem.

For the Hamiltonian family use $H_\theta = \beta \cos\theta Z$, no jumps, and the known input $|+\rangle$. Twice the output coherence is $e^{-2ih\cos\theta}$. It remains observable in an ordinary channel, so a global phase convention is unnecessary. Symmetrization at $\theta, -\theta$ gives a complex polynomial of degree at most $2Q$ approximating e^{-2ihx} on $[-1, 1]$.

Here is a direct precision estimate. If a polynomial has degree less than n and approximates $e^{-i\tau x}$ to error $\delta < 1$, its n-th forward difference on an interval of length $0 < \ell \le 2$ gives

$$\delta \ge \left|\sin \frac{\tau \ell}{2n}\right|^n . \tag{67}$$

Choosing $\ell = \pi n/\tau$ when allowed proves $n > 2\tau/\pi$. Choosing $\ell = \min(2, n/\tau)$ gives

$$\log(1/\delta) \le n \log\bigl(4 \max\{1, n/\tau\}\bigr). \tag{68}$$

Splitting $n \ge \log(1/\delta)$ and its complement yields $n = \Omega(\Phi_{\log(1/\delta)}(\tau))$. With $\tau = 2h$, $\delta = 4\varepsilon$ and $n = 2Q + 1$, this proves (the additive one is absorbed because a zero-query algorithm cannot approximate the entire $h \ge 1$ family at this accuracy)

$$Q = \Omega(\Phi_\lambda(h)), \qquad h \ge 1. \tag{69}$$

The same coherence on the designated good block proves the propagator version without postselection. These are familiar polynomial-method bounds; the direct derivation fixes their joint parameter dependence for the present contract.

Combining the bounds and saturating a loose cap

The function Φ_λ is increasing and $\Phi_\lambda(2x) \le 2\Phi_\lambda(x)$. Put $u = \sqrt{ar}$, $v = \sqrt{a\lambda}$. The two logarithms $\log(e + \sqrt{z})$ and $\log(e + z)$ are within a factor two, so $\Phi_\lambda(v) = \Theta(d_\lambda(a))$. If $\lambda \le r \le a$, then $v \le u$ and $u \ge \lambda$, hence $\Phi_\lambda(u) \le 2u$. If $\lambda \ge r$, then $u \le v$. Consequently, for $a \ge r \ge 1$,

$$\Phi_\lambda\left(h + \sqrt{a(r+\lambda)}\right) = \Theta\bigl(\max\{\sqrt{ar}, d_\lambda(a), \Phi_\lambda(h)\ [h \ge 1]\}\bigr). \tag{70}$$

For $h < 1$ the dissipative terms already dominate it. Different hard subclasses under the same supplied upper budgets suffice for this worst-case maximum; their lower bounds have not been multiplied. This proves the lower part of Theorem 2, and also the lower part of Theorem 10 after omitting u.

There is a useful extension to every cap $r \ge 1$. The uniform total rate $G_t \le \alpha^2 I$ implies, for $z \ge 1$,

$$\mathbb{E}_\rho z^{N_T} \le \exp(2a(z-1)). \tag{71}$$

Indeed, differentiating the weighted positive trajectory sum gives at most $\alpha^2(z-1)$ times its trace. This proof is uniform over inputs and references and uses no actual-rate oracle. Taking $z = 2$, the cutoff $r_c = \lceil 4(a+\lambda) \rceil$ gives $\Pr(N_T > r_c) \le e^{-2\lambda}$. The algorithm may replace r by $r_{\text{eff}} = \min(r, r_c)$, with the channel tail cost already paid in (56).

Supplementary Corollary 12 (All-cap query characterization)**.** *For hard-capped channels under the stated access and regularity assumptions, every $a \ge 1$, integer $r \ge 1$, $h \ge 0$, and $0 < \varepsilon \le 1/100$ have worst-case query complexity*

$$\Theta\left(\Phi_\lambda\left(h + \min\{a, \sqrt{a(r+\lambda)}\}\right)\right). \tag{72}$$

An implementation uses at most $O(\min(r, a+\lambda))$ record slots, with the routing and precision costs of Supplementary Note 5.

Proof. The original compiler gives $O(\Phi_\lambda(h + \sqrt{a(r+\lambda)}))$. The rate cutoff above gives $O(\Phi_\lambda(h + \sqrt{a(a+\lambda)}))$, up to constants. If $\lambda \le a$, the latter is $O(\Phi_\lambda(h+a))$. If $\lambda \ge a$, the inequality $\log(e + \lambda/a) \le 2\sqrt{\lambda/a}$ implies $\sqrt{a\lambda} \le 2\lambda/\log(e+\lambda/a)$; equivalently the same comparison

follows from (70) at the real value $r = a$. It again gives $O(\Phi_\lambda(h+a))$. Selecting the smaller of the two certified cutoffs proves the upper bound.

For $r \le a$, (70) supplies the lower bound. For $r > a$, use the event family with cap $\lfloor a \rfloor \ge a/2$, which is admissible in the at-most-r class and gives $\Omega(a)$. Together with (66) and (69), this gives $\Omega(\Phi_\lambda(h+a))$. The last comparison follows from $d_\lambda(a) = O(\Phi_\lambda(a))$ and $\Phi_\lambda(a) \le a + d_\lambda(a)$. This proves (72). □

This saturation recovers the known general action-based order when the cap is uninformative. It is a consequence of the event-sensitive construction and a rate cutoff, not a further speedup for unrestricted channels. Sub-unit actions, zero-event promises, fixed system dimension and exact constraints on integrated physical norms are outside this lower-bound characterization.

Supplementary Note 9. Activity certificates and physical examples

Absorbing channels and grade-lowering loss

Split the system into a transient sector and a sink, and supply arbitrary rectangular factors $F_\mu(t)$ from the former to the latter. Set

$$L_\mu^{\mathrm{tot}}(t) = \begin{pmatrix} 0 & 0 \\ F_\mu(t) & 0 \end{pmatrix}, \qquad H^{\mathrm{tot}}(t) = \mathrm{diag}(H(t), 0). \tag{73}$$

After one jump the state lies in the sink, where both the Hamiltonian and all jumps vanish. Thus every two-jump word is zero, also on superpositions of the sectors. Write $K' = -(\sum_\mu F_\mu^\dagger F_\mu/2 + iH)K$. The channel acts on a block operator as

$$\begin{pmatrix} K(T)XK(T)^\dagger & K(T)Y \\ ZK(T)^\dagger & W + \displaystyle\int_0^T \sum_\mu F_\mu(t)K(t)XK(t)^\dagger F_\mu(t)^\dagger\, dt \end{pmatrix}. \tag{74}$$

Here the input is $\left(\begin{smallmatrix} X & Y \\ Z & W \end{smallmatrix}\right)$. The off-diagonal blocks show that the chronological propagator survives in system coherences after tracing out the environment. With the supplied rectangular factors, the embedding costs a constant number of queries. Obtaining a block encoding from an arbitrary channel oracle would require a further reduction beyond (74).

For an explicit noncommuting example, take $F(t) = \sqrt{\gamma}|s\rangle\langle v(t)|$ and $H = 0$, with a live space of dimension at least two and an absorbing sink s. When v takes values $|0\rangle$ and $|+\rangle$, the two dissipative parts $D = F^\dagger F/2$ satisfy

$$[D_0, D_+] = \frac{\gamma^2}{8}\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \neq 0. \tag{75}$$

Smooth oracle-controlled interpolation gives a nonunital, time-dependent example outside the stationary Hermitian-jump class. Its input includes neither a diagonalizing frame nor a no-jump propagator. This example identifies the covered class; a fixed known two-dimensional instance carries no lower-bound claim.

Supplementary Proposition 13 (Loss filtration)**.** *Let $N = \sum_{n=0}^r nP_n$. Suppose $H(t)$ and G_t commute with N, and every jump lowers its grade filtration by at least one. Then the process has at most r jumps on every input. A sufficient special case is $[N, H(t)] = 0$ and $[N, L_\mu(t)] = -k_\mu L_\mu(t)$, with integers $k_\mu \ge 1$.*

Proof. The no-event generator preserves every grade. A chronological word with j jumps maps grade at most n to grade at most $n - j$, so it vanishes for $j > r$. In the special case, $[N, L_\mu^\dagger L_\mu] = 0$, proving the required condition on G_t. The argument is an operator support statement and includes entangled references. □

Interactions within each grade may be arbitrary and number preserving; collective lowering operators may vary in time. The grade decomposition must be supplied. Knowing only that a jump operator is nilpotent is insufficient: for

$$L = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \qquad H = 0, \tag{76}$$

one has $L^3 = 0$, but $L^\dagger L$ mixes the two upper grades and

$$L^2 e^{-sL^\dagger L/2} L|2\rangle = -\frac{2}{\sqrt{5}} e^{-3s/4} \sinh(\sqrt{5}s/4)|0\rangle \neq 0 \tag{77}$$

for $s > 0$. Continuity then gives nonzero three-jump amplitude on a positive-volume set of strictly ordered times. The intervening no-event evolution has defeated the nilpotency of the jump operator.

Weak repumping without a hard cap

Bound the total jump count by combining the pump count with the grade budget. Let N have nonnegative integer grades on a finite-dimensional payload, and let r_0, b be known nonnegative integers. The initial state has support on grades at most r_0. Separate downward jumps L_μ from pumps P_ν, and assume

$$\begin{aligned} [N, H(t)] &= 0, & [N, L_\mu(t)] &= -k_\mu L_\mu(t), \\ k_\mu &\geq 1, & [N, P_\nu(t)] &= \ell_\nu P_\nu(t), \\ 0 \leq \ell_\nu &\leq b, & \sum_\nu P_\nu(t)^\dagger P_\nu(t) &\leq p(t) I. \end{aligned} \tag{78}$$

All charges k_μ, ℓ_ν are integers. Neutral jumps count as pumps and consume the same rate budget. Supply the deterministic bound $\rho \geq \int_0^T p(t)dt$, together with a factor stack for both downward jumps and pumps satisfying the earlier access, normalization and regularity assumptions. The bound ρ is part of the input certificate; its efficient recovery from a general black box is not assumed.

Supplementary Theorem 14 (Certified weak-pump simulation)**.** *Under* (78)*, let U_T count pumps and N_T all events. Uniformly on the promised input space and references,*

$$\begin{aligned} N_T &\leq r_0 + (b+1)U_T, \\ \Pr(U_T \geq n) &\leq \Pr(\mathrm{Poisson}(\rho) \geq n). \end{aligned} \tag{79}$$

Choose an integer $m \geq 0$ with $\Pr(\mathrm{Poisson}(\rho) > m) \leq \varepsilon/8$, and put $r = r_0 + (b+1)m$. Then Theorem 2 applies with this r. For $\rho > 0$, a sufficient elementary choice has

$$m = \left\lceil 8\left(\rho + \frac{\Lambda}{\log(e + \Lambda/\rho)}\right)\right\rceil - 1, \quad \Lambda = \log(8/\varepsilon). \tag{80}$$

For $\rho = 0$, take $m = 0$. At fixed accuracy and bounded b, $r_0 + \rho = O(1)$ gives $O(1 + A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}}})$ raw queries even when the number of possible physical jumps is unbounded.

Proof. Track the grade along a nonzero trajectory word. The effect of each charge-homogeneous jump commutes with N, so no-event evolution leaves the grade unchanged. A word with D downward and U pump events, starting at grade $q \leq r_0$, therefore satisfies

$$D \leq \sum k_\mu \leq q + \sum \ell_\nu \leq r_0 + bU. \tag{81}$$

Adding the pump count gives the first inequality in (79). The support argument also covers superpositions, without an initial grade measurement.

To bound the pump count, let $\sigma_n(t)$ collect the positive unnormalized states of trajectories with exactly n pumps, summed over all downward events. The count blocks obey

$$\dot{\sigma}_n = \mathcal{L}_t(\sigma_n) - \mathcal{J}_{\mathrm{up},t}(\sigma_n) + \mathcal{J}_{\mathrm{up},t}(\sigma_{n-1}). \tag{82}$$

The count-block mass is $a_n = \operatorname{Tr}\sigma_n$, and its outgoing pump flux is $f_n = \operatorname{Tr}(G_{\mathrm{up}}\sigma_n)$. Thus $\dot{a}_n = f_{n-1} - f_n$ and $0 \le f_n \le p(t)a_n$. For a threshold n_*, compare this flux with the backward tail probability of a Poisson process:

$$u(t,n) = \Pr\left(n + \mathrm{Poisson}\left(\int_t^T p(s)ds\right) \ge n_*\right). \tag{83}$$

Increasing the count increases the tail probability, so $\Delta u(t,n) = u(t,n+1) - u(t,n)$ is nonnegative. The backward equation is $\partial_t u = -p(t)\Delta u$. Combining it with the count flux gives

$$\frac{d}{dt}\sum_n u(t,n)a_n(t) = \sum_n \Delta u(t,n)[f_n - p(t)a_n] \le 0. \tag{84}$$

Only counts below the threshold contribute nonzero increments. At the final time the test function is the threshold indicator; at the initial time its expectation is the Poisson tail. The inequality therefore proves stochastic domination. A reference can be included in every count block.

It remains to justify the explicit cutoff. The Poisson Chernoff inequality is $\Pr(U_T \ge n) \le e^{-I_\rho(n)}$ for $n > \rho$, with $I_\rho(n) = n\log(n/\rho) - n + \rho$. We bound its unique root above ρ at value Λ by $8[\rho + \Lambda/\log(e + \Lambda/\rho)]$. Set $z = \Lambda/\rho$. If $z \le 1$, the root is at most $e\rho$. If $z \ge 1$, set $u = \log(e+z)$ and $y = 8z/u$. The inequalities $z \ge e^u/(e+1)$ and $\log(7u/8) \le 7u/8 - 1$ give

$$\log y - 1 - u/8 \ge \log(7/(e+1)) > 0. \tag{85}$$

It follows that $y(\log y - 1) \ge yu/8 = z$, as required for the root bound. The integer threshold in (80) therefore controls $\Pr(U_T > m)$. Since $N_T > r \Rightarrow U_T \ge m+1$, this also controls the total-event tail. The error bound (56) completes the simulation claim. □

For a smaller cutoff one can compute an exact Poisson quantile, paying its scalar preprocessing cost. A certified upward-rounded Chernoff threshold is another option. The comparison in Supplementary Note 6 imposes the stop before perturbing the signal data; those perturbations may therefore violate the grading or pump certificate. The theorem gives an upper bound in ρ and b, without matching lower bounds in these parameters.

A spin example uses $N = \sum_j |1\rangle\langle 1|_j$, arbitrary number-preserving interactions, collective loss, and local pumps $P_j(t) = \sqrt{\gamma_j(t)}\,\sigma_j^+$. Each pump has $b = 1$; a valid total-rate certificate is $\rho \ge \int_0^T \sum_j \gamma_j(t)dt$, with the sum taken over every site. This choice excludes coherent raising terms in H and neutral jumps outside the pump budget. An infinite bosonic ladder would need domain control and a truncation whose cost is included.

A positive-operator interface for further classes

One can also certify the event tail by bounding an exponential moment with a positive observable. Use the counting-field generator[6]

$$\mathcal{L}_{\theta,t}(X) = -i[H_t, X] + e^\theta \mathcal{J}_t(X) - \tfrac{1}{2}\{G_t, X\}. \tag{86}$$

The resulting evolution is completely positive, with trace $\mathbb{E}e^{\theta N_T}$. Supply a positive observable satisfying

$$V_\theta \ge I, \qquad \mathcal{L}^*_{\theta,t}(V_\theta) \le c_\theta(t)V_\theta, \qquad P_0 V_\theta P_0 \le v_\theta P_0. \tag{87}$$

Gronwall's inequality for the expectation of this observable gives

$$\sup_{\operatorname{supp}\rho\subseteq P_0} \Pr_{\rho}(N_T \geq n) \leq \inf_{\theta>0} v_\theta \exp\left(\int_0^T c_\theta(t)dt - \theta n\right). \tag{88}$$

To see this, differentiate $\operatorname{Tr}(V_\theta \rho_\theta(t))$ and apply (87). Positivity permits Gronwall's inequality, while $V_\theta \geq I$ bounds the trace needed for Markov's inequality. Tensoring with an arbitrary reference preserves each step.

The grade certificate (78) corresponds to $V_\theta = e^{\theta N}$. Its downward-jump coefficients $e^{\theta(1-k_\mu)} - 1 \leq 0$ can be dropped in an upper bound, and the pump contribution is at most $p(t)(e^{(b+1)\theta} - 1)V_\theta$. Hence

$$\mathbb{E}e^{\theta N_T} \leq \exp\Big(\theta r_0 + \rho(e^{(b+1)\theta} - 1)\Big), \tag{89}$$

in agreement with (79). The operator certificate must be constructed separately. Only its tail bound is passed to the compiler; neither V_θ nor its square root is prepared.

Supplementary References

[1] Combettes, P. L. & Yamada, I. Compositions and convex combinations of averaged nonexpansive operators. *J. Math. Anal. Appl.* **425**, 55–70 (2015). https://doi.org/10.1016/j.jmaa.2014.11.044.

[2] Chen, B., Gao, M., Wang, X. & Zhou, S. Time-dependent Hamiltonian simulation with optimal query complexity. Preprint at https://arxiv.org/abs/2608.06094v1 (2026).

[3] Chen, B., Gao, M., Ji, Z., Li, T., Wang, X. & Zhou, S. Gate-efficient implementation of the query-optimal time-dependent Hamiltonian simulation. Preprint at https://arxiv.org/abs/2608.30629v1 (2026).

[4] Lee, T., Mittal, R., Reichardt, B. W., Špalek, R. & Szegedy, M. Quantum query complexity of state conversion. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, 344–353 (2011). https://doi.org/10.1109/FOCS.2011.75.

[5] Aggarwal, A. & Alman, J. Optimal-degree polynomial approximations for exponentials and Gaussian kernel density estimation. In *37th Computational Complexity Conference*, LIPIcs **234**, 22:1–22:23 (2022). https://doi.org/10.4230/LIPIcs.CCC.2022.22.

[6] Garrahan, J. P. & Lesanovsky, I. Thermodynamics of quantum jump trajectories. *Phys. Rev. Lett.* **104**, 160601 (2010). https://doi.org/10.1103/PhysRevLett.104.160601.