Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantum simulation of open systems with bounded jump activity

Jacob Kitchen

jacob.kitchen@hotmail.com

Unaffiliated https://orcid.org/0009-0007-7319-0238

Article

Keywords:

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11628481/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Quantum simulation of open systems with bounded jump activity

Jacob Kitchen
Unaffiliated
Corresponding author: jacob.kitchen@hotmail.com

Abstract. Open quantum systems can accumulate a large dissipative action while undergoing few quantum jumps. General simulation bounds based on generator strength and evolution time do not distinguish such dynamics from sustained activity. Here we show that a supplied activity bound reduces the quantum query cost without requiring the intervening evolution to commute. With coherent access to Hamiltonian and jump factors, fixed accuracy and bounded activity give square root dependence on dissipative normalization action. Matching lower bounds hold for the reduced channel under the same oracle model. The construction extends to certified jump tails and gives a normalization-one encoding of a noncommuting accretive propagator without an activity assumption. Explicit bounds account for additional circuit, precision and discretization costs. These results identify jump activity as a resource for simulating absorption, finite excitation loss and weakly repumped dynamics.

Introduction

Quantum simulation of open systems must account for both coherent evolution and irreversible events. A Lindblad equation describes these processes through a Hamiltonian and a collection of jump operators. Recent quantum algorithms attain optimal query complexity for general Lindblad dynamics, with a leading cost linear in a normalization action that combines Hamiltonian strength, dissipation strength and evolution time[1–3]. This action is an appropriate general resource measure, but it does not distinguish repeated dissipative activity from evolution that exhausts a finite supply of excitations.

An absorbing system illustrates the distinction. After one decay into an inert sink, it cannot jump again, however long the evolution continues. A system initially containing finitely many excitations can undergo several losses and then reach the same situation. Before each event, its coherent and dissipative evolution may still vary in time and fail to commute. A useful simulation guarantee should be able to exploit an event bound without requiring the intervening propagator to be explicitly solvable.

There are several precedents for reducing the dependence on dissipative evolution. A single Hermitian jump and Choi-commuting dissipators admit fast-forwarding constructions[4]. Solvable and nilpotent operator algebras can yield finite Kraus expansions[5;6]. Quantum trajectory algorithms obtain favourable precision dependence for stationary generators whose total loss operator is proportional to the identity[7]. These results rely on different structural and access assumptions. In particular, a finite expansion does not by itself provide an efficient oracle implementation of its chronological no-jump propagators.

The closely related problem of non-unitary propagation also depends on the input representation. Linear combinations of Hamiltonian simulation provide a general approach to accretive dynamics. Low and Somma identify access to the square root of the dissipative part as a route to a cost that is additive in square root dissipative action and Hamiltonian action[8]. Amplitude–phase separation and implementations of the associated dissipative families give further constructions[9;10]. Comparisons between these methods must distinguish access to a factor from access to its square, and an operator encoding from preparation of a normalized output state.

Here we show that jump activity controls the query cost of Lindblad simulation under coherent access to the Hamiltonian and jump factors. A weighted causal query compression construction gives a square root dependence on dissipative normalization action when the event count and accuracy are fixed. We prove matching worst case bounds for the reduced channel in the stated oracle model, recover the general action dependence when the event bound is uninformative, and extend the upper bound to certified event tails. The same construction gives a normalization-one encoding of a time-dependent accretive propagator without an activity promise. Absorbing dynamics, finite excitation loss and weak repumping supply concrete activity certificates.

Results

Activity and coherent factor access

Consider a finite-dimensional system obeying

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

The supplied jump map specifies the counting process. Let N_T denote its number of jumps before time T. An activity certificate is either a hard cap, $N_T \le r$ on every input, or a uniform tail bound $\sup_\rho \Pr_\rho(N_T > r) \le p_{\mathrm{tail}}$. A promise can instead hold on a specified input subspace, in which case the simulation guarantee is restricted to that subspace and arbitrary entangled references.

The input provides unitary encodings of the rectangular stack $L(t) = \sum_\mu |\mu\rangle \otimes L_\mu(t)$, divided by a known α, and of $H(t)$, divided by a known β. Thus $\|L(t)\| \le \alpha$ and $\|H(t)\| \le \beta$. Controlled calls, adjoints and coherent time selection on a chosen dyadic grid are available. A call to either selected unitary or its adjoint is one raw query. The two action parameters are

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

These are budgets determined by the supplied oracle normalizations. They need not equal integrated physical norms. Constructing or improving an encoding, loading data and implementing the coherent selectors are separate costs. In particular, a classical list of jump encodings is not already a coherent encoding of their stack.

For continuum channel simulation, assume known regularity bounds

$$\begin{aligned}\|H(t) - H(s)\| &\le K_H |t-s|,\\ \|\mathcal{J}_t - \mathcal{J}_s\|_\diamond &\le K_J |t-s|.\end{aligned} \tag{3}$$

An effective modulus of continuity can replace these Lipschitz bounds. Regularity is required of the physical jump map, rather than of a particular choice of Kraus phases or label basis. Discontinuous changes of that basis are therefore allowed. The mesh affects address width and arithmetic cost; it does not require one sequential query per time step. Formal input and output definitions appear in Supplementary Note 1.

Channel complexity at bounded activity

Write $\lambda = \log(C/\varepsilon)$, where C is a universal constant, and define

$$\begin{aligned}\Phi_\lambda(x) &= x + \frac{\lambda}{\log(e + \lambda/x)} \quad (x > 0),\\ \Phi_\lambda(0) &= 0.\end{aligned} \tag{4}$$

For $0 < \varepsilon < 1/8$, an integer cutoff $r \ge 0$, and $p_{\mathrm{tail}} \le \varepsilon/8$, the channel generated by Equation (1) can be simulated with

$$\begin{aligned}Q &= O(1 + \Phi_\lambda(B)),\\ B &= A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}}(r+\lambda)}.\end{aligned} \tag{5}$$

The error is at most ε in diamond norm, without the conventional factor of one half. The algorithm is trace preserving, accepts one unknown system input, and discards its finite compiler registers at the end. It uses neither postselection nor a reflection about the input state. An input entangled with an arbitrary reference obeys the same guarantee.

For hard caps, $A_{\mathrm{D}} \ge r \ge 1$, and $\varepsilon \le 1/100$, the dependence in Equation (5) is optimal up to constants in the worst case over admissible raw-factor oracle instances. This statement allows both system dimension and oracle registers, including the jump-label register, to grow. At fixed accuracy it gives $\Theta(A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}} r})$ queries. The lower bound applies to the ordinary output channel after all environmental records have been discarded. It is not a lower bound on the cost of printing those records.

A loose cap need not increase the cost above the general action-based order. The total rate bound $G_t \le \alpha^2 I$ supplies an alternative cutoff of order $A_{\mathrm{D}} + \lambda$. For $A_{\mathrm{D}} \ge 1$, any integer hard cap $r \ge 1$, and $\varepsilon \le 1/100$, the complete worst case characterization is

$$\begin{aligned}Q &= \Theta\big(\Phi_\lambda(B_{\mathrm{cap}})\big),\\ B_{\mathrm{cap}} &= A_{\mathrm{H}} + \min\left\{A_{\mathrm{D}}, \sqrt{A_{\mathrm{D}}(r+\lambda)}\right\}.\end{aligned} \tag{6}$$

Zero total action is the identity and requires no queries. The optimality statement does not cover sub-unit dissipative budgets, a zero-event promise, fixed system dimension, or exact constraints on integrated physical norms. The upper bound in Equation (5) has the broader domain stated above.

Figure 1 shows how the explicit cost function changes with the event cap. It evaluates the asymptotic expression with its unknown constants omitted; it is not a measured query count or runtime comparison. The formal theorems and proofs are in Supplementary Note 1 and Supplementary Note 8.

How the construction uses activity

Causal query compression converts a unitary transducer into an ordinary quantum circuit. The transducer formalism separates a public input from private states used to certify its action[11]. Reuse circuits, polynomial filters and compact time routing provide the underlying compression machinery[1–3;12;13]. The present change is a factor-driven network with an adjustable positive weight q. It changes the balance between carry and bath ports while leaving the physical short-time map unchanged.

The network keeps a finite event history. It permits the first $r+1$ events and freezes the state after the over-

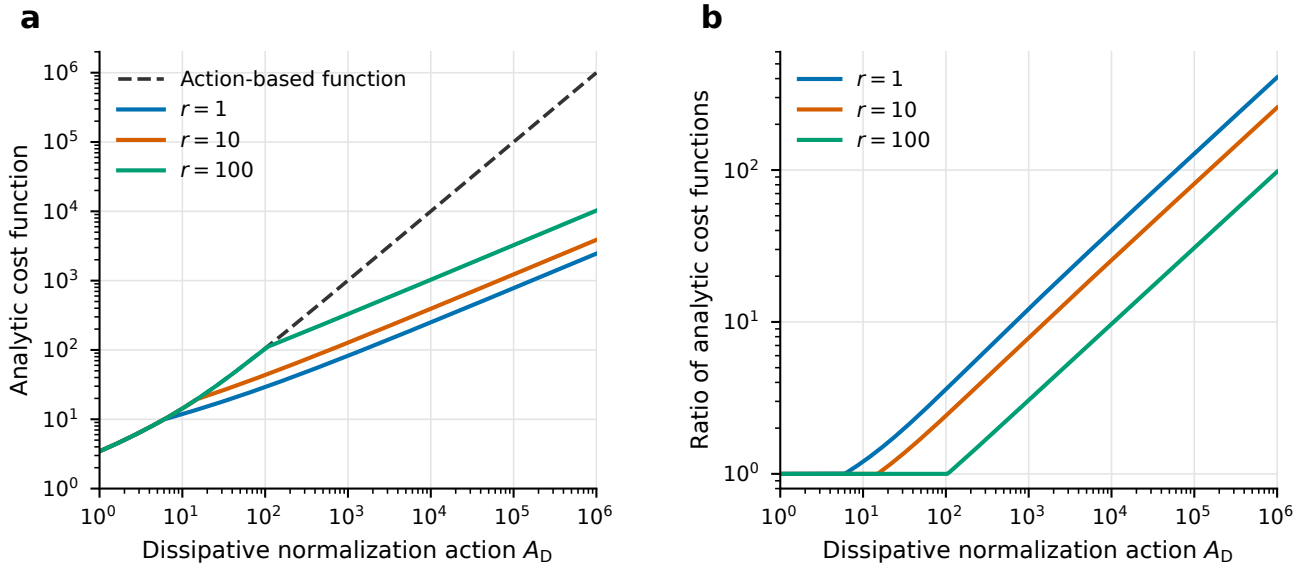


Figure 1. Activity dependence of the query bound. a Explicit bound functions for zero Hamiltonian action and $\lambda = 5$. The solid curves evaluate $\Phi_5(\min\{a, \sqrt{a(r+5)}\})$, with $a = A_D$; the dashed curve is $\Phi_5(a)$. **b** Ratio of the action-based function to each activity-dependent function. The curves coincide until $a = r + 5$, after which the activity bound has square root leading dependence on a. The ratios concern these analytic functions only. Unknown asymptotic constants, data loading, history routing, gate synthesis and measurement costs are omitted. No measured algorithmic speedup or hardware performance is represented. Source values are in Supplementary Data 1.

flow event. In particular, it does not stop immediately after event r, because the following no-jump evolution remains part of the desired channel. If Γ denotes its private witness map, the energy bound is

$$\|\Gamma\|^2 \le \frac{A_D}{q} + A_H + q(r+2). \tag{7}$$

This is an operator bound on the entire public input space, including occupied bath inputs. It is stronger than a bound proved only for the initialized root history. That distinction is needed when the compiler amplifies the full public unitary before selecting a column or discarding records.

Energy alone does not control the dependence on precision. The private transfer also obeys a norm bound inside a disk after a diagonal similarity in time and history depth. This estimate controls the generally nonnormal transfer directly. The similarity and witness are proof objects and are not prepared or implemented. Balancing $q \asymp \sqrt{A_D/(r+\lambda)}$, then retaining the exact contraction rate of the inherited filter $(1+z^2)/2$, gives Equation (5). The construction combines reuse lengths with bounded coefficients and allocates one final approximation budget. Supplementary Note 2–Supplementary Note 4 give the certificate and its proof; Supplementary Note 5 specifies the circuits.

Absorption and finite excitation loss

The cap-one class already includes noncommuting time dependence. Split the system into a live sector and an inert sink. A jump factor $F_\mu(t)$ maps the live sector into the sink, while the Hamiltonian acts only within the live sector. Every two-jump trajectory vanishes. On the live sector, however, the no-jump propagator remains the chronological solution generated by $\sum_\mu F_\mu(t)^\dagger F_\mu(t)/2 + iH(t)$. It also appears in coherences between the live sector and the sink after the environment is traced out. The reduced channel can therefore retain the difficulty of noncommuting propagation even though at most one event occurs.

For example, let $F(t) = \sqrt{\gamma}|s\rangle\langle v(t)|$, where s is an absorbing sink and $v(t)$ is a smoothly varying vector in a live space of dimension at least two. At times when v equals $|0\rangle$ and $|+\rangle$, the corresponding positive loss operators fail to commute. These channels are nonunital and lie outside a stationary Hermitian-jump construction. The example demonstrates the scope of the theorem; a fixed, explicitly known two-dimensional instance is not itself evidence of oracle hardness.

More generally, let $N = \sum_{n=0}^{r} nP_n$ be an excitation grade. If $H(t)$ and G_t preserve every grade and each jump lowers the grade filtration by at least one, then no trajectory can contain more than r jumps. A sufficient condition is $[N, H(t)] = 0$ and $[N, L_\mu(t)] = -k_\mu L_\mu(t)$, with integers $k_\mu \ge 1$. Arbitrary number-preserving interactions and time-dependent collective losses are allowed. A finite spin system with loss operators built from lowering operators supplies a direct example. The finite-dimensional domain and any promised input subspace must be specified as part of the problem.

A tail certificate extends the construction to weak repumping. Assume a grade-preserving Hamiltonian and jump operators of definite grade, as specified in Supplementary Note 9. Suppose the initial grade is at most r_0, each pump raises it by at most b, and the total pump rate has a known integral bound ρ. The number

of pumps is stochastically dominated by a Poisson variable with mean ρ. If its upper tail above m is at most $\varepsilon/8$, then $r = r_0 + (b+1)m$ is a valid event cutoff. For local spin pumps, $b = 1$ and an admissible ρ is the sum of their integrated rate bounds over every site. Thus bounded $r_0 + \rho$, bounded b and fixed accuracy give $O(1 + A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}}})$ raw queries even when arbitrarily many events have nonzero probability. Coherent raising terms and unbudgeted neutral jumps are excluded from this certificate. Supplementary Note 9 proves these statements and gives a more general positive-operator tail certificate.

Accretive propagation from factors

The no-event block yields a second output problem. Define

$$K(T) = \mathcal{T} \exp\left[- \int_0^T A(t)\, dt \right], \qquad A(t) = \tfrac{1}{2} L(t)^\dagger L(t) + iH(t). \tag{8}$$

Assume the factor access above and known Lipschitz regularity of $L(t)^\dagger L(t)/2 + iH(t)$. A unitary circuit has a designated block approximating $K(T)$, with normalization one and operator error ε, using

$$O\Big(1 + \Phi_\lambda\Big(A_{\mathrm{H}} + \sqrt{A_{\mathrm{D}}\lambda}\Big)\Big) \tag{9}$$

raw queries. No event cap, spectral gap, mixing-time estimate or initial-state overlap promise is needed. For $A_{\mathrm{D}} \geq 1$ and $\varepsilon \leq 1/100$, this order is worst case optimal under the same normalization-budget convention.

The compiler obtains this result from the depth-zero stopped network. It first amplifies a unitary on the whole public space and only then selects the root carry block. Applying amplification directly to the nonunitary $K(T)$ would not justify the result. Preparing $K(T)|\psi\rangle / \|K(T)|\psi\rangle\|$ is an additional task: the designated block succeeds with probability approximately $\|K(T)|\psi\rangle\|^2$, and state preparation and amplification must be charged. The full proof and approximation bounds are in Supplementary Note 7.

Comparison with existing algorithms

Table 1 separates the closest constructions by their output and assumptions. For general Lindblad simulation, the action-based query order is recovered by Equation (6) when the cap is loose. The improvement here is in the dependence on a certified activity bound; it is not a universal improvement for unrestricted channels. The inherited filter, reuse identities and scalar-clock circuit remain the ones used in earlier causal compression algorithms[1–3;12;13].

The comparison with fast-forwardable dissipators is structural. Shang et al. establish square root scaling for a single Hermitian jump and Choi-commuting classes[4]. The present activity condition allows nonunital absorption and general noncommuting time dependence. The advantage here depends on a small event cutoff, whereas those constructions exploit algebraic structure and can apply when jump activity is sustained. Finite Kraus expansions for solvable jump algebras[5;6] can be more effective when their constituent operations are explicitly available. A bare event cap, or nilpotency of a jump operator alone, does not provide that access. Indeed, intervening no-jump evolution can defeat the implication from nilpotency to a finite event cap, as shown in Supplementary Note 9.

The propagator theorem addresses the square-root-access direction discussed by Low and Somma[8]: dissipative and Hamiltonian costs enter additively, including for noncommuting time dependence. Its precision term and output normalization are explicit in Equation (9). This does not make encodings of the dissipative operator and its factor interchangeable. Amplitude–phase separation[9] provides a related decomposition, while the root-oracle analysis of Wang et al. charges the implementation of the controlled dissipative family as well as the number of calls to that family[10]. Those costs must be combined before comparing raw queries. Finally, the present bound does not improve every precision regime of the restricted stationary trajectory algorithms[7].

Discussion

The distinction between normalization action and jump activity leads to a different resource measure for open-system simulation. A factor oracle can describe evolution over a long horizon while a supplied excitation or pump budget limits its activity. Under that certificate, the channel query cost interpolates between square root dependence on dissipative action and the general action-based order. The lower bound remains valid when the environment is discarded, and the no-event construction gives a parallel result for noncommuting accretive propagation.

These are worst case statements about a specified oracle interface. Their normalizations are supplied upper budgets rather than optimally normalized physical operators. In the event lower-bound family, for example, the physical loss action can be of order r while the supplied dissipative budget is much larger. Resolving small hidden amplitudes in that encoding is part of the query problem. Readable classical formulas or more informative input oracles can permit algorithms outside this lower-bound argument. The results therefore do not assert elementary-gate hardness for every Hamiltonian with an explicit formula.

Circuit cost also matters independently of query count. Constructing the factor stack, preparing labels, computing coherent addresses and scanning event records may dominate the implementation. Observable estimation adds its own sampling or estimation cost, and full tomography is outside the output guarantee. The resource bounds in Supplementary Note 5 make these contributions explicit. They support an assessment for a chosen physical model, without claiming a runtime advantage for every member of the class.

Table 1. Closest algorithms and the assumptions that differ. Here $a = A_{\mathrm{D}}$, $h = A_{\mathrm{H}}$, and $\lambda = \log(C/\varepsilon)$. Orders written in this notation are compared only where the access models can be matched. A primitive family call must be expanded into its own oracle cost before comparison. The statements summarize the cited constructions and do not exclude future refinements.

Construction	Output and assumptions	Relation to the present result
General Lindblad compression[1–3]	Channel simulation from coherent generator encodings; no activity promise required.	General action-based order $O(1+\Phi_\lambda(h+a))$, after matching normalizations. The present bound adds a dependence on certified activity.
Fast-forwardable dissipators[4]	A single Hermitian jump or Choi-commuting dissipators; square root scaling in Hamiltonian simulation time for the covered families.	A direct scaling precedent with different hypotheses. The present class includes nonunital absorption and noncommuting time dependence.
Finite Kraus constructions[5;6]	Solvable or nilpotent algebraic structure, with implementable terms in the resulting expansion.	A finite expansion alone does not supply the chronological no-jump operations from an unknown raw factor. Explicitly solvable cases can be cheaper.
Stationary trajectory algorithms[7]	For total loss proportional to the identity, trajectory methods attain additive precision dependence in jump-operator queries.	The present result allows different loss structure and time dependence; it does not improve every restricted precision regime.
Linear non-unitary simulation[8]	Accretive propagation; square-root access is identified as a route to additive dissipative and Hamiltonian costs.	The no-event theorem realizes that dependence with explicit precision cost for noncommuting time dependence. Normalized-state output costs remain separate.
Amplitude–phase and root-oracle constructions[9;10]	Stationary or piecewise constructions and explicit implementations of controlled dissipative families.	Count both the family invocations and their primitive factor-query cost. The present theorem treats the time-dependent factor contract directly.

The activity certificates identify useful settings in which to make that assessment. Absorption and finite excitation loss require no algebraic solution of the no-jump dynamics, while weak incoherent pumping replaces an exact cap with a controlled tail. Extending these certificates to coherent pumping requires a bound on grade changes during no-event evolution. Matching lower bounds for soft-tail parameters, characterizations at fixed dimension and more efficient record routing remain open. These questions concern the range and implementation of the activity-dependent theorem, rather than its reduction to a particular numerical example.

Methods

Transducer and finite circuit construction

For a cell of duration δ, put $A = L^\dagger L/2 + iH$ and $R_\delta = (I + \delta A/2)^{-1}$. The rational carry and jump operators are

$$M_\delta = 2R_\delta - I, \quad J_\delta = \sqrt{\delta}\, L R_\delta, \qquad M_\delta^\dagger M_\delta + J_\delta^\dagger J_\delta = I. \tag{10}$$

Known rotations and a Hermitian involution built from the raw factor and Hamiltonian encodings realize a Cayley transducer for this cell. The resolvent identifies its target analytically; the circuit does not assume an inverse oracle. The weight q appears in the carry and bath couplings but cancels from the physical cell. All oracle garbage is retained.

Compose the cells on a public direct sum indexed by valid event histories. The chronological product is unitary on the entire public space. Stack the local private witnesses to obtain the operator Γ in Equation (7). A balance of incoming and outgoing bath norms proves this estimate even for initially occupied history wires. The private compression is then analysed in bright coordinates as a direct sum of scalar paths. Weighting a vertex at time j and depth n by $\exp(-sj/J - \nu n)$ yields a disk bound after similarity. The proof uses a quadratic-form inequality for the Cayley admittance, rather than a spectral-radius estimate.

The averaged filter is $f(z) = (1+z^2)/2$. A bounded signed combination of reuse averages gives an approximation $\widetilde{V}$ to the full public unitary V. The linear combination has a block encoding of normalization $3/2$. Attenuation makes the implemented block $\widetilde{V}/2$, after which one oblivious-amplification step produces a clean initialized isometry with error linear in $\left\|\widetilde{V} - V\right\|$. Reflections act on zero compiler registers tensored with the complete public space. No witness preparation or unknown input reflection is used. Supplementary Note 2 gives the exact certificate, including its normal-

ization, and Supplementary Note 3–Supplementary Note 5 specify its realization.

Stopping and error allocation

Stopping is imposed on the physical process before its coefficients are rounded. A counter retains events through $r+1$ and makes the overflow sector absorbing. The original and stopped channels share the same no-overflow completely positive map. Their overflow maps have the same effect operator, bounded by $p_{\mathrm{tail}} I$ on the promised input space. Their difference in diamond norm is therefore at most $2p_{\mathrm{tail}}$. This linear tail error uses the reduced channel; comparison of coherent trajectory amplitudes generally gives a square root probability loss.

Let J be the number of uniform cells and $\gamma = \beta + \alpha^2$. A one-cell channel comparison has error at most $4\gamma^2\delta^2$. Duhamel comparison of the frozen generators and telescoping of trace-preserving cell products yield a grid error at most $[K_H + K_J + 4\gamma^2]T^2/J$. If the sampled factor and Hamiltonian signal blocks have operator errors ζ_F and ζ_H, and the compiled isometry has error η, the total error is bounded by

$$\begin{aligned}\left\|\mathcal{E}_{T:0} - \widetilde{\mathcal{E}}\right\|_{\diamond} \le\ & 2p_{\mathrm{tail}} \\ & + \frac{[K_H + K_J + 4\gamma^2]T^2}{J} \\ & + 2A_{\mathrm{H}}\zeta_H + 8A_{\mathrm{D}}\zeta_F + 2\eta.\end{aligned} \tag{11}$$

Each contribution is allocated a fixed fraction of ε. The rounded signal data need not preserve the original grading: the counter already enforces stopping. Errors in the complete controlled oracle unitaries, including garbage blocks, have a separate hybrid budget of order ε/Q per call. Signal-block accuracy alone does not control them. Supplementary Note 6 gives all constants and mesh conditions.

Gates and memory beyond raw queries

The history register stores at most $r+1$ ordered event slots, each containing a time and jump label. It uses $O((r+1)\log[J(m+1)])$ qubits. A rate cutoff can reduce the slot count to $O(\min(r, A_{\mathrm{D}} + \lambda))$ in the all-cap regime, up to the overflow slot. Head-and-suffix routing implements the free scalar paths without iterating over all cells. Binomial coefficient states and uniform reuse intervals are prepared reversibly using logarithmic count registers and polynomial-precision arithmetic.

The resulting gate bound is the raw query count multiplied by the actual costs of coherent source calls, address and embedding operations, a linear scan of the stored event slots, and polynomial-precision arithmetic. The full expression, including finite input length, guard bits and synthesis precision, is given in Supplementary Note 5. No efficient loading circuit for an arbitrary table is assumed. These resource estimates establish a finite circuit implementation, without converting a query theorem into an unconditional elementary-gate speedup.

Lower-bound reductions

Three stationary families establish the optimal query order. A hidden membership oracle specifies which live states can decay into sinks. The reduced output distinguishes presence from absence of a mark with constant bias. This bounds overlaps of every possible simulator purification and permits a nonnegative adversary argument[14], giving $\Omega(\sqrt{A_{\mathrm{D}} r})$ queries. The construction uses growing system and jump-label registers.

A cap-one family has a factor proportional to $\cos\theta$. Its survival probability is $e^{-2A_{\mathrm{D}}\cos^2\theta}$. Symmetrizing the output probability of a query algorithm produces a polynomial approximation problem. The exponential approximation theorem of Aggarwal and Alman[15], with interval length $2A_{\mathrm{D}}$, supplies the dissipative precision bound. A diagonal Hamiltonian family supplies the Hamiltonian and precision contribution through an observable off-diagonal coherence. The maximum of these separate lower bounds matches Equation (5) in its stated regime; the lower bounds are not multiplied. A uniform count-rate estimate gives the saturation in Equation (6). Supplementary Note 8 contains all reductions.

Activity certificates and analytic illustration

The grade certificate follows by applying the no-jump propagator and each jump to grade filtrations. For weak pumps, the number of all events is at most the initial grade plus $b+1$ times the pump count. A positive count-block flux comparison bounds the latter by a Poisson process with the supplied rate integral. More general certificates use the counting-field generator[16] and a positive observable satisfying an operator differential inequality. Supplementary Note 9 proves the tail bounds and gives explicit cutoffs.

Figure 1 is a deterministic illustration of Equation (6). It sets $A_{\mathrm{H}} = 0$, $\lambda = 5$, and evaluates $\Phi_5(\min\{a, \sqrt{a(r+5)}\})$ on a logarithmic grid in $a = A_{\mathrm{D}}$. The reference curve is $\Phi_5(a)$. The plotted ratios compare these functions with their unspecified asymptotic constants set equal. They do not estimate a speedup factor for implemented algorithms, and $\lambda = 5$ is an illustrative parameter rather than a calibrated physical error tolerance.

Use of artificial intelligence tools

ChatGPT (OpenAI) was used to assist with prose editing and presentation. The mathematical ideas, results and proofs are the author's own. The author takes responsibility for the final manuscript.

Data availability

This work contains theoretical results and no experimental dataset. The statements and derivations are provided in the article and Supplementary Notes 1–9.

The analytic values underlying Figure 1 are provided in Supplementary Data 1.

Author contributions

J.K. is responsible for the theoretical results and the manuscript.

Corresponding author

Jacob Kitchen: jacob.kitchen@hotmail.com.

Funding

This research received no funding.

Competing interests

The author declares no competing interests.

References

[1] Chen, B., Gao, M., Wang, X. & Zhou, S. Query-optimal and gate-efficient Lindbladian simulation. Preprint at https://arxiv.org/abs/2609.18757v1 (2026).

[2] Kitchen, J. Causal query compression for Lindblad dynamics: optimal queries and nearly linear local simulation. Preprint at https://arxiv.org/abs/2609.31312v1 (2026).

[3] Wang, C. & Ye, C. Query-optimal quantum simulation of Lindblad evolution. Preprint at https://arxiv.org/abs/2609.17490v2 (2026).

[4] Shang, Z.-X., Guo, N., Rebentrost, P., Aspuru-Guzik, A., Li, T. & Zhao, Q. Fast-forwardable Lindbladians imply quantum phase estimation. Preprint at https://arxiv.org/abs/2510.06759v2 (2025).

[5] Burdine, C. & Blair, E. P. Trotterless simulation of open quantum systems for NISQ quantum devices. *Adv. Quantum Technol.* **8**, 2400240 (2025). https://doi.org/10.1002/qute.202400240.

[6] Burdine, C., Bauer, N., Siopsis, G. & Blair, E. P. Efficient simulation of open quantum systems on NISQ trapped-ion hardware. *Adv. Quantum Technol.* **8**, 2400606 (2025). https://doi.org/10.1002/qute.202400606.

[7] Borras, E. & Marvian, M. Quantum simulation algorithms based on quantum trajectories. *Quantum* **10**, 2063 (2026). https://doi.org/10.22331/q-2026-04-13-2063.

[8] Low, G. H. & Somma, R. D. Optimal quantum simulation of linear non-unitary dynamics. Preprint at https://arxiv.org/abs/2508.19238v2 (2025).

[9] Hu, Q. & Jin, S. Quantum simulation of non-unitary dynamics via amplitude-phase separation. Preprint at https://arxiv.org/abs/2602.09575v2 (2026).

[10] Wang, C., Zhuang, X.-N., Dou, M., Chen, Z.-Y. & Guo, G.-P. Poisson-compiled quantum singular value transformation for power-exponential dissipation. Preprint at https://arxiv.org/abs/2608.04263v1 (2026).

[11] Belovs, A., Jeffery, S. & Yolcu, D. Taming quantum time complexity. *Quantum* **8**, 1444 (2024). https://doi.org/10.22331/q-2024-08-23-1444.

[12] 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).

[13] 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).

[14] 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.

[15] 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.

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

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementaryreading.pdf
- supplementarydata1.csv