One-Clean-Qubit Simulation of Quantum Catalytic Logspace with Measurements

Abstract

We prove that polynomial-time quantum catalytic logspace computations with intermediate measurements and resets can be simulated using one clean qubit and otherwise maximally mixed qubits. The source may use logarithmic-space classical control and polynomially many catalytic qubits, whose arbitrary unknown state must be restored exactly, including its correlations with an inaccessible reference. The simulator has polynomial size, is deterministic-logspace-uniform, and is unitary until a final measurement with inverse-polynomial signed acceptance bias. Thus $\mathsf{BQCL} \subseteq \mathsf{DQC1}$ under this bias convention, extending the unitary containment of Buhrman et al. (TQC, 2025). The main difficulty is to avoid storing a polynomial-length measurement history in initialized memory. We express the source's acceptance bias as a normalized trace involving the inverse of a channel-history matrix. Although the history matrix can be poorly conditioned, complete positivity bounds the Frobenius norm of its inverse. A nuclear-norm regularization estimate then controls the required trace with no dependence on the exponential catalyst dimension. We implement the resulting bounded spectral filter by coherent phase estimation and extract its relevant block through a reflected unitary trace test. The logarithmic auxiliary space is included in the maximally mixed input at only an inverse-polynomial loss in bias.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenghua Liu and Hanyu Li.

1 Introduction

Catalytic space separates initialized work memory from borrowed memory whose unknown contents must be restored. In the classical setting, Buhrman et al. [1] showed that this additional memory can support computations not known to fit in the initialized workspace alone. Quantum catalytic space imposes a stronger restoration requirement: a borrowed register must be returned in its original state for every quantum input to that register. This also preserves its entanglement with any inaccessible reference. The model asks how much computation a large register can support without consuming the information it contains.

Buhrman et al. [2] introduced quantum catalytic space and proved that its unitary logspace computations admit one-clean-qubit simulations. They left open the corresponding containment when intermediate measurements are allowed. The usual deferred-measurement construction does not settle this question. It stores measurement outcomes coherently, requiring a fresh initialized register for a potentially polynomial-length history. Resets pose a related difficulty: unlike unitary conjugations, they need not be contractions in the Hilbert–Schmidt norm.

We prove the containment for the measured model. The simulator does not reproduce the source state or retain its measurement outcomes. Instead, it recovers the sign of the source's acceptance bias from a normalized trace. Exact restoration provides an identity on the entire catalyst operator space, not merely on its diagonal coordinates. This identity supplies the trace multiplicity needed to offset the size of the catalyst.

1.1 Model and result

We use the polynomial-length, logspace-uniform general-circuit formulation of quantum catalytic space in [2, Definitions 15, 18–24 and 30]. For an input x of length n, the source has an initialized work register W of $s = O(\log n)$ qubits and a catalytic register C of $c = \mathrm{poly}(n)$ qubits. Its live classical controller, including all retained measurement outcomes, counts toward the work-space bound. The computation has a worst-case polynomial time bound and may use computational-basis measurements, resets, and classically controlled constant-local gates. There is no uncharged classical transcript available for later queries. The input is classical and is available to the logspace circuit generator.

Writing Φ_x for the overall channel before the final output is read, we require a work state η_x such that, for every reference register R and every state ρ_{CR},

$$(\Phi_x \otimes \mathrm{id}_R)\big(|0^s\rangle\langle 0^s|_W \otimes \rho_{CR}\big) = \eta_x \otimes \rho_{CR}. \tag{1}$$

Lemma 2 shows that this product form follows already from exact restoration of the catalyst marginal for every catalyst state. In particular, it remains valid after the live controller is included in W. A designated work qubit is the output; its probability of being 1 is denoted by p_x. The bounded-error class BQCL uses $p_x \geq 2/3$ on yes instances and $p_x \leq 1/3$ on no instances; QCL is the exact version.

For $\mathsf{DQC1}$ we use the signed-bias convention of [2, Definitions 49–51]. A circuit starts with one qubit in $|0\rangle$ and polynomially many qubits in the maximally mixed state, applies a polynomial-size logspace-uniform unitary circuit, and measures the clean probe at the end. There is a polynomial q such that acceptance has probability at least $1/2 + 1/q(n)$ on yes instances and at most $1/2 - 1/q(n)$ on no instances. We do not assert constant bias within a single one-clean-qubit circuit.

We fix the same effective elementary-gate convention for the source and the simulator. Constant-local unitaries are described using phase and rotation gates and classical reversible gates; adjoints, complex conjugates, and controlled versions have logspace-generated descriptions. This closure requirement is part of the gate-description convention. Additional angles introduced by the simulation are computed to inverse-polynomial accuracy in logarithmic space. Exact restoration is assumed for the original source channel; approximation is used only in the target construction. This distinction is relevant to the gate-set issues discussed in [2, Section 3.3].

Theorem 1. *Every polynomial-time, deterministic-logspace-uniform quantum catalytic computation in the model above, with bounded error, $O(\log n)$ initialized work qubits, and polynomially many catalytic qubits, has a polynomial-size logspace-uniform simulation using one clean qubit and otherwise maximally mixed qubits. The simulation is unitary until its final measurement and has inverse-polynomial signed acceptance bias. Consequently, under the stated circuit and bias conventions,*

$$\mathsf{BQCL} \subseteq \mathsf{DQC1}.$$

The containment also holds for the exact class QCL.

1.2 Method and related work

The proof uses a history-matrix inverse to represent a product of channel matrices. Its conditioning may be exponentially poor. Approximating the entire inverse in operator norm would therefore give an unsuitable simulation. What we need, however, is only a particular normalized trace. Every nonzero block of the inverse is the matrix of a channel, so complete positivity bounds the inverse's Frobenius norm. A regularization estimate in nuclear norm then controls the trace error at an inverse-polynomial regularization scale. The catalyst dimension cancels from this bound. This use

of channel structure, together with the catalytic trace identity, is the main step specific to the measured catalytic model.

Measurement elimination in ordinary space-bounded quantum computation was established by Fefferman and Remscrim [3]. For a computation using space S and time t, their construction uses $O(S + \log t)$ space and $\mathrm{poly}(t2^S)$ time. Applying that bound with $S = s + c$ does not provide a polynomial-time one-clean-qubit simulation when c is polynomial. The present argument instead treats the initialized and catalytic parts separately and uses exact restoration to control a normalized trace rather than the whole computation.

The circuit ingredients for processing this trace are established ones. Normalized unitary trace estimation is native to the one-clean-qubit model [4, 5]. Cade and Montanaro [6] gave one-clean-qubit algorithms for smooth spectral functions of log-local Hamiltonians. Edenhofer, Hasegawa, and Le Gall [7] treated spectral sums in a block-encoding input model, including trace extraction with logarithmically many block ancillas. More recently, Ji et al. [8] studied DQC1-completeness of normalized spectral traces under approximation-theoretic hypotheses. These results concern trace estimation once a suitable matrix representation has been obtained; the catalytic reduction and its dimension-independent regularization bound supply that representation here.

Our implementation uses linear combinations of unitaries and the compression-to-walk construction underlying qubitization and singular value transformation [9, 10]. We give a direct coherent phase-estimation implementation of the required filter to keep its logarithmic auxiliary space and deterministic logspace uniformity explicit. A reflection then extracts the desired ancillary block from a unitary trace. Thus the block ancillas need not be physically initialized in the final circuit.

2 The catalytic trace identity

We first express an adaptive computation as a fixed channel sequence. A computational-basis measurement is implemented by resetting a work bit, copying the measured basis value into it, and dephasing that bit. Later controlled gates read the resulting classical configuration. Irreversible controller updates use reversible Boolean gates and resets on logarithmic work space. By scanning the polynomially many tape addresses, the emitted schedule can be made independent of measurement outcomes. Appendix A gives the implementation, including conditional resets and the accounting for temporary work bits.

This is the mixed-state circuit formalism of Aharonov, Kitaev, and Nisan [11], with the work-space accounting made explicit. After this conversion and padding, write

$$\Phi_x = \Phi_T \circ \cdots \circ \Phi_1, \qquad T = \mathrm{poly}(n). \tag{2}$$

Each elementary map is a constant-local unitary conjugation, a one-qubit dephasing, or a one-qubit reset. In particular, it is completely positive and trace preserving (CPTP) on the whole computational Hilbert space, including controller configurations unreachable from the clean input. We may take $s \geq 1$ and $T \geq 1$ and pad their bounds as functions of n. For small inputs, all asymptotic bounds can be read with n replaced by $\max\{n, 2\}$. Set

$$K = 2^s, \qquad D = 2^c, \qquad d = KD.$$

Thus K is polynomial, whereas D may be exponential.

Lemma 2 (Exact restoration)**.** *Let $\Psi : C \to WC$ be a CPTP map such that $\mathrm{Tr}_W \Psi(\rho) = \rho$ for every state ρ on C. There is a state η on W such that $\Psi(X) = \eta \otimes X$ for every operator X on C. Consequently, Ψ also preserves arbitrary correlations of C with an inaccessible reference.*

Proof. Since states span the operator space, $\mathrm{Tr}_W \Psi = \mathrm{id}_C$. Let $|\Omega\rangle = D^{-1/2} \sum_{i=0}^{D-1} |i\rangle_C |i\rangle_R$ be the normalized maximally entangled vector, and put $\Omega_{CR} = |\Omega\rangle\langle\Omega|$. The state $\tau = (\Psi \otimes \mathrm{id}_R)(\Omega_{CR})$ has CR marginal Ω_{CR}. Put $P = I_W \otimes \Omega_{CR}$. Then $\mathrm{Tr}((I-P)\tau) = 0$. Positivity implies $\tau^{1/2}(I-P) = 0$, since the squared Frobenius norm of this operator is $\mathrm{Tr}((I-P)\tau)$. Hence $\tau = P\tau P$. The range of P is $W \otimes \mathrm{span}\{|\Omega\rangle\}$, so $\tau = \eta \otimes \Omega_{CR}$ for a state η. Comparing coefficients of $|i\rangle\langle j|_R$ in the two Choi states gives $\Psi(|i\rangle\langle j|) = \eta \otimes |i\rangle\langle j|$ for all i, j. Linearity proves the claim, and tensoring this identity with id_R proves the reference statement. □

Applying the lemma to $\Psi(X) = \Phi_x(|0^s\rangle\langle 0^s| \otimes X)$ proves (1) from marginal restoration alone. It also shows why including the final controller and temporary registers in the work output does not impose an additional catalytic condition.

Vectorize operators by $\mathrm{vec}(|i\rangle\langle j|) = |i\rangle|j\rangle$, and order the doubled registers as $W, \overline{W}, C, \overline{C}$. Bars label the second operator-coordinate copy; they do not denote additional input information. A unitary conjugation by G has matrix $G \otimes \overline{G}$, up to this fixed reordering. Let S_t be the d^2-dimensional Liouville matrix of Φ_t, and set $S = S_T \cdots S_1$. Exact restoration gives

$$S(|\omega\rangle \otimes I_{D^2}) = |\eta_x\rangle \otimes I_{D^2}, \qquad |\omega\rangle = |0^{2s}\rangle, \quad |\eta_x\rangle = \mathrm{vec}(\eta_x). \tag{3}$$

This is an identity on all D^2 catalyst operator coordinates, including the off-diagonal ones.

Define the output observable and signed acceptance bias by

$$O = (|1\rangle\langle 1| - |0\rangle\langle 0|)_{\mathrm{out}} \otimes I_{W \setminus \mathrm{out}}, \qquad e_x = \mathrm{Tr}(O\eta_x) = 2p_x - 1, \qquad |z\rangle = \frac{\mathrm{vec}(O)}{\sqrt{K}}.$$

The vector $|z\rangle$ is normalized because $\mathrm{Tr}(O^2) = K$. It can be prepared by creating s Bell pairs and applying $XZX = -Z$ to the output half in the first copy, where X, Z are Pauli matrices. Retaining this sign is important in the later trace test. Equation (3) yields

$$\frac{1}{D^2} \mathrm{Tr}[(|\omega\rangle\langle z| \otimes I_{D^2})S] = \frac{e_x}{\sqrt{K}}. \tag{4}$$

Indeed, the trace over each catalyst coordinate is $\langle z \mid \eta_x\rangle = \mathrm{Tr}(O\eta_x)/\sqrt{K}$. For a bounded-error source, $e_x \geq 1/3$ on yes instances and $e_x \leq -1/3$ on no instances; for an exact source, $e_x \in \{-1, 1\}$.

The factor D^2 in (4) is specific to restoration of the full operator space. Merely preserving the maximally mixed state does not suffice: the completely depolarizing channel has Liouville trace 1, whereas the identity channel has Liouville trace D^2. Preserving only the computational-basis states similarly controls just D diagonal coordinates. The stronger multiplicity in (4) is what permits the dimension cancellation below.

3 Regularizing the channel history

We place the channel product in a corner of an inverse history matrix. The inverse need not have polynomial operator norm. Its blocks, however, satisfy a common Frobenius bound, and that is enough to approximate the trace in (4). For a matrix, $\|\cdot\|$ denotes operator norm, $\|\cdot\|_F$ Frobenius norm, and $\|\cdot\|_1$ nuclear norm.

Lemma 3 (CPTP Frobenius bound)**.** *If S_Ψ is the Liouville matrix of a CPTP map on a d-dimensional system, then $\|S_\Psi\|_F^2 \leq d^2$.*

Proof. The unnormalized Choi matrix $\mathcal{J}(\Psi) = \sum_{i,j} \Psi(|i\rangle\langle j|) \otimes |i\rangle\langle j|$ is positive semidefinite by complete positivity [12] and has trace d by trace preservation. Its entries are a reshuffling of the entries of S_Ψ, so the two Frobenius norms agree. If $\mu_i \geq 0$ are its eigenvalues, then

$$\|S_\Psi\|_F^2 = \mathrm{Tr}(\mathcal{J}(\Psi)^2) = \sum_i \mu_i^2 \leq \left(\sum_i \mu_i\right)^2 = d^2. \qquad \square$$

Let J be the least power of two with $J \geq T + 1$. On a J-state clock and the Liouville register, define

$$L = \sum_{t=1}^{T} |t\rangle\langle t-1| \otimes S_t, \qquad A = I - L, \qquad B = A^{-1} = \sum_{r=0}^{T} L^r. \tag{5}$$

The inverse formula follows from $L^{T+1} = 0$. In particular, $B_{T,0} = S$. Apart from its J diagonal identity blocks, the only nonzero blocks of B are $B_{t,r} = S_t \cdots S_{r+1}$ for $0 \leq r < t \leq T$. Each such product is the Liouville matrix of a CPTP map. Since there are at most J^2 nonzero blocks, Lemma 3 gives

$$\|B\|_F^2 \leq J^2 d^2 = J^2 K^2 D^2. \tag{6}$$

This bound does not require a polynomial lower bound on the smallest singular value of A.

Lemma 4 (Nuclear-norm regularization)**.** *For an invertible matrix A and $\lambda > 0$, put*

$$B_\lambda = A^\dagger (AA^\dagger + \lambda^2 I)^{-1}.$$

Then

$$\left\| A^{-1} - B_\lambda \right\|_1 \leq \frac{\lambda}{2} \left\| A^{-1} \right\|_F^2. \tag{7}$$

Proof. If $A = U\Sigma V^\dagger$ is a singular value decomposition, then $A^{-1} - B_\lambda$ is V times the diagonal matrix with entries

$$\frac{1}{\sigma} - \frac{\sigma}{\sigma^2 + \lambda^2} = \frac{\lambda^2}{\sigma(\sigma^2 + \lambda^2)} \leq \frac{\lambda}{2\sigma^2}$$

times $U^\dagger$. The entries are nonnegative. The inequality is $2\lambda\sigma \leq \sigma^2 + \lambda^2$; summing over the singular values proves (7). $\square$

Define the contraction

$$Q = |0\rangle\langle T|_{\mathrm{clock}} \otimes |\omega\rangle\langle z| \otimes I_{D^2}. \tag{8}$$

Equations (4) and (5) imply

$$\frac{\mathrm{Tr}(QB)}{D^2} = \frac{e_x}{\sqrt{K}}. \tag{9}$$

Using $|\,\mathrm{Tr}(QM)| \leq \|Q\|\,\|M\|_1$, we obtain

$$\left| \frac{\mathrm{Tr}(QB_\lambda)}{D^2} - \frac{e_x}{\sqrt{K}} \right| \leq \frac{\|B - B_\lambda\|_1}{D^2} \leq \frac{\lambda}{2} J^2 K^2. \tag{10}$$

The exponential factor D^2 in the Frobenius bound cancels against the trace normalization supplied by exact restoration. Hence an inverse-polynomial λ suffices even when the inverse itself is ill-conditioned. The bounded matrix λB_λ, rather than B in operator norm, is the object implemented next.

4 The one-clean-qubit simulation

A block encoding of a contraction M is a unitary U_M such that $(\langle 0^a| \otimes I)U_M(|0^a\rangle \otimes I) = M$. The all-zero ancillas specify a matrix block. They are not assumed to be initialized in the final simulator; the trace readout will extract that block from a maximally mixed input.

4.1 Encoding the history matrix

Every elementary channel matrix S_t has a three-ancilla block encoding of $S_t/2$. For a unitary channel, $S_t = G \otimes \overline{G}$ is unitary. For one-qubit dephasing, $S_t = (II + ZZ)/2$, where juxtaposition denotes a tensor product of Pauli matrices. The reset $X \mapsto |0\rangle\langle 0| \operatorname{Tr} X$ has Liouville matrix

$$S_{\text{reset}} = |00\rangle\langle 00| + |00\rangle\langle 11| = \frac{1}{4}(II + IZ + ZI + ZZ + XX + iXY + iYX - YY). \tag{11}$$

The eight signed Pauli terms in (11), averaged with equal weights, therefore encode $S_{\text{reset}}/2$. For a unitary S_t, use four copies of S_t and the four terms $I, -I, I, -I$; their average is $S_t/2$. For dephasing, use $II, II, ZZ, ZZ, II, -II, II, -II$. An inactive clock value uses four copies each of I and $-I$, encoding zero. The signs and factors of i are phases of the selected unitaries, not probabilities.

Preparing three selector qubits with Hadamard gates, applying the selected term, and undoing the preparation realizes each average. This is the constant-size linear-combination construction of [10, Lemma 52]. Multiplex the encoding of $S_{r+1}/2$ on clock value $r < T$ and the encoding of zero on $r \geq T$, then cyclically increment the clock. The resulting unitary U_L encodes $L/2$. The zero blocks exclude both unused transitions and clock wraparound.

An additional selector, prepared in $\sqrt{1/3}|0\rangle + \sqrt{2/3}|1\rangle$, selects I or $-U_L$. Undoing this preparation gives a unitary U_A with all-zero ancillary block

$$A/3 = (I - L)/3. \tag{12}$$

The multiplexers are implemented by enumerating clock values and source gate indices. Their controls act on $O(\log n)$ bits and admit ancilla-free polynomial-size implementations, as described in Appendix A. Thus U_A has polynomial size and a logspace-uniform description. It uses only $O(\log n)$ qubits beyond the doubled data register.

Add a selector qubit h and form the Hermitian unitary

$$U = |0\rangle\langle 1|_h \otimes U_A + |1\rangle\langle 0|_h \otimes U_A^\dagger.$$

Let Π project onto the all-zero block-encoding ancillas, leaving h and the history/data register unrestricted. If their number is a, then the compression of U to this subspace is

$$H = (\langle 0^a| \otimes I)U(|0^a\rangle \otimes I) = \frac{1}{3}\begin{pmatrix} 0 & A \\ A^\dagger & 0 \end{pmatrix}. \tag{13}$$

In particular, $\|H\| \leq 1$. This dilation turns the regularized inverse into a scalar function of a Hermitian contraction.

4.2 Coherent spectral filtering

Lemma 5 (Coherent spectral filter)**.** *Let $0 < \delta, \varepsilon \leq 1$, with δ^{-1} and ε^{-1} bounded by polynomials in n and with logspace-computable descriptions. There is a polynomial-size, logspace-uniform unitary whose all-zero ancillary block is within operator norm ε of*

$$f_\delta(H) = \frac{\delta H}{H^2 + \delta^2 I}.$$

The ancillary register has $O(\log n)$ qubits.

Proof. Put $\mathcal{R} = 2\Pi - I$ and $V = \mathcal{R}U$. Since $U = U^\dagger$,

$$\frac{V + V^\dagger}{2} = \Pi U \Pi - (I - \Pi)U(I - \Pi). \tag{14}$$

This operator is block diagonal with respect to Π, with compression H. Define the periodic function $g_\delta(\theta) = f_\delta(\cos\theta)$, where $f_\delta(t) = \delta t/(t^2 + \delta^2)$. On a V-eigenvector with eigenvalue $e^{i\theta}$, the Hermitian operator $(V + V^\dagger)/2$ has eigenvalue $\cos\theta$. Functional calculus in (14) therefore gives

$$(\langle 0^a| \otimes I) g_\delta(V) (|0^a\rangle \otimes I) = f_\delta(H), \tag{15}$$

where $g_\delta(V)$ denotes the operator taking value $g_\delta(\theta)$ at eigenphase θ. The scalar bounds

$$|g_\delta(\theta)| \le \frac{1}{2}, \qquad |g'_\delta(\theta)| \le \frac{1}{\delta}$$

follow from $2\delta|t| \le t^2 + \delta^2$ and $|f'_\delta(t)| = \delta|\delta^2 - t^2|/(t^2 + \delta^2)^2 \le 1/\delta$.

Apply coherent phase estimation to V with an m-qubit phase register. For phase label j, apply to another qubit the real rotation whose zero-to-zero amplitude is $g_\delta(2\pi j/2^m)$, and then uncompute phase estimation. Such a rotation is $R_y(2\arccos g_\delta(2\pi j/2^m))$, with $R_y(\alpha) = e^{-i\alpha Y/2}$. For a V-eigenvector of phase θ, the resulting all-zero phase/rotation block has amplitude

$$\sum_{j=0}^{2^m-1} \Pr[j \mid \theta] g_\delta(2\pi j/2^m). \tag{16}$$

The probabilities appear because phase estimation is uncomputed; no measurement or postselection is used.

Choose circular phase error $O(\varepsilon\delta)$ and failure probability $O(\varepsilon)$. The Lipschitz and range bounds imply a uniform $O(\varepsilon)$ error in (16). The explicit estimate in Appendix B gives error at most $\varepsilon/8$ with

$$m = O\left(\log \frac{1}{\delta\varepsilon^2}\right).$$

Phase estimation uses $2^m - 1$ calls to V when controlled powers are expanded into repetitions, hence polynomially many calls. Approximating the label-dependent rotations and elementary target operations within the remaining error budget gives operator-norm error at most ε: the phase-average estimate holds uniformly on every eigenspace of V, and compression cannot increase operator norm.

For uniformity, enumerate the 2^m phase labels in the classical circuit description. The rotation amplitudes are

$$\frac{\delta\cos(2\pi j/2^m)}{\cos^2(2\pi j/2^m) + \delta^2}.$$

The denominator is at least δ^2 and the amplitude has magnitude at most $1/2$. Computing cosine, this quotient, and its inverse cosine to sufficient fixed precision takes logarithmic space and polynomial time; Appendix A gives explicit series and precision bounds. The circuit itself needs no arithmetic workspace for these values: it contains one classically specified, label-controlled rotation for each j. All its large controls have logarithmic length. Finally compress the original a block ancillas as in (15). The total ancillary space is $O(\log n)$. □

Set $\delta = \lambda/3$. Squaring the block matrix in (13) shows that the lower-left selector block of the filter is

$$\langle 1|_h f_{\lambda/3}(H)|0\rangle_h = \lambda A^\dagger (AA^\dagger + \lambda^2 I)^{-1} = \lambda B_\lambda. \tag{17}$$

Thus one inverse-polynomial scaling factor suffices for the whole channel product, rather than one such factor per channel.

4.3 Trace extraction and decision bias

Let U_f be the filter circuit from Lemma 5. Separate its D^2-dimensional register $C\overline{C}$ from all the remaining registers. The latter have $k = O(\log n)$ qubits, including the doubled work register, clock, selector h, and every block, phase, and rotation ancilla. There is an explicitly specified unitary F on these small registers satisfying

$$F|0^k\rangle = |1\rangle_h|T\rangle_{\text{clock}}|z\rangle_{W\overline{W}}|0\rangle_{\text{block/phase/rotation ancillas}}.$$

It uses basis-state preparation and the Bell-pair preparation of $|z\rangle$. Define $W_* = F^\dagger U_f$ and

$$a_x = \frac{1}{D^2}\operatorname{Re}\operatorname{Tr}\Big[(\langle 0^k|\otimes I)W_*(|0^k\rangle\otimes I)\Big]. \tag{18}$$

The selected input is $|0\rangle_h|0\rangle_{\text{clock}}|\omega\rangle$ and the selected output is $|1\rangle_h|T\rangle_{\text{clock}}|z\rangle$. Consequently, (17) and (10) imply

$$\left|a_x - \lambda\frac{e_x}{\sqrt{K}}\right| \le \frac{\lambda^2}{2}J^2K^2 + \varepsilon. \tag{19}$$

Here the filter error remains at most ε after taking the selected block, and a D^2-dimensional matrix of operator norm at most ε has normalized trace of magnitude at most ε.

The following trace cancellation is a direct block-readout construction. It gives the logarithmic-ancilla trace reduction used in one-clean-qubit spectral estimation, without initializing those ancillas; compare [7, Theorem 4.6].

Lemma 6 (Ancillary-block trace readout)**.** *For W_* and a_x as above, a one-clean-qubit circuit has final probe expectation $a_x/2^k$. Its other qubits are maximally mixed.*

Proof. Let $P = |0^k\rangle\langle 0^k|\otimes I_{D^2}$. Introduce a maximally mixed qubit b and define the unitary

$$\mathcal{V} = |0\rangle\langle 0|_b\otimes W_* + |1\rangle\langle 1|_b\otimes(2P - I)W_*.$$

Its normalized trace satisfies

$$\frac{\operatorname{Tr}\mathcal{V}}{2^{k+1}D^2} = \frac{\operatorname{Tr}(W_*) + \operatorname{Tr}((2P-I)W_*)}{2^{k+1}D^2} = \frac{\operatorname{Tr}(PW_*)}{2^kD^2}. \tag{20}$$

Prepare the clean probe in $|+\rangle$, apply controlled-$\mathcal{V}$, and apply a Hadamard to the probe. For a maximally mixed target of dimension $N = 2^{k+1}D^2$, the off-diagonal entry of the probe state before the last Hadamard is $\operatorname{Tr}(\mathcal{V}^\dagger)/(2N)$. Its final Pauli-$Z$ expectation is therefore $\operatorname{Re}\operatorname{Tr}(\mathcal{V})/N$, which is $a_x/2^k$ by (20). This is the one-clean-qubit trace test [4].

The reflection involves only the k small-register qubits. Write

$$P_0 = |0^k\rangle\langle 0^k| = 2^{-k}\sum_{S\subseteq[k]} Z_S, \qquad Z_S = \prod_{j\in S} Z_j,$$

where Z_j acts on the jth small-register qubit and $Z_\varnothing = I$. Since P_0 is a projector, $2P_0 - I = -e^{i\pi P_0}$. The exponential factors into 2^k commuting Pauli rotations. Each Pauli-string rotation has an ancilla-free implementation, and $2^k = \text{poly}(n)$. In a controlled implementation, the leading minus sign is retained as a phase on the control; the empty-set term is also retained. Adding the controls on b and the probe leaves polynomial size and logarithmic-space uniformity unchanged, as detailed in Appendix A. □

Choose

$$\lambda = \frac{1}{100J^2K^3}, \qquad \delta = \lambda/3, \qquad \varepsilon = \frac{\lambda}{200K}. \tag{21}$$

All three reciprocals are polynomially bounded. Equation (19) then gives

$$\left| a_x - \lambda \frac{e_x}{\sqrt{K}} \right| \leq \frac{\lambda}{100K} \leq \frac{\lambda}{100\sqrt{K}}. \tag{22}$$

As $1/3 - 1/100 > 1/4$, we have

$$a_x \geq \frac{\lambda}{4\sqrt{K}} \quad \text{on yes instances}, \qquad a_x \leq -\frac{\lambda}{4\sqrt{K}} \quad \text{on no instances}.$$

Accepting when the final probe is 0, Lemma 6 gives

$$p_{\text{yes}} \geq \frac{1}{2} + \gamma, \qquad p_{\text{no}} \leq \frac{1}{2} - \gamma, \qquad \gamma = \frac{\lambda}{2^{k+3}\sqrt{K}}. \tag{23}$$

A final NOT converts this to the convention that output 1 means acceptance, without another qubit.

For completeness, the construction tolerates further approximation of target gates. Let $N_g = \text{poly}(n)$ bound the number of elementary gates in the complete unitary trace circuit, and put $\xi = \lambda/(2^{k+6}K)$. Approximate each gate requiring numerical synthesis in operator norm by at most ξ/N_g, retaining any exact source primitives as such. Telescoping products bounds the total unitary error by ξ. For any input state, the change in a measurement probability is at most $2\xi \leq \gamma/4$. Thus the weaker gap $\gamma/2$ is valid for the synthesized circuit as well. Every required accuracy is still inverse-polynomial. Approximating these target operations does not alter the source channel to which the exact restoration hypothesis was applied.

The final circuit has one clean probe and $2c + k + 1$ maximally mixed qubits. The k small-register qubits include every auxiliary register used for block encoding and filtering; the additional mixed qubit is b. No initialized measurement record or hidden clean workspace is used. All operations before the final probe measurement are unitary. The channel multiplexer, the repeated walk calls in phase estimation, the phase-label rotations, and the small-register reflections each have polynomial size. Their gate descriptions are generated with logarithmic-space counters and arithmetic, as justified in Appendix A. Padding the source and target register bounds as functions of n gives a fixed reciprocal-polynomial lower bound on the bias. This proves Theorem 1; exact source acceptance probabilities satisfy the same separation from 1/2.

References

[1] Harry Buhrman, Richard Cleve, Michal Koucký, Bruno Loff, and Florian Speelman. Computing with a full memory: Catalytic space. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pp. 857–866, 2014. doi:10.1145/2591796.2591874.

[2] Harry Buhrman, Marten Folkertsma, Ian Mertz, Florian Speelman, Sergii Strelchuk, Sathyawageeswar Subramanian, and Quinten Tupker. Quantum catalytic space. In *20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025)*, LIPIcs **350**, pp. 11:1–11:24, 2025. doi:10.4230/LIPIcs.TQC.2025.11.

[3] Bill Fefferman and Zachary Remscrim. Eliminating intermediate measurements in space-bounded quantum computation. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pp. 1343–1356, 2021. doi:10.1145/3406325.3451051. Full version: arXiv:2006.03530.

[4] E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters* **81**(25), 5672–5675, 1998. doi:10.1103/PhysRevLett.81.5672.

[5] Dan Shepherd. Computation with unitaries and one pure qubit. arXiv:quant-ph/0608132, 2006. arXiv:quant-ph/0608132.

[6] Chris Cade and Ashley Montanaro. The quantum complexity of computing Schatten p-norms. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*, LIPIcs **111**, pp. 4:1–4:20, 2018. doi:10.4230/LIPIcs.TQC.2018.4.

[7] Roman Edenhofer, Atsuya Hasegawa, and François Le Gall. Dequantization and hardness of spectral sum estimation. arXiv:2509.20183v1, 2025. arXiv:2509.20183v1.

[8] Zhengfeng Ji, Tongyang Li, Changpeng Shao, Xinzhao Wang, and Yuxin Zhang. DQC1-completeness of normalized trace estimation for functions of log-local Hamiltonians. Accepted to the 67th IEEE Symposium on Foundations of Computer Science (FOCS 2026). Preprint, arXiv:2604.01519, 2026. arXiv:2604.01519.

[9] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum* **3**, 163, 2019. doi:10.22331/q-2019-07-12-163.

[10] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pp. 193–204, 2019. arXiv:1806.01838.

[11] Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC 1998)*, pp. 20–30, 1998. doi:10.1145/276698.276708.

[12] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications* **10**(3), 285–290, 1975. doi:10.1016/0024-3795(75)90075-0.

A Uniform implementation

A.1 Compiling the adaptive source

A computational-basis measurement of a qubit q with a retained outcome is the channel obtained by resetting a work bit b to 0, applying $\mathrm{CNOT}_{q\to b}$, and dephasing b. On an input operator X this gives $\sum_{u\in\{0,1\}} P_u X P_u \otimes |u\rangle\langle u|_b$, where $P_u = |u\rangle\langle u|_q \otimes I$ on the measured system and its other registers. This is exactly the measurement instrument, including its effect on correlations. Overwriting an outcome is a reset on its work bit.

A logspace classical controller has polynomially many configurations. Its live tape, internal state, and input and quantum-tape head positions occupy $O(\log n)$ bits. During one simulated transition, keep its old configuration fixed. Scan all relevant gate types and tape addresses, using equality tests on the old configuration to select the prescribed operation. The addresses occupy logarithmic space because the input and catalytic tape lengths are polynomial. Exactly the selected operation acts on each classical branch; the configuration is not updated midway through the scan.

Conditional measurements are handled by conditionally copying the addressed basis value into a zero work bit and then dephasing that bit. For a conditional reset, first reset a temporary work qubit to 0, conditionally swap it with the addressed qubit, and reset the temporary qubit again. On the inactive branch the addressed qubit is unchanged; on the active branch it is reset. The condition is classical on all reachable configurations. These implementations are nevertheless compositions of unitaries, dephasings, and resets on the entire Hilbert space, so they define CPTP maps even on unreachable inputs.

Compute the new classical configuration into a second logarithmic register using reversible Boolean gates, the old configuration, the outcome bits, and the read-only input. Uncompute the temporary equality flags, reset the old configuration and expired outcomes, and exchange the roles of the two configuration registers. Invalid encodings may be assigned a fixed halting transition. Halting configurations use identity operations, permitting padding to the worst-case polynomial time bound. All temporary registers are logarithmic and can be reset and reused in the source. The scan, the controller transition, and their local-gate decompositions have logspace-generated polynomial-length descriptions. This proves the fixed-channel representation (2) without assuming any initialized history outside the work-space bound.

A.2 Controls and circuit generation

The target construction can implement all its controls without initialized scratch qubits. For a specified ℓ-bit control string, let

$$P_{\mathrm{ctrl}} = \prod_{j=1}^{\ell} \frac{I + \sigma_j Z_j}{2}, \qquad \sigma_j \in \{-1, 1\}.$$

A rotation conditioned on this string has generator $P_{\mathrm{ctrl}} \otimes A$, where A is a one-qubit Pauli matrix. Expanding the product gives 2^ℓ mutually commuting Pauli strings, so exponentiating the sum gives a product of Pauli-string rotations. Each such rotation is implemented by changing the Pauli bases to Z, computing parity onto one of its participating qubits with CNOT gates, rotating that qubit, and undoing the parity computation and basis changes. This uses no extra qubit. Purely diagonal phase gates are treated by the same expansion. In particular, a controlled NOT is exactly

$$\exp\left[-\frac{i\pi}{2} P_{\mathrm{ctrl}} \otimes (X - I)\right],$$

which keeps the relative phase that would be lost by replacing X with a Pauli rotation alone.

In every application, $\ell = O(\log n)$, so enumerating the subsets and their signed coefficients costs polynomial time and logarithmic space. Cyclic clock increments can be decomposed into successive controlled bit flips. Phase-label controls, the history selectors, and the reflections use the same construction. All scalar phases are retained when a unitary is subsequently controlled. Additional constant-local source primitives are handled under the gate-description closure convention stated in the introduction. Their controls from the clock, selectors, and phase estimation together still have logarithmic length.

An indexed source gate can be recovered by replaying the source's logspace generator with a gate counter. The inverse of a polynomial gate list is generated by replaying this procedure in reverse index order and taking each gate's adjoint. Complex conjugation is gatewise. Nested loops for the clock multiplexer and the phase-estimation powers require only a fixed number of polynomially bounded counters. They therefore preserve deterministic logspace uniformity.

A.3 Angles and precision

Only the extra target angles require numerical evaluation by the compiler. The fixed selector weights, the Fourier phases, the reflection phases, and the phase-filter rotations can all be evaluated with $O(\log n)$ bits of precision. To make the latter claim explicit, reduce the argument of cosine to $[-\pi, \pi]$ and use its Taylor series. After r terms the error is at most $\pi^{2r}/(2r)!$. For inverse-polynomial error, $r = O(\log n)$ suffices. The terms can be generated one at a time by their recurrence, maintaining only the current term and partial sum. A logspace-computable value of π can be obtained, for example, from $\pi = 16 \arctan(1/5) - 4 \arctan(1/239)$ and the geometrically convergent arctangent series.

If $t = \cos\theta$ is perturbed by ν, the filter value $\delta t/(t^2 + \delta^2)$ changes by at most $|\nu|/\delta$. The quotient can also be evaluated directly with guard bits, since its denominator is at least δ^2. Both losses are polynomial. After evaluating the quotient, clip its approximation to $[-1/2, 1/2]$; this does not increase its error relative to the exact filter value. On this interval, $|\frac{d}{du} \arccos u| \le 2/\sqrt{3}$. Evaluate $\arccos u = \pi/2 - \arcsin u$ using

$$\arcsin u = \sum_{r=0}^{\infty} \frac{\binom{2r}{r}}{4^r (2r+1)} u^{2r+1}.$$

For $|u| \le 1/2$, the absolute tail after r terms is bounded by a constant times 4^{-r}. Again logarithmically many terms suffice, and the summands a_r obey $a_{r+1} = a_r u^2 (2r+1)^2/[2(r+1)(2r+3)]$ with $a_0 = u$. Thus only the previous summand, the partial sum, and an index are needed.

Fixed-precision addition, multiplication, and division on $O(\log n)$-bit registers can be performed by schoolbook algorithms using $O(\log n)$ work bits and polynomial time. The series require only a constant number of such registers and a logarithmic counter. Polynomial error amplification and logarithmically many rounding steps are absorbed by choosing a sufficiently large constant in the $O(\log n)$ precision bound. The number of controlled rotations emitted is polynomial, so allocating reciprocal-polynomial error to each gate also keeps the aggregate error within Lemma 5 and the budget following (23). This numerical work is performed by the classical circuit generator, not by a hidden clean quantum register. No eigenvalue computation or phase-finding procedure on a classically supplied matrix is needed.

B A phase-estimation error bound

For an eigenphase $2\pi\varphi$ of V, phase estimation with $M = 2^m$ labels has amplitudes

$$\alpha_j(\varphi) = \frac{1}{M} \sum_{r=0}^{M-1} e^{2\pi i r(\varphi - j/M)}.$$

If the circular distance between φ and j/M is $0 < \Delta \leq 1/2$, the geometric-sum formula and $|\sin(\pi\Delta)| \geq 2\Delta$ give

$$|\alpha_j(\varphi)|^2 \leq \frac{1}{4M^2\Delta^2}.$$

Fix $0 < \zeta \leq 1/2$ with $M\zeta \geq 1$. On either side of φ around the circle, the quantities $M\Delta$ for labels with $\Delta > \zeta$ form a truncated sequence whose successive terms differ by 1 and whose first term is at least $M\zeta$. Thus the failure probability is at most

$$\frac{1}{2} \sum_{r=0}^{\infty} \frac{1}{(M\zeta + r)^2} \leq \frac{1}{2} \left(\frac{1}{(M\zeta)^2} + \frac{1}{M\zeta} \right) \leq \frac{1}{M\zeta}.$$

The middle inequality follows by bounding the sum after its first term by the integral of $u \mapsto (M\zeta + u)^{-2}$ over $[0, \infty)$.

Take $\zeta = \varepsilon\delta/(32\pi)$ and let M be the least power of two at least $4096/(\delta\varepsilon^2)$. On the event of circular error at most ζ, the change in g_δ is at most $2\pi\zeta/\delta = \varepsilon/16$. On the complementary event it is at most 1, while that event has probability at most $1/(M\zeta) \leq \pi\varepsilon/128 < \varepsilon/16$. The error in the ideal phase average (16) is therefore at most $\varepsilon/8$, uniformly over all eigenphases. This leaves room for the rotation and elementary-gate errors in Lemma 5. The choices give $m = O(\log(1/(\delta\varepsilon^2)))$ and polynomially many walk calls.