A distillation–teleportation protocol for fault-tolerant QRAM

Alexander M. Dalzell,[1] András Gilyén,[2] Connor T. Hann,[1] Sam McArdle,[1] Grant Salton,[1,3] Quynh T. Nguyen,[4] Aleksander Kubica,[1,5] Fernando G.S.L. Brandão[1,6]

[1] *AWS Center for Quantum Computing, Pasadena, CA, USA*
[2] *HUN-REN Alfréd Rényi Institute of Mathematics, Budapest, Hungary*
[3] *Amazon Quantum Solutions Lab, Seattle, WA, USA*
[4] *School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA*
[5] *Yale Quantum Institute & Department of Applied Physics, New Haven, CT, USA*
[6] *Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA*

Abstract

We present a protocol for fault-tolerantly implementing the logical quantum random access memory (QRAM) operation, given access to a specialized, noisy QRAM device. For coherently accessing classical memories of size 2^n, our protocol consumes only $\mathrm{poly}(n)$ fault-tolerant quantum resources (logical gates, logical qubits, quantum error correction cycles, etc.), avoiding the need to perform active error correction on all $\Omega(2^n)$ components of the QRAM device. This is the first rigorous conceptual demonstration that a specialized, noisy QRAM device could be useful for implementing a fault-tolerant quantum algorithm. In fact, the fidelity of the device can be as low as $1/\mathrm{poly}(n)$. The protocol queries the noisy QRAM device $\mathrm{poly}(n)$ times to prepare a sequence of n-qubit QRAM *resource states*, which are moved to a general-purpose $\mathrm{poly}(n)$-size processor to be encoded into a QEC code, distilled, and fault-tolerantly teleported into the computation. To aid this protocol, we develop a new gate-efficient streaming version of quantum purity amplification that matches the optimal sample complexity in a wide range of parameters and is therefore of independent interest.

The exponential reduction in fault-tolerant quantum resources comes at the expense of an exponential quantity of purely classical complexity—each of the n iterations of the protocol requires adaptively updating the 2^n-size classical dataset and providing the noisy QRAM device with access to the updated dataset at the next iteration. While our protocol demonstrates that QRAM is more compatible with fault-tolerant quantum computation than previously thought, the need for significant classical computational complexity exposes potentially fundamental limitations to realizing a truly $\mathrm{poly}(n)$-cost fault-tolerant QRAM.

Contents

1 Introduction

The development of fast and large-scale random access memory (RAM) has played an indispensable role in the development of conventional computing. Early RAM devices assisted in the first demonstrations of stored-program electronic computers [1–3], and today, the availability of efficient high-speed RAM enables data-intensive computing applications in areas like machine learning [4, 5]. At an abstract level, RAM performs the following operation: take as input an n-bit address x specifying a location in memory, and retrieve one of the 2^n data items $f(x)$, labeled by x. In practice, the access time for RAM is astonishingly fast: modern RAM chips can achieve latency of 10 nanoseconds or faster.[1] Furthermore, the RAM runtime is independent of the location of the data within the memory, and the latency can remain nearly unchanged even as the overall size of the memory is scaled up.

The idea of quantum random access memory (QRAM) [7, 8] is to achieve something similar even when the n-qubit address register is in a quantum superposition $\sum_x \alpha_x|x\rangle$ of all 2^n addresses. For simplicity, we consider the most basic version of QRAM: applying a phase $(-1)^{f(x)}$ onto basis state $|x\rangle$, where the 2^n binary values $f(0), f(1), \ldots, f(2^n-1)$ are stored in classical memory.

$$\text{QRAM operation:} \qquad \sum_{x\in\{0,1\}^n} \alpha_x|x\rangle \quad \overset{V(f)}{\longmapsto} \quad \sum_{x\in\{0,1\}^n} (-1)^{f(x)}\alpha_x|x\rangle. \tag{1}$$

We refer to the n-bit Boolean function $f\colon \{0,1\}^n \to \{0,1\}$ as the classical *dataset* or *data table* we want to query. For each f, the n-qubit unitary $V(f)$ that implements the QRAM operation is diagonal in the computational basis, with diagonal ± 1 entries determined by f. We note that the controlled $V(f)$ operation[2] can be used to implement the more familiar formulation of QRAM, which reads the classical data into an ancilla register as $|x\rangle|0\rangle \mapsto |x\rangle|f(x)\rangle$.[3]

Whereas a single classical RAM query can access at most one entry of a data table, a single QRAM query suffices to create a superposition over all 2^n entries. The assumption that QRAM is cheap and available underlies a number of proposed quantum algorithms (see Refs. [8–11] for relevant surveys), which leverage this ability to offer up to exponential speedups over their classical counterparts. Often, the need for QRAM in these algorithms is contained within an unspecified oracle or data access assumption. For instance, quantum machine learning algorithms for support vector machines [12], Gaussian process regression [13], and recommendation systems [14] require only a polylogarithmic (in the size of the dataset) number of queries to an oracle that accesses (in superposition) the entries of a classical matrix or vector. Similarly, quantum algorithms for solving differential equations [15–20] discretize the equations and invert the resulting linear systems [21], in some cases incurring only a polylogarithmic (in the size of the linear system) number of queries to the classical data defining the instance, such as object geometries and boundary conditions. As a final example, quantum algorithms for solving optimization problems like semidefinite and linear programs [22–29], with applications in logistics and finance [30, 31], require coherent oracle access to the classical matrices defining the optimization problem. In all of these areas, the claimed speedup is typically dependent upon the assumption that—at least at an abstract level—the cost of QRAM is similar to that of RAM.

> **Cheap QRAM assumption.** *For an arbitrary data table f, the computational cost of implementing the unitary operation $V(f)$ from Eq. (1) is* $\mathrm{poly}(n)$.

Here, the term *computational cost* is intentionally vague—depending on the context, it might refer to circuit depth, physical runtime, energy dissipated, or some other metric—one must define it more precisely before justifying the assumption (see discussion in Ref. [8]). Focusing on physical runtime/latency as a metric, the assumption of $\mathrm{poly}(n)$ cost is roughly valid in the case of RAM: one can write down classical circuits for RAM that have $O(n)$ depth, and in practice actual RAM chips maintain extremely fast latency even at very large scale. However, for QRAM, the validity of this assumption has been the source of significant controversy [8, 9, 32, 33]. At the root of the issue is the fact that, unlike RAM, QRAM must be implemented in such a way that information about *which* address is being queried is not leaked to the environment, which would lead to decoherence. Strategies for preventing this decoherence without also reducing QRAM's relative power have so far proved to be elusive.

[1] RAM latency is a complex topic, due to the many kinds of memory that are used in a computer, each type offering benefits and drawbacks on a number of dimensions including speed, physical size per bit, volatility, and price. See Ref. [6] and the associated datasheet found there for a concrete example of an asynchronous static RAM chip achieving 10 nanoseconds latency on 2^{20} memory locations (each storing 8 bits).

[2] Controlled $V(f)$ is equivalent to (non-controlled) $V(\hat{f})$ for a dataset $\hat{f}$ with $n+1$ address bits; see Appendix A.

[3] In some places in the literature [8], the operation $|x\rangle|0\rangle \mapsto |x\rangle|f(x)\rangle$ is referred to as QRACM, where the additional C emphasizes the fact that the data $f(x)$ are classical and thus the states $|f(x)\rangle$ are computational basis states. This distinguishes QRACM from its generalization, QRAQM, where each $|f(x)\rangle$ can be an arbitrary (possibly multiqubit) quantum state. In this paper, we do not consider QRAQM, and we refer to QRAM interchangeably with QRACM.

One might try to justify the cheap QRAM assumption by writing down an $O(n)$-depth quantum circuit for the n-qubit unitary $V(f)$ [34–37], and then running that circuit on a general-purpose fault-tolerant quantum processor; assuming gates can be implemented in parallel, $O(n)$ latency is achievable. This strategy—referred to as "circuit QRAM" in Ref. [8]—has significant drawbacks. In particular, it requires $\Omega(2^n)$ logical ancilla qubits and $\Omega(2^n)$ classical co-processors to control the system and perform active error correction on all its components in parallel. Each logical ancilla may require dozens or hundreds of physical qubits, leading to an extremely large device footprint, a conclusion that is further exacerbated by the presence of a large number of magic state factories for implementing in parallel the non-Clifford T or Toffoli gates in the circuit, of which there must be at least $\Omega(\sqrt{2^n})$ [38]. One estimate for a surface code approach found that *quadrillions* of physical qubits would be needed for querying an 8-gigabyte memory [34]. The opportunity cost of these quantum and classical resources is steep. For example, the $O(2^n)$ classical co-processors can perform complex tasks like sparse matrix-vector multiplication for $2^n \times 2^n$ matrices in $\text{poly}(n)$ time [8, 9, 33]. Consequently, for circuit QRAM, the cheap QRAM assumption is only justifiable in a cost model that essentially precludes the possibility of quantum advantage in many proposed applications.

Ideally, the QRAM operation would instead be carried out by a specialized hardware element, separate from the main general-purpose quantum processor, mirroring how RAM is performed in a different way than computation on the main CPU. While the physical size of the QRAM hardware element would scale as $\Omega(2^n)$—this is necessary simply to store the dataset f—the runtime could be as little as $O(n)$. Furthermore, since the device is specialized for QRAM, it could in principle be performed *passively* and *ballistically* [8], that is, implemented automatically by natural evolution of the system while requiring at most $\text{poly}(n)$ external interventions from classical control and dissipating at most $\text{poly}(n)$ energy.[4] Constructing such a device is a formidable engineering challenge; there are currently no fully convincing proposals on how it could be done, but nothing rules it out in theory;[5] see Fig. 1a for an abstract picture of how such a device might be structured.

Yet, even if a passive, physical QRAM device did exist, it is unclear how it could actually be useful to a fault-tolerant quantum computation (FTQC). The ability to apply the physical QRAM operation at computational cost $\text{poly}(n)$ is not sufficient to justify the cheap QRAM assumption, even if there are no errors in the QRAM device itself (which is not realistic anyway). The problem is that, in FTQC, we need to perform the logical QRAM operation, denoted by $\overline{V(f)}$, onto an address register encoded into some quantum error-correcting (QEC) code. Naively, we could implement $\overline{V(f)}$ by un-encoding the n logical qubits into n physical qubits, running the physical qubits through the physical QRAM device, and re-encoding the output. However, the un-encoding and re-encoding processes introduce uncorrectable errors, and any noise in the physical QRAM device will also propagate into logical errors on the re-encoded state. An alternative would be to find a QEC code where the logical $\overline{V(f)}$ is a transversal gate, meaning it can be implemented as a tensor product of $\text{poly}(n)$ physical $V(f)$ gates without the need for un-encoding and re-encoding. Unfortunately, there are known challenges to finding such codes [8]. A general n-qubit QRAM gate is in the n-th level of the Clifford hierarchy (see Section 2.2 and Appendix D), and all known examples of codes supporting transversal implementation of a gate in the n-th level have $O(2^n)$ qubits [41–43]. In fact, there is at least one example of a gate in the n-th level—the single-qubit $\pi/2^n$ rotation gate—where a matching lower bound of $\Omega(2^n)$ qubits has been shown for a strong form of transversality [44], leading one to speculate that a similar lower bound may hold for QRAM, as well.

Our main contribution is to devise a protocol that implements the logical operation $\overline{V(f)}$ fault tolerantly, using $\text{poly}(n)$ queries to a noisy device that can implement the physical QRAM operation with at least $1/\text{poly}(n)$ fidelity, as well as $\text{poly}(n)$ fault-tolerant operations on a general-purpose quantum processor—exponentially lower than the number of fault-tolerant quantum operations required for circuit QRAM. The protocol generalizes well-known distillation–teleportation protocols for non-Clifford gates like the T gate and the CCZ gate. First, the physical $V(f)$ gate is used to prepare many copies of a faulty physical n-qubit QRAM resource state. Next, the physical resource states are encoded into a QEC code (the protocol is agnostic to which one) and distilled into a single high-fidelity logical resource state. Finally, the high-fidelity logical resource state is teleported into the computation to enact the logical QRAM gate, up to a correction which can be computed classically—our protocol outsources this calculation to a classical processor as depicted in Fig. 1b. The required correction is a different logical QRAM

[4] Reading from a classical RAM can be viewed as a passive operation: although the circuit for RAM has $\Omega(2^n)$ gates/components, these gates are etched onto the chip and are performed without any external intervention—one simply needs to set the voltages on the n input pins specifying the desired address. It is possible to design a RAM circuit that dissipates only $O(n)$ energy, although since this does not represent a bottleneck in practical systems, actual RAM chips are better modeled as dissipating $O(\sqrt{2^n})$ energy (memory is laid out in 2D and in practice an entire row/column is activated, rather than just a single memory cell) [8, 39].

[5] We ignore speed-of-light constraints, which we expect only to be relevant at large QRAM size [40]. At large enough scale, the speed of light would prevent both RAM and QRAM from achieving query latency $O(n)$, since the dataset of size 2^n must be embedded in 2 or at most 3 spatial dimensions, and the time needed for information to travel across the device would be at least $\Omega(2^{n/3})$ (in the case of a 3D embedding).

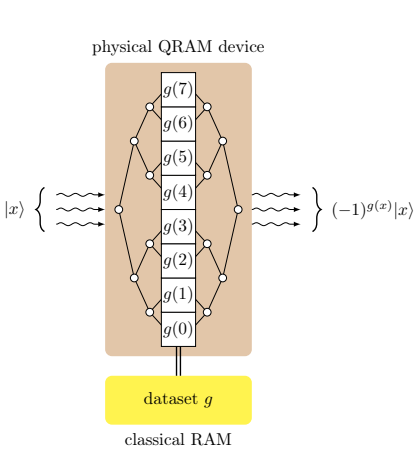

(a) Possible structure of physical QRAM device passively implementing $V(g)$ from Eq. (1)

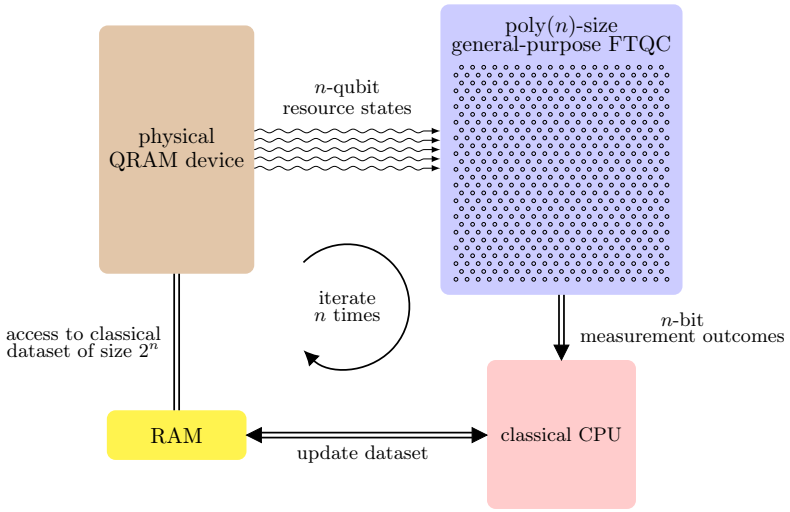

(b) Setup of our protocol

Figure 1: (a) Ideally, the physical QRAM operation $V(g)$ of Eq. (1) is performed passively by a specialized device. We may imagine, for example, encoding the address state $\sum_x \alpha_x|x\rangle$ into the polarization states of n photons, and then sending them into a pre-manufactured device, where they return having picked up a -1 phase only on branches of the superposition where $g(x) = 1$ [8, 45]. This might be accomplished by placing the classical bits $g(0), g(1), \ldots, g(2^n - 1)$ at the leaves of a binary tree, selectively routing the photons to the correct leaf based on their polarization, picking up a phase if a photon exists at location x and $g(x) = 1$, and then unrouting the n photons. (b) Our protocol utilizes a specialized, physical QRAM device, which is separate from the general-purpose fault-tolerant quantum processor. The QRAM device is used to create QRAM resource states on n physical qubits which are moved onto the main processor. The main processor encodes, distills, and teleports these resource states, generating classical n-bit measurement outcomes, which are sent to a classical CPU. The classical CPU performs a calculation to update the dataset stored in classical memory (RAM), which is queried by the physical QRAM device at the next iteration of the protocol.

gate $\overline{V(f')}$, where f' is determined by f and random measurement outcomes obtained during the teleportation procedure. The correction $\overline{V(f')}$ is then implemented in the same way, requiring a correction of its own, $\overline{V(f'')}$, where f'' is again dependent on f' and random measurement outcomes. We show that after iterating this process for n rounds, no further correction is necessary. This is a consequence of the fact that, despite its exponential circuit complexity, the unitary $V(f)$ lies in the n-th level of the Clifford hierarchy [46] for every f, which implies that the first correction $V(f')$ is in the $(n-1)$-th level, the second correction $V(f'')$ is in the $(n-2)$-th level, and so on. Our insights are (i) to notice that these corrections always lie within the family of QRAM gates of Eq. (1), allowing for a straightforward recursive implementation, and (ii) to devise a method for preparing the high-fidelity encoded resource states, completing the end-to-end workflow for fault-tolerant $\overline{V(f)}$. A no-go theorem in Ref. [8] ruled out a wide class of QRAM distillation–teleportation protocols; our protocol sidesteps this theorem by being adaptive and querying the physical QRAM on different datasets (f, f', f'', etc.) in each round. We provide a more complete informal overview of the protocol in Section 2 and a detailed error analysis of each step in Section 4.

By showing how to perform logical $\overline{V(f)}$ using $\mathrm{poly}(n)$ calls to physical QRAM, our protocol salvages the potential utility of the specialized, faulty QRAM device, and it encourages a model of quantum computation where QRAM is performed separately from the main quantum processing unit.

Our protocol also *partially* justifies the cheap QRAM assumption, provided that a passive QRAM device can be constructed. Indeed, if noisy physical QRAM has computational cost $\mathrm{poly}(n)$, then the quantum resources required to implement fault-tolerant QRAM via our protocol also scales only as $\mathrm{poly}(n)$. The main caveat is that running our protocol requires a non-negligible amount of adaptive *classical* computation of complexity $O(2^n)$ to compute the required correction operations (and "reload" the passive QRAM device, so that it has access to the new classical dataset at the next round of the protocol), although this complexity may be amenable to some degree of parallelization. We explore the nuances of this caveat in Section 5. In our protocol, this adaptive classical computation and QRAM reloading appears necessary in order to avoid revealing which address is being queried even while using a noisy QRAM device. In Section 7, we pose the question of whether this reflects an inevitable

limitation of fault-tolerant QRAM or whether a stronger justification of the cheap QRAM assumption, where both quantum and classical resources are poly(n), may be possible.

In any case, our protocol can be viewed as trading $O(2^n)$ quantum resources for $O(2^n)$ classical resources. That is, our protocol does not require the $O(2^n)$ actively error-corrected quantum resources incurred in circuit QRAM (fault-tolerant quantum gates, ancilla qubits, magic state factories, control wiring, classical co-processors, etc.). Instead, it requires poly$(n)2^n$ purely classical resources in addition to only poly(n) fault-tolerant quantum resources and poly(n) queries to a faulty QRAM device. There may be applications where such a tradeoff is beneficial, since quantum devices will be significantly slower and more expensive than classical devices for the foreseeable future.

2 Overview of protocol

2.1 Warm-up: distillation–teleportation protocol for the T gate

In many schemes for FTQC, it is relatively cheap to implement logical Clifford gates (e.g., they can often be done transversally). On the other hand, non-Clifford logical gates like the T gate and the CCZ gate are more expensive; these gates can instead be performed using distillation–teleportation protocols. In this section, we review such a protocol for the T gate, a diagonal gate mapping $|0\rangle \mapsto |0\rangle$ and $|1\rangle \mapsto \mathrm{e}^{\mathrm{i}\pi/4}|1\rangle$. Although the CCZ gate is actually a special case of the QRAM operation from Eq. (1) and thus more directly related to our protocol, the T gate provides a gentler introduction because it is a single-qubit gate.

We discuss distillation and gate teleportation separately, beginning with gate teleportation. Henceforth, we denote logical states and operations with an overline, for example, we denote the logical T gate by $\overline{T}$. Teleporting $\overline{T}$ into a quantum computation requires the preparation of a resource state, also known as a magic state, which is $\overline{T}$ applied to the equal superposition state:

$$|\overline{T}\rangle = \overline{T}|\overline{+}\rangle = \frac{1}{\sqrt{2}}(|\overline{0}\rangle + \mathrm{e}^{\mathrm{i}\pi/4}|\overline{1}\rangle)\,, \tag{2}$$

where $|\overline{0}\rangle$ and $|\overline{1}\rangle$ denote the encoded computational basis for some QEC code (here we are agnostic to which code), and $|\overline{+}\rangle = \frac{1}{\sqrt{2}}(|\overline{0}\rangle + |\overline{1}\rangle)$. The $\overline{T}$ gate can then be applied to an arbitrary quantum state $|\overline{\alpha}\rangle$ (on one logical qubit) by entangling $|\overline{\alpha}\rangle$ with $|\overline{T}\rangle$ and making a (logical) measurement, as follows:

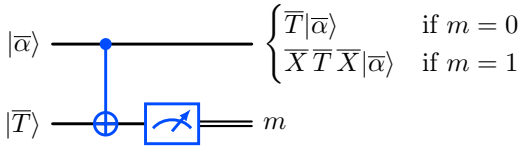

(3)

The logical CNOT gate (a Clifford gate) and the single-qubit logical measurement are both performed fault-tolerantly within the QEC code to ensure negligible chance of logical error. Direct computation verifies that the single-qubit measurement outcome $m \in \{0, 1\}$ is uniformly random, regardless of the state $|\overline{\alpha}\rangle$. If the measurement outcome is $m = 0$, the gate $\overline{T}$ is exactly implemented on the top wire. However, if the outcome $m = 1$ is obtained, the wrong phase was applied to the state, equivalent to the gate $\overline{X}\,\overline{T}\,\overline{X}$ instead of $\overline{T}$ (where X, Y, Z denote the Pauli operators). To fix this, one must apply a *correction* operation when the measurement outcome is 1. The correction required to undo the erroneous $\overline{X}\,\overline{T}\,\overline{X}$ gate and re-do the $\overline{T}$ gate is the $\overline{T}\,\overline{X}\,\overline{T}^\dagger\overline{X}$ gate, which is equal to the phase gate $\overline{S} = \overline{T}^2$, up to a global phase. Crucially, the phase gate is a Clifford gate, and thus the logical $\overline{S}$ can typically be implemented fault tolerantly in a more direct fashion. The full gate teleportation circuit with the correction is then given by:

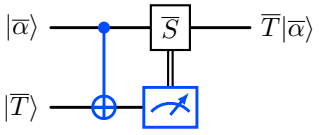

(4)

The benefit of performing the $\overline{T}$ gate via gate teleportation is that the difficulty is reduced to preparing a high-fidelity $|\overline{T}\rangle$ state. This state can be prepared through a multistep process of physical preparation, encoding, and then distillation. For concreteness, one can consider magic state injection schemes for the surface code [47–50].

Here, the first step is to prepare the $|T\rangle$ state on a single physical qubit. Next, an encoding procdure is performed, that is, $|T\rangle$ is mapped to $|\overline{T}\rangle$, which is encoded in a $d \times d$ surface code patch. This can be realized by, for instance, preparing a product state and performing appropriate stabilizer measurements. This procedure is generally not fault-tolerant; if the underlying hardware has error rate p, the logical error on the prepared state is $O(p)$, but the logical error can be kept independent of how large one makes the code distance d. Some of the possible logical errors are heralded—they can be detected by applying certain checks, in which case the procedure can be restarted from scratch, improving the postselected fidelity. The final step is magic state distillation, whereby multiple noisy $|\overline{T}\rangle$ states are consumed to produce a smaller number of higher-fidelity $|\overline{T}\rangle$ states. For example, the 15-to-1 magic state distillation protocol uses 15 input magic states of error rate p_{in}, succeeds with probability $1 - O(p_{\text{in}})$, and conditioned on success, produces a single output magic state of error rate $p_{\text{out}} = O(p_{\text{in}}^3)$ [50–52]. By recursively applying this protocol, one can distill $|\overline{T}\rangle$ states with arbitrarily low error rate even when all physical components have noise rate $p = O(1)$, provided that p is below a certain threshold for state distillation.

In some instances, one may not want to perform the corrective $\overline{S}$ gate directly.[6] In this case, another option is to perform the $\overline{S}$ gate also via gate teleportation, using the resource state

$$|\overline{S}\rangle = \overline{S}|\overline{+}\rangle = \frac{1}{\sqrt{2}}(|\overline{0}\rangle + \mathrm{i}|\overline{1}\rangle)\,. \tag{5}$$

The conditional correction required when teleporting $\overline{S}$ is the gate $\overline{S}\,\overline{X}\,\overline{S}^\dagger\overline{X} \propto \overline{S}^2 = \overline{Z}$. While implementing the Pauli $\overline{Z}$ gate is typically easy for FTQC schemes, it could in principle also be implemented by teleporting the resource state

$$|\overline{-}\rangle = \overline{Z}|\overline{+}\rangle = \frac{1}{\sqrt{2}}(|\overline{0}\rangle - |\overline{1}\rangle)\,. \tag{6}$$

Teleporting the $|\overline{-}\rangle$ state requires no correction, regardless of the measurement outcome, since $\overline{Z}\,\overline{X}\,\overline{Z}\,\overline{X} \propto \overline{\mathbb{I}}$, where $\overline{\mathbb{I}}$ is the (logical) identity operator. Following this strategy, we can write the following circuit, which implements the $\overline{T}$ gate via three successive teleportations, where the second and third teleportations are applied only if all prior measurement outcomes are 1.

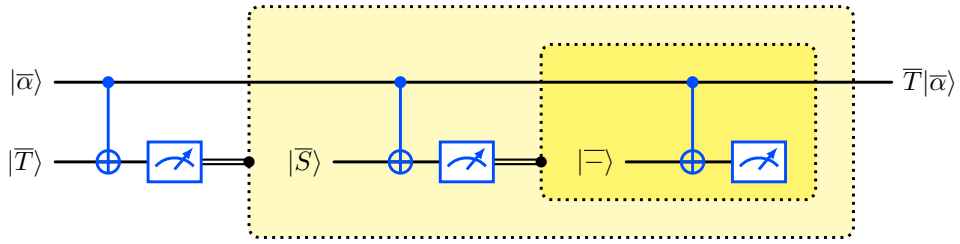

(7)

This approach may strike the reader as unnecessary, but designing the procedure in this iterative way will mirror the structure of our full protocol for QRAM.

2.2 Teleportable gates and the Clifford hierarchy

Not all gates can be teleported in the manner of circuit (4). The key reason it works is that the correction operation, $\overline{S}$, is a Clifford gate. In general, if one attempts to teleport a diagonal single-qubit logical gate $\overline{G}$, the conditional correction is $\overline{G}\,\overline{X}\,\overline{G}^\dagger\,\overline{X}$. Early work on gate teleportation [58] characterized the set of teleportable gates. It identified a hierarchy of teleportable gates known as the Clifford hierarchy. Focusing here on logical gates for consistency with the rest of this section, the logical Clifford hierarchy is a sequence of sets $\mathcal{C}_k$ for $k = 1, 2, \ldots$, where $\mathcal{C}_1$ is the set of logical Pauli gates, and $\mathcal{C}_k$ is defined recursively by

$$\mathcal{C}_k = \{\overline{G} \colon \overline{G}\,\overline{P}\,\overline{G}^\dagger \in \mathcal{C}_{k-1} \text{ for all } \overline{P} \in \mathcal{C}_1\}\,. \tag{8}$$

That is, the k-th level of the Clifford hierarchy are gates that, under conjugation, transform Pauli gates into gates in the $(k-1)$-th level. We may recognize $\mathcal{C}_2$ as the set of gates that transform Paulis to Paulis—that is, the set of

[6] For some codes, such as the color code [53], the $\overline{S}$ gate is transversal; for the surface code, however, it is fold-transversal [54, 55], and therefore more challenging to implement. One may also prefer to use autocorrected gadgets [56, 57] that avoid direct implementation of $\overline{S}$.

Clifford gates. The $\overline{T}$ gate lies in $\mathcal{C}_3$ because $\overline{T}\,\overline{X}\,\overline{T}^\dagger = \mathrm{e}^{-\mathrm{i}\pi/4}\overline{S}\,\overline{X}$ is Clifford, $\overline{T}\,\overline{Y}\,\overline{T}^\dagger = \mathrm{e}^{-\mathrm{i}\pi/4}\overline{S}\,\overline{Y}$ is Clifford, and $\overline{T}\,\overline{Z}\,\overline{T}^\dagger = \overline{Z}$ is Pauli (and therefore Clifford).

Focusing here on single-qubit diagonal gates, if a gate $\overline{G}$ lies in $\mathcal{C}_k$, then the teleportation procedure calls to use the state $\overline{G}|\overline{+}\rangle$ as a resource state. The conditional correction $\overline{G}\,\overline{X}\,\overline{G}^\dagger\,\overline{X}$ is also diagonal and lies in $\mathcal{C}_{k-1}$. As pointed out already in Ref. [58], this immediately yields a recursive procedure for implementing any gate in $\mathcal{C}_k$: prepare $\overline{G}|\overline{+}\rangle$; teleport; if outcome 1 is obtained, classically compute the required correction $\overline{G}' = \overline{G}\,\overline{X}\,\overline{G}^\dagger\,\overline{X} \in \mathcal{C}_{k-1}$; prepare $\overline{G}'|\overline{+}\rangle$; teleport; if outcome 1 is obtained, compute the required correction $\overline{G}'' = \overline{G}'\,\overline{X}\,\overline{G}'^\dagger\,\overline{X} \in \mathcal{C}_{k-2}$, etc. Each correction is one level lower in the hierarchy than the last. After enough rounds, no further correction will be required, as in circuit (7).

2.3 Teleporting the QRAM gate

The teleportation strategy for single-qubit diagonal gates can also be applied to multi-qubit diagonal gates, such as the QRAM unitary $\overline{V(f)}$ from Eq. (1). We define *QRAM resource states* analogously to the resource state $|T\rangle$ (cf. Eq. (2)).

$$\text{QRAM resource state:} \qquad |\Psi(f)\rangle = V(f)|+\rangle^{\otimes n} = \frac{1}{\sqrt{2^n}} \sum_{x\in\{0,1\}^n} (-1)^{f(x)}|x\rangle\,. \tag{9}$$

We denote the encoded logical QRAM resource state by $|\overline{\Psi(f)}\rangle$.

Assuming that we can prepare the encoded resource state $|\overline{\Psi(f)}\rangle$, then we may teleport the QRAM gate into an arbitrary encoded n-qubit state $\sum_x \alpha_x|\overline{x}\rangle$ by making an entangled measurement, as in the following circuit (cf. circuit (3)).

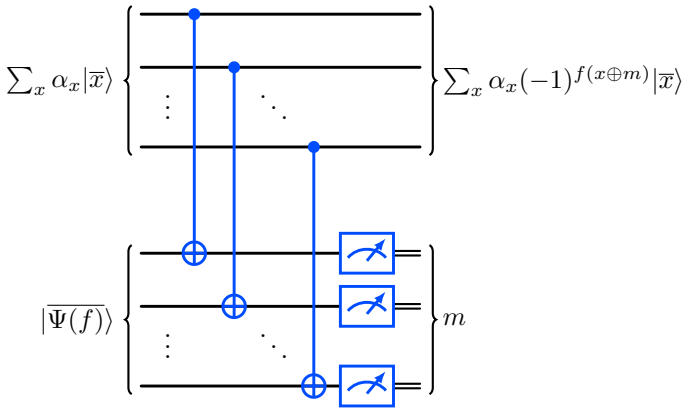

(10)

Here we emphasize that in the context of our protocol, the circuit is a logical circuit: all qubits are logical qubits, and both the upper and lower sets of n logical qubits are constructed out of $n' > n$ physical qubits using some QEC code. For example, one could choose to encode each logical qubit into its own $d \times d$ surface code patch, giving $n' = nd^2$ for that example. The n logical CNOT gates and n logical single-qubit measurements in circuit (10) are performed fault tolerantly within this code, allowing us to neglect the chance of logical error.

Each possible n-bit measurement outcome $m \in \{0,1\}^n$ is obtained with uniform probability $1/2^n$, regardless of the state $\sum_x \alpha_x|\overline{x}\rangle$. If $m = 0^n$ is obtained, then by direct calculation (see Section 4.5), we can verify that the gate $\overline{V(f)}$ has been correctly applied, yielding the state $\sum_x \alpha_x(-1)^{f(x)}|\overline{x}\rangle$. However, most of the time, we obtain a nonzero measurement outcome $m \neq 0^n$, in which case the phase $(-1)^{f(x)}$ is applied onto the basis state $|\overline{x \oplus m}\rangle$ rather than $|\overline{x}\rangle$, yielding the state $\sum_x \alpha_x(-1)^{f(x\oplus m)}|\overline{x}\rangle$. Here and throughout, $\oplus$ denotes bitwise addition, modulo 2.

To correct for this, we need to apply the phase $(-1)^{f(x\oplus m)\oplus f(x)}$ onto the basis state $|\overline{x}\rangle$ for each x; that is, we need to implement the correction operation $\overline{V(f')}$, where f' is a Boolean function defined by the rule

$$f'(x) = f(x) \oplus f(x \oplus m)\,. \tag{11}$$

The function f' depends on m, and thus it can only be determined after the teleportation of $|\overline{\Psi(f)}\rangle$ has been performed. To tie back to the case of a single-qubit diagonal gate $\overline{G}$ discussed in Section 2.2, where the conditional

correction was $\overline{G}' = \overline{G}\,\overline{X}\,\overline{G}^\dagger\,\overline{X}$, we can note that $\overline{V(f)}^\dagger = \overline{V(f)}$ and rewrite

$$\overline{V(f')} = \overline{V(f)}\,\overline{X}^m\,\overline{V(f)}^\dagger\,\overline{X}^m\,, \tag{12}$$

where X^m denotes the n-qubit Pauli operator with Pauli-X in positions where $m_i = 1$ and identity operator $\mathbb{I}$ in positions where $m_i = 0$, such that $\overline{X}^m|\overline{x}\rangle = |\overline{x \oplus m}\rangle$.

It now suffices to observe that for any f, the n-qubit unitary $\overline{V(f)}$ is in the n-th level of the logical Clifford hierarchy [8, 46]; we provide a self-contained proof of this in Appendix D. This guarantees that the correction $\overline{V(f')}$ will be in the $(n-1)$-th level. As explained in Appendix D, the reason this holds is related to the degree of the Boolean functions f and f', when they are expanded as a polynomial of their n input bits. Specifically, we may observe that the highest-degree monomials in the expansion of $f(x)$ are the same as those in the expansion of $f(x \oplus m)$. Thus, when $f'(x)$ is defined as $f(x) \oplus f(x \oplus m)$, the highest-degree monomials all cancel out, leaving only monomials of a lower degree. That is, the degree of f' is smaller than the degree of f by at least one. We use this fact to prove in general that if a Boolean function h has degree d, then $\overline{V(h)} \in \mathcal{C}_d$. In particular, since $\overline{V(f)} \in \mathcal{C}_n$ (the maximum possible degree of any function is n), we have that $\overline{V(f')} \in \mathcal{C}_{n-1}$.

Our protocol proposes to implement the correction $\overline{V(f')}$ in the same fashion as $\overline{V(f)}$: by preparing the resource state $|\overline{\Psi(f')}\rangle$ and teleporting as in circuit (10). This will also produce a correction, associated with a Boolean function f'' of degree $n-2$. As we iterate, we descend the Clifford hierarchy, and the degree of our correction function is reduced. Once we have performed n rounds of teleportation, our correction function has degree zero. If a Boolean function h is constant, this implies that $\overline{V(h)} \propto \overline{\mathbb{I}}$; thus, once we have reduced the correction function to degree zero, we may cease iterating the protocol.

Later, in Section 4.5, we perform a more complete analysis of the teleportation channel; for example, we quantify the error in the teleportation channel when an imperfect resource state is teleported instead of $|\overline{\Psi(f)}\rangle$.

2.4 Preparing the encoded QRAM resource state

The analysis above shows how we can implement the logical $\overline{V(f)}$ gate, provided that we can adaptively prepare the resource states $|\overline{\Psi(g)}\rangle$, up to low error, for any particular Boolean function g. At first glance, this seems like a tall task. There are 2^{2^n} different states that we may need to prepare. By a simple counting argument, the quantum circuit complexity of at least one of these states is at least $\Omega(2^n/n)$. The innovation of our protocol is to outsource this complexity to a single-purpose, faulty (and ideally passive) QRAM device, which may be able to exploit the unique structure of QRAM to implement $V(g)$ cheaply, but imperfectly.

We propose a three-step procedure for preparing these states, analogous to the preparation of the $|\overline{T}\rangle$ state: physical preparation, encoding, and distillation.

- **Physical preparation**: we assume that we have access to a QRAM device that can implement an approximation to $V(g)$ at the physical level, as discussed in Section 1 and Fig. 1. By running this device on the initial input state $|+\rangle^{\otimes n}$, we produce the physical resource state $|\Psi(g)\rangle$ of Eq. (9). The device can be faulty. In fact, our protocol can succeed as long as the device produces states that have at least $1/\mathrm{poly}(n)$ minimum fidelity with respect to $|\Psi(g)\rangle$.
- **Encoding**: The physical n-qubit state is not protected by a QEC code, and thus it is vulnerable to error. We immediately encode it into (an approximation of) the logical state $|\overline{\Psi(g)}\rangle$ using some number $n' > n$ of physical qubits on our main quantum processor. This step incurs some additional logical error because encoding arbitrary states is not fully fault tolerant. However, for topological codes like the surface code, there exist effective methods for encoding a physical qubit into a logical qubit [49]. The logical error due to encoding is $O(p)$—independent of the code distance—where p is the physical error rate. In Section 4.2, we use the general results of Ref. [59] to formalize the error in this step. Since the state $|\overline{\Psi(g)}\rangle$ is an n-qubit state, we expect the total logical error incurred from encoding to be $O(np)$, although for the case of general codes, we can only show $O(n\sqrt{p})$. The physical error rate must be $p = O(1/n)$ or $p = O(1/n^2)$, so that the total error from encoding remains $O(1)$, but for relevant sizes of n (e.g., $n = 43$ already corresponds to one terabyte of QRAM), the $p = O(1/n)$ condition is already met on devices that exist today.
- **Distillation**: Distillation procedures [51, 52] for the $|\overline{T}\rangle$ (or $|\overline{\mathrm{CCZ}}\rangle$) state leverage the existence of QEC codes where T (or CCZ) is transversal. The overhead, that is, the number of noisy copies of $|\overline{T}\rangle$ needed to distill one $\varepsilon_{\mathrm{dist}}$-good copy of $|\overline{T}\rangle$ is $\mathrm{polylog}(1/\varepsilon_{\mathrm{dist}})$, and this can be improved to $O(1)$ overhead using high-rate codes [60–62]. For the $V(g)$ gate, we do not know of any suitable codes that would enable this kind of approach. However, we can still distill $|\overline{\Psi(g)}\rangle$ using state-agnostic *quantum purity amplification* methods [63–

69], which take many copies of an arbitrary mixed state $\overline{\rho}$ and produce one $\varepsilon_{\text{dist}}$-good copy of the pure state $|\overline{\Xi}\rangle\langle\overline{\Xi}|$, where $|\overline{\Xi}\rangle$ is the principal component (i.e., top eigenvector) of $\overline{\rho}$. These methods do not leverage or learn any properties of $|\overline{\Xi}\rangle$, and it is known that the optimal overhead achievable in such settings is $\Theta(1/\varepsilon_{\text{dist}})$ [68]. In Section 4.4, we discuss several specific state-agnostic approaches. We first consider the iterated swap test purification method studied in Refs. [66, 67, 69, 70], which is appealing for its simplicity. In the regime where the physical preparation and encoding steps prepare states with high (but still imperfect) fidelity, the iterated swap test approach is nearly optimal. On the other hand, as the fidelity of the undistilled input states decreases, the overhead of the iterated swap test rapidly increases, scaling exponentially in the inverse input fidelity. To alleviate this issue, we propose a new gate-efficient state-agnostic quantum purity amplification procedure based on quantum principal component analysis [71, 72], which achieves nearly optimal sample complexity even in the regime of low input fidelity, while still being compatible with the streaming model (i.e., where the undistilled input states are processed one at a time, rather than all at once as in the known sample-optimal protocol [69]).

To apply these state-agnostic distillation approaches within our protocol, it must be the case that the state that is output by the physical preparation and encoding processes has the ideal resource state $|\overline{\Psi(g)}\rangle$ as its principal component. Evaluating this assertion requires specifying a noise model in our abstract QRAM device and in our main quantum processor. We suppose that our main processor is subject to circuit-level stochastic noise. For the QRAM device, the only assumption we make is that the noise is independent of the dataset, in the sense that, for dataset g, it enacts the n-qubit quantum channel $\mathcal{N}_2 \circ \mathcal{V}(g) \circ \mathcal{N}_1$, where $\mathcal{N}_{1,2}$ are g-independent noise channels, and $\mathcal{V}(g) = V(g)[\cdot]V(g)^\dagger$ is the ideal QRAM channel. In this case, we can ensure that the principal component of the state we prepare is $|\overline{\Psi(g)}\rangle$ by performing a *partial Clifford-twirl* of the unitary $V(g)$. This method leverages the fact that for any Clifford circuit C formed from Z, X, CZ, and CX (i.e., CNOT) gates, we have $|\overline{\Psi(g)}\rangle = C|\overline{\Psi(g_C)}\rangle$ for some dataset g_C; the idea is to choose a random C, compute the dataset g_C, query g_C with the QRAM device, and then apply $\overline{C}$ fault-tolerantly to restore $|\overline{\Psi(g)}\rangle$. Partial Clifford twirling is not necessary under the stronger assumption that the noise in the QRAM device naturally guarantees that the ideal resource state is the principal component.

It is important that our protocol can work even when the QRAM device has low (at least inverse polynomial) fidelity. Given the engineering challenges associated with building a reliable physical QRAM device, it is much easier to imagine realizing our protocol in practice, especially as n grows, if the physical QRAM device need only have a small correlation with the correct output. Along these lines, another key benefit of a distillation–teleportation approach to fault-tolerant QRAM is that one always has the option to restart the preparation, encoding, and distillation procedure if an error is detected. For instance, if the physical QRAM device recognizes certain errors (e.g., photon loss), one can simply postselect on these events not occurring, improving the effective fidelity of the device from the perspective of our protocol.

2.5 Full summary of protocol and statement of results

To summarize, our main result is a protocol for implementing the logical QRAM operation $\overline{V(f)}$, up to arbitrarily high fidelity, using many queries to a device that can perform the physical QRAM operation $V(g)$ (for any/all g) with a lower nonzero fidelity. As discussed in Section 1, the physical QRAM operation could be accomplished with a dedicated subcomponent of the larger quantum device specialized for QRAM, which need not be capable of universal fault-tolerant quantum computation.

The protocol to implement $\overline{V(f)}$ cycles at most n times through five steps discussed in the previous subsections: (i) physical preparation, (ii) encoding, (iii) distillation, (iv) teleportation, and (v) adaptive classical computation of the correction. Step (v) uses measurement outcomes from step (iv) to transform the dataset according to the *update rule* (UR) of Eq. (11), prior to returning to step (i). The entire protocol is depicted in Fig. 2, where each of the five steps is shown in a different color. A more detailed specification and formal error analysis of each step is provided in Section 4. We arrive at the following statement of the cost of implementing $\overline{V(f)}$.

Theorem 1 (Main result (informal))**.** *For any data table f with 2^n entries, and any error parameter $\varepsilon > 0$, the protocol performs the logical QRAM operation $\overline{V(f)}$ up to error ε (in diamond distance), under the assumption that the physical QRAM device implementing physical $V(f)$ has noise independent of f. The quantum resources required are:*

- $\text{poly}(n)/\varepsilon$ *calls to a device that performs the physical $V(g)$, for various g (determined adaptively) with any nonzero minimum fidelity $F \geq 1/\text{poly}(n)$.*
- $\text{poly}(n)/\varepsilon$ *calls to a* $\text{poly}(n)$*-cost fault-tolerant encoding procedure that encodes n-qubit physical states into a suitable QEC code capable of FTQC, while incurring at most $O(1)$ logical error.*

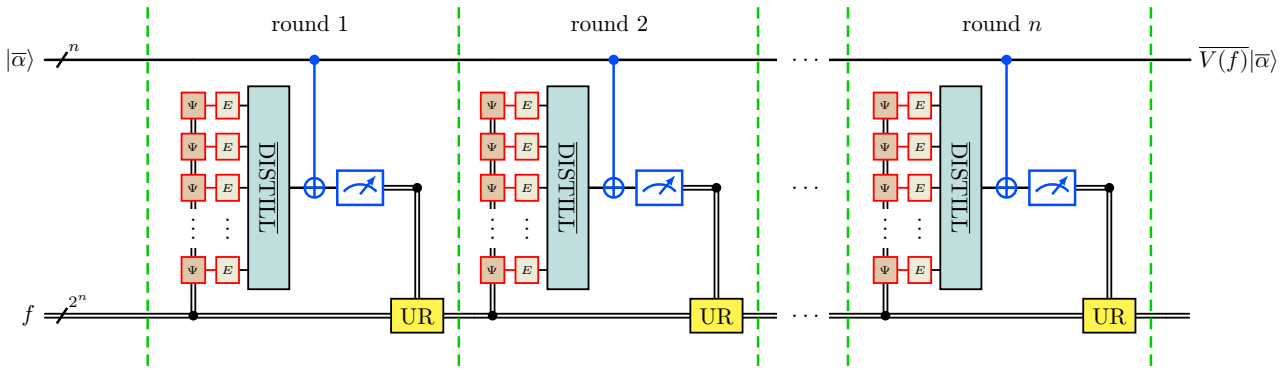

Figure 2: Quantum circuit depiction of the protocol for implementing the logical diagonal QRAM operation $\overline{V(f)}$ (see Eq. (1)) fault tolerantly for a data table f. The protocol cycles through n rounds, and each round has five steps: preparation, encoding, distillation, teleportation, and classical update, depicted in different colors. All gates are fault-tolerant, logical gates, except the query to the noisy QRAM device Ψ and the encoding step E, outlined in red. The solid black wires represent encoded logical quantum registers of n logical qubits, the red wires represent unencoded quantum registers of n physical qubits, and the double black lines represent classical registers. For simplicity, we have not depicted the twirling step in the figure (which may not be necessary in practice), where prior to each application of Ψ, the dataset is modified by an independently chosen random transformation, which is corrected for after E has been applied; see circuit (42).

- $\mathrm{poly}(n)/\varepsilon$ *fault-tolerant one- and two-qubit logical gates, single-qubit logical* $|\overline{0}\rangle$ *state preparations, and single-qubit logical measurements.*

The classical resources required are:

- n *applications of the classical update rule, each of which has* $O(2^n)$ *complexity in a standard RAM model.*
- $\mathrm{poly}(n)/\varepsilon$ *twirling operations on the dataset, each of which has complexity* $\mathrm{poly}(n)2^n$ *in a standard RAM model.*

After each update rule and twirling operation, the physical QRAM device must be "reloaded" or otherwise given access to the updated classical dataset.

The main implication of this result is the following: suppose that a quantum algorithm calls the QRAM operation $V(f)$ at most $T = \mathrm{poly}(n)$ times, and suppose that one has access to a QRAM device that approximately performs the physical QRAM operation with at least $1/\mathrm{poly}(n)$ fidelity, at computational cost $\mathrm{poly}(n)$ (similar to the cost of RAM). Then, one may take $1/\varepsilon = O(T) = \mathrm{poly}(n)$, and conclude that the algorithm can be implemented fault-tolerantly using only $\mathrm{poly}(n)$ quantum resources. As a result, our protocol provides a step toward justifying the cheap QRAM assumption, and it provides a method of fault-tolerantly implementing quantum algorithms that depend on QRAM.

Even in a cost model where noisy physical QRAM incurs computational cost $\Omega(2^n)$—for instance, if the physical QRAM has $\Omega(2^n)$ active gates each requiring $\Omega(1)$ energy input—our protocol still provides the benefit that the exponential quantum complexity is contained entirely to physical quantum operations that can be optimized specifically to perform QRAM. There is no need for an exponential amount of QEC and the associated overheads it incurs.

2.5.1 Caveat: classical complexity

The main caveat of our protocol is that it requires a non-negligible amount of purely classical adaptive computation. In particular, after receiving random measurement outcome m, the protocol requires replacing the value $g(x)$ with the value $g(x) \oplus g(x \oplus m)$ for all 2^n addresses x of the dataset, as in Eq. (11). While computing the new value is easy for any individual x, the sheer number of different x means the complexity—in terms of classical circuit size or RAM calls to the dataset g—is at least $\Omega(2^n)$.

However, we must recall that naively, the logical QRAM requires $\Omega(2^n)$ fault-tolerant quantum resources, if implemented as a fault-tolerant circuit. One unit of fault-tolerant quantum resources, such as one fault-tolerant quantum gate, is expected to be several orders of magnitude more expensive in terms of both financial cost and computational runtime than one unit of classical computation, such as a classical gate or floating point operation [73]. Thus, trading $\Omega(2^n)$ quantum for $\Omega(2^n)$ classical resources may lead to an overall cheaper and faster computation.

Furthermore, we expect that although it formally has $\Omega(2^n)$ computational cost, the complexity of the classical update rule has significantly better constant prefactors than the classical computation required to power active QEC of an entire QRAM circuit. As mentioned previously, circuit QRAM with $\mathrm{poly}(n)$ latency would require a fault-tolerant quantum computer with $\Omega(2^n)$ logical qubits. Such a device would likely require $\Omega(2^n)$ full-fledged *classical* chips to be co-located with the logical qubits, in order to process in parallel the QEC syndrome data generated by the computation in real time. For example, in the surface code operating at a 1 MHz QEC cycle rate, the amount of syndrome data generated by 2^{20} logical qubits, each encoded into its own patch at code distance 11, would be more than 15 terabytes per second. Specialized classical decoding algorithms must be run continuously to identify and correct errors as they occur. In contrast, for a datset of size 2^{20} bits, the classical update rule in our protocol is a single structured transformation of a 120 kilobyte dataset. The only interaction between this dataset and the quantum processor is the reloading of the physical QRAM device with the updated dataset.

Additionally, because of the structure of the classical update rule $g(x) \mapsto g(x) \oplus g(x \oplus m)$, it is conceivable that dedicated classical chips could be built to parallelize the process of performing the update and the reloading of the QRAM device. In Section 5, we analyze the complexity of the update rule, and illustrate how it can be implemented with a classical circuit of depth $\mathrm{poly}(n)$, although embedding this circuit into 2 or 3 spatial dimensions leads to asymptotically growing wire density. We show how in a model of parallel computation, the update rule is equivalent (up to $\mathrm{poly}(n)$ factors) to performing sparse matrix-vector multiplication, with a particularly close connection to the (fast) Walsh–Hadamard transform. This fact helps to understand the expected difficulty of parallelization and it clarifies the opportunity cost of these classical resources.

2.5.2 Comments on scalability

An additional caveat is the fact that our protocol is likely not to be fully scalable for indefinitely large n. This stems from two aspects, the physical QRAM device, and the encoding step.

First, the physical size of an (ideally passive) physical QRAM device would need to grow as $\Omega(2^n)$, yet we need it to produce physical QRAM resource states with at least $1/\mathrm{poly}(n)$ fidelity. Thus, the fidelity of the individual device components needs to improve as n grows. It is known that certain architectural approaches to QRAM, namely, bucket brigade QRAM, possess a certain noise resilience property: the overall infidelity of the physical QRAM operation scales as $O(qn^2)$ [36], where q is the error rate of the individual router components that compose the device. This is exponentially better than $O(q2^n)$, which would be the naive expectation, given the exponential number of error-prone routers in the device. This noise resilience is a crucial fact for the possibility of practical QRAM. If this kind of scaling is achieved, then the physical per-component error rate q must decrease asymptotically roughly as $O(1/n^2)$ to be useful for our protocol.

Second, the physical QRAM resource state is an n-qubit state, and in a noise model where each operation on our main quantum processor fails with probability p, the encoding of this n-qubit physical state into an n-qubit logical state necessarily incurs at least $\Omega(np)$ logical error. The formal analysis, later, shows how $O(n\sqrt{p})$ can be achieved regardless of the choice of QEC code. Either way, p must decrease as $O(1/n)$ or as $O(1/n^2)$ to keep this error of total size $O(1)$.

While any practical implementation of this protocol will certainly need to pay close attention to error rates at every step, this is not a hugely debilitating conceptual issue for QRAM. This is because we do not ever expect to need to build a QRAM device for very large values of n. For example, typical RAM devices in classical computers are of size roughly 10 gigabytes, corresponding to only $n = 36$. The physical error rate in state-of-the-art quantum devices in several different platforms already achieves p in the range of 10^{-3}–10^{-2}. Improving by roughly an order of magnitude to $p = q = 10^{-4}$ would be sufficient to enable our protocol at size $n = 36$, assuming that the dominant error contribution scales as $4qn^2$ (where the presumed constant prefactor of 4 is chosen to align with Ref. [36, Eq. (28)]).

2.5.3 Comments on applications

Our investigation has been primarily motivated by the goal of evaluating the viability of QRAM as a primitive for fault-tolerant quantum computation in an abstract sense. Nonetheless, in Section 6, we consider whether our protocol could provide a practical advantage over alternative methods in several concrete applications. Generally, although our protocol achieves asymptotically polynomial $\mathrm{poly}(n)/\varepsilon$ complexity, we find that this version of the protocol struggles to provide an immediate advantage. For example, in quantum machine learning scenarios, the $\Omega(2^n)$ cost of the classical update rule makes it difficult to find examples where an end-to-end speedup persists over alternative classical methods. The observation in Section 5 that the update rule is similar in power to a sparse

matrix-vector multiplication clarifies that, in our search for superpolynomial quantum speedups, we must only target problems where the ability to perform classical $2^n \times 2^n$ sparse matrix-vector multiplications is not already sufficient to solve the problem in $\text{poly}(n)$ time, which considerably reduces the set of candidates. See Section 6.2 for comments on possible scenarios where this conclusion may be avoided.

On the other hand, in scenarios like quantum chemistry and cryptanalysis where QRAM is utilized—in that context often referred to as a "quantum lookup table" or "quantum read-only memory"—the $\Omega(2^n)$ classical cost is tolerable. In fact, in these instances, it is typically already being proposed to implement QRAM with a fully error-corrected quantum circuit of depth $\Omega(2^n)$ [74]. Our protocol could allow this exponential fault-tolerant complexity to be offloaded to a specialized physical QRAM device and a classical computer. However, in our preliminary resource analysis at relevant system sizes in Section 6, the $\text{poly}(n)/\varepsilon$ cost is still too large to provide an actual advantage. Part of the issue is that if the QRAM operation is called T times, one must take $1/\varepsilon = \Omega(T)$, and hence the total cost of implementing all T QRAM queries scales as T^2. The discovery of a distillation protocol for QRAM resource states with overhead $\text{polylog}(1/\varepsilon)$ instead of $1/\varepsilon$ would be extremely beneficial in this calculation.

2.5.4 Extension to multiple output bits

In Appendix A, we explain how our protocol can be straightforwardly extended to the case where b classical bits are stored at each of 2^n addresses, and one wishes to coherently read all b bits into a separate bus register. That is, we show how to fault-tolerantly perform the operation $\overline{U(f)}$ that implements $|\overline{x}\rangle|\overline{u}\rangle \mapsto |\overline{x}\rangle|\overline{u \oplus f(x)}\rangle$, with $f\colon \{0,1\}^n \to \{0,1\}^b$ here denoting a function with b output bits. The strategy is to observe that conjugating $U(f)$ by a Hadamard transform on the bus register yields a diagonal unitary with ± 1 on the diagonal that may be viewed as a generalization of $V(f)$ from Eq. (1). The unitary acts on $n+b$ qubits rather than n qubits, but importantly, the degree of the Boolean function is at most $n+1$, which can be much smaller than $n+b$. The protocol proceeds identically to how it is described in the main text, except that the resource states are larger, requiring $n+b$ qubits, which leads to greater gate complexity overhead when performing distillation and teleportation.

2.6 Relation to prior work

The idea of QRAM was first formalized by Giovannetti, Lloyd, and Maccone (GLM) in Refs. [7, 75] (although some primitive versions of QRAM had been sketched earlier, see e.g. Ref. [76, Chapter 6]). These works first introduced the idea of a dedicated QRAM hardware element—a device specially designed for QRAM and separate from the main quantum processor—by proposing implementations based on optical and atomic hardware. Many other proposals have followed, including proposals based on superconducting circuits [77–80], photonic systems [81, 82], and neutral atom arrays [83] (see Ref. [8] for a more detailed review). We highlight that notions of teleportation-based QRAM [81] and QRAM resource states [83]—albeit resource states of size exponential in n—have previously been proposed. Unfortunately, all of these proposed QRAM implementations face daunting practical challenges, and most are not passive, meaning they require active control over $\Omega(2^n)$ quantum components, which would undermine the cheap QRAM assumption. (As Ref. [8] notes, the proposal of Ref. [80] is a noteworthy example of a passive implementation, although it faces challenges of scalability and practicality.) To our knowledge, there has not yet been a proposed implementation of a physical QRAM device that is simultaneously practical, scalable, and passive.

The initial GLM QRAM papers also sparked a long-running debate about the practicality of QRAM and validity of the cheap QRAM assumption, especially in relation to error correction and fault tolerance. In particular, GLM proposed a specific QRAM architecture—the "bucket brigade" QRAM—that they argued was intrinsically robust to errors. This claim was initially met with some skepticism (see, e.g., Ref. [32]), but the robustness was later proven in Ref. [36], which showed that the overall error of a bucket-brigade QRAM query scaled only with $\text{poly}(n)$, despite the fact that the QRAM itself is comprised of $\Omega(2^n)$ error-prone components. This robustness is key to the viability of our own proposal, since the passive QRAM device in Fig. 1a could indeed have only moderate overall error rates compatible with our distillation–teleportation scheme.

Even with some intrinsic robustness against errors, QRAM is still likely to require QEC in most applications. As mentioned in Section 1, fault-tolerant implementations of QRAM based on circuit decompositions of the unitary $V(f)$ involve $\Omega(2^n)$ qubits and $\Omega(\sqrt{2^n})$ non-Clifford gates, and face serious questions of practicality at large-scales. To our knowledge, the survey of QRAM in Ref. [8] was the first to consider the possibility of achieving a fault-tolerant QRAM by a method other than circuit QRAM. They proved several no-go theorems that present barriers to finding a code where QRAM is transversal. They also proved a no-go theorem ruling out certain distillation–teleportation protocols. Specifically, they considered protocols that have a distillation phase that queries the physical QRAM gate $V(f)$ up to Q times to prepare a resource state $\chi(f)$, followed by a teleportation channel where $\chi(f)$ interacts with

an arbitrary state $|\overline{\alpha}\rangle\langle\overline{\alpha}|$ in an attempt to prepare $\overline{V(f)}|\overline{\alpha}\rangle\langle\overline{\alpha}|\overline{V(f)}$. They showed that for this setup, $Q \geq \Omega(2^{2n})$ queries are required to succeed with high fidelity on all possible choices of $|\overline{\alpha}\rangle\langle\overline{\alpha}|$. Translating their logic into our language, they observed that regardless of the protocol and the function f, one can always find an f' which differs from f at only a few addresses, but where the resource states $\chi(f)$ and $\chi(f')$ are $O(\sqrt{Q}/2^n)$-close. This implies that if $Q \ll 2^{2n}$, then the teleportation channel cannot produce well-distinguishable outputs when f is queried compared to when f' is queried (data processing inequality). Yet, if we suppose that f and f' differ at even one address j while agreeing on some other address k, then when $|\overline{\alpha}\rangle = \frac{1}{\sqrt{2}}(|\overline{j}\rangle + |\overline{k}\rangle)$, the $\overline{V(f)}$ and $\overline{V(f')}$ gates should lead to distinguishable orthogonal states, a contradiction.

Each of the n rounds within our protocol individually fits into the framework of the no-go theorem of Ref. [8]. Our protocol circumvents this result because it adaptively updates the QRAM function being queried in each round, based on measurement outcomes obtained in prior rounds. If two functions f and f' are different, even at a single address, there will be at least one round where the resource states being teleported by our protocol are far away from each other; in fact, if f and f' differ at exactly one address in round $r = 1$, then they will differ at exactly 2^{r-1} addresses in each round $r = 2, 3, 4, \ldots, n$, provided that all of the n-bit random measurement outcomes obtained up until round r form a linearly independent set.

Certain elements of our protocol also connect with prior work outside the context of QRAM. For example, the task of quantum purity amplification has been extensively studied, and we comment more on this in Section 4.4. Additionally, while the n-qubit states $|\Psi(f)\rangle$ from Eq. (9) have—to the best of our knowledge—not previously been proposed as QRAM resource states, they have been utilized in other contexts, often by the name of "phase states." For example, phase states have been studied as pseudorandom quantum states in the context of cryptography [84, 85], as targets for quantum state tomography [86], and as a mechanism for showing search-to-decision reductions in quantum complexity theory [70].

3 Error models and physical protocol requirements

The theory of FTQC shows how a quantum processor can perform an arbitrary quantum computation through a sequence of noisy physical operations on a set of physical qubits, provided that the physical noise is sufficiently uncorrelated and its rate p is below a constant threshold [87–89]. Our protocol augments this by assuming that we also have access to a physical QRAM device that can perform an approximate $V(g)$ gate on n physical qubits, for any function g, as illustrated in Fig. 1b. We require that g can be modified from one query to the next via classical communication with the device. Also, we require that the quantum information contained in the n physical qubits output by the device can be transported into n physical qubits on the main quantum processor without significant degradation of its fidelity, whether by physically moving the qubits output by the device to the main processor, or by some other means.

In this section, we specify the noise models we consider for each of these two components, and we define the setup for the FTQC part of our protocol on the main processor. The protocol is agnostic to many of the details here, including which QEC code is used, and we attempt to keep it as general as possible.

3.1 Error model of the physical QRAM device

Without a more concrete implementation in mind, we cannot fully model the noise in the device. However, we consider a general noise model that assumes only that the noise is independent of g, the function being queried. This noise model was also employed for some of the results in Ref. [8].

Definition 1 (Dataset-independent QRAM noise)**.** *A physical QRAM device that implements channel $\widetilde{\mathcal{V}}(g)$ on input g is said to have dataset-independent noise if there exists $\mathcal{N}_1$ and $\mathcal{N}_2$ independent of g for which*

$$\widetilde{\mathcal{V}}(g) = \mathcal{N}_2 \circ \mathcal{V}(g) \circ \mathcal{N}_1 \,, \tag{13}$$

where $\mathcal{V}(g) = V(g)[\cdot]V(g)^\dagger$ is the ideal unitary channel.

We defend the plausibility of this noise model by appealing to the presumed structure of shallow-depth physical QRAM implementations. One imagines that the bits of classical memory corresponding to the values $g(0)$, $g(1)$,…, $g(2^n - 1)$ are distributed in memory cells over 1D or 2D space—for example, the illustration in Fig. 1a distributes them in 1D space. As discussed in Ref. [8], at a high level, a poly(n)-depth QRAM implementation requires a routing step, a readout step, and an unrouting step. In the routing step, the n-qubit address information $|x\rangle$ is used to (coherently) activate a path to the memory cell corresponding to address x. In the readout step, a qubit must

interact with the classical bit of information $g(x)$ stored in that cell, gaining a -1 phase if and only if $g(x) = 1$. Then, in the unrouting step, the activated routers must be coherently reset before being traced out to ensure the overall operation maintains coherence between different $|x\rangle$.

The important fact to notice is that the routing and unrouting steps are not at all dependent on the dataset g. Thus, to justify the validity of Definition 1 in this model, any noise that occurs during the routing step could be propagated backward to the beginning of the circuit and contribute to $\mathcal{N}_1$, while any noise in the unrouting step could be propagated forward to the end of the circuit and contribute to $\mathcal{N}_2$.

The only step that can be dependent on g is the readout step, but this step is generally considered to be simpler to implement than the routing step [8]. For example, in the bucket-brigade QRAM circuit of Ref. [36, Figure 10], the readout step is performed in a single circuit layer by a set of parallel single-qubit Pauli-X gates: at memory cell x, an X gate is applied if $g(x) = 1$, and an identity gate $\mathbb{I}$ is applied if $g(x) = 0$. (This classically controlled X gate would ideally be applied passively via interaction with a non-volatile memory storing $g(x)$.) Since the identity $\mathbb{I}$ and Pauli-X gates are simple, it may be plausible that, in some implementations, the noise can be independent of which of them is applied, justifying Definition 1 for the physical QRAM device. Furthermore, we note that even if the noise is not identical for the $\mathbb{I}$ and X gates, it is plausible that the dataset-independent noise property *effectively* holds in the context of our protocol, thanks to the protocol's *partial Clifford twirling* (Section 4.3) that effectively randomizes the data being queried—intuitively this randomization should remove dependence of g from the average noise channel.

We leave to future work the task of more rigorously showing that certain microscopic (i.e., component-level) noise models lead to dataset-independent noise in the form of Definition 1. However, we note that the set of noise processes that fall into this category can include some counterintuitive members. For example, we may suppose that a physical QRAM device is constructed from a depth-n binary tree of routing elements, but that one of the routers in the tree is "dead." The QRAM attempts to activate a path through the tree to a particular address x at one of the leaves, and if this path passes through the dead router, it causes catastrophic failure of the device, leading the device to instead output the maximally mixed state $\mathbb{I}/2^n$. Let $\mathcal{X} \subset \{0,1\}^n$ denote the set of addresses that cause the dead router to activate when they are queried, and let $\Pi_{\mathcal{X}} = \sum_{x\in\mathcal{X}} |x\rangle\langle x|$ be the projector onto these address states. Then, we may write the channel implemented by the noisy device as

$$\widetilde{\mathcal{V}}(g)[\rho] = V(g)(\mathbb{I} - \Pi_{\mathcal{X}})\rho(\mathbb{I} - \Pi_{\mathcal{X}})V(g)^\dagger + \mathrm{tr}(\Pi_{\mathcal{X}}\rho)\frac{\mathbb{I}}{2^n} \tag{14}$$

$$= (\mathbb{I} - \Pi_{\mathcal{X}})V(g)\rho V(g)^\dagger(\mathbb{I} - \Pi_{\mathcal{X}}) + \mathrm{tr}(\Pi_{\mathcal{X}}\rho)\frac{\mathbb{I}}{2^n}\,, \tag{15}$$

where the second equality follows since $V(g)$ is a diagonal unitary, and thus it commutes with the diagonal projector $\mathbb{I} - \Pi_{\mathcal{X}}$. We may then rewrite the noisy channel $\widetilde{\mathcal{V}}(g)$ in a dataset-independent fashion as $\widetilde{\mathcal{V}}(g) = \mathcal{N}_2 \circ \mathcal{V}(g) \circ \mathcal{N}_1$, with $\mathcal{N}_1 = \mathcal{I}$ (the identity channel) and

$$\mathcal{N}_2[\rho] = (\mathbb{I} - \Pi_{\mathcal{X}})\rho(\mathbb{I} - \Pi_{\mathcal{X}}) + \mathrm{tr}(\Pi_{\mathcal{X}}\rho)\frac{\mathbb{I}}{2^n}\,. \tag{16}$$

Furthermore, in the context of our protocol, the device is always queried on the input state $\rho = |+\rangle\langle+|^{\otimes n}$. Thus, this particular noise process has fidelity given by

$$\mathrm{tr}\Big(|\Psi(g)\rangle\langle\Psi(g)|\widetilde{\mathcal{V}}(g)[|+\rangle\langle+|^{\otimes n}]\Big) = \frac{(2^n - |\mathcal{X}|)(2^n - 1 - |\mathcal{X}|)}{2^{2n}} + \frac{1}{2^n} \geq 1 - \frac{2|\mathcal{X}|}{2^n}\,. \tag{17}$$

That is, if the dead router is deep in the binary tree and $|\mathcal{X}| \ll 2^n$, then the fidelity of the device remains close to 1.

The fact that our protocol can work in such a scenario is counterintuitive because the dead router would seem to completely block access to the information $g(x)$ for any $x \in \mathcal{X}$ and thus make it impossible to correctly implement the QRAM when the address register has high overlap with the support of $\Pi_{\mathcal{X}}$. As we will show in Section 4.3, our protocol handles this with partial Clifford twirling, a technique that scrambles the dataset g into a new dataset g_C so that the information $g(x)$ is contained in $g_C(y)$, and the location y is uniformly random, and in particular, the probability that $y \in \mathcal{X}$ is $|\mathcal{X}|/2^n$. Thus, for every x, the dead router only compromises the information $g(x)$ with small probability, even for $x \in \mathcal{X}$.

3.2 Error model of the main quantum processor

Our main quantum processor acts on a set of noisy physical qubits with a quantum circuit, that is, a sequence of noisy physical quantum operations including initializations, 1- and 2-qubit gates, and measurements, as well as classical computation and adaptive classical feedback. We will assume that our main processor is subject to circuit-level stochastic noise, defined below.

Definition 2 (Circuit-level stochastic noise, Section 2.5 of Ref. [59])**.** *A physical implementation of a quantum circuit V is said to be subject to parameter-p stochastic noise if the following holds.*

- *Purely classical components are implemented perfectly without any errors.*
- *Each quantum component $\mathcal{P}$, including (classically controlled-)gates, qubit initializations, measurements, is realized by $\widetilde{\mathcal{P}} = (1-p)\mathcal{P} + p\mathcal{N}_{\mathcal{P}}$, where $\mathcal{N}_{\mathcal{P}}$ is a quantum channel of the same input and output registers as $\mathcal{P}$.*

Circuit-level stochastic noise is a special case of the more standard model called local-stochastic noise model [89], which additionally allows some correlations between the gate faults. We choose circuit-level stochastic noise over local stochastic noise because this allow us to analyze a fault-tolerant logical state preparation procedure using the results of Ref. [59], for which the theorems require independent noise. However, we will require that the QEC code and FTQC scheme are able to correct against the more general local stochastic noise.

3.3 Fault-tolerant quantum computation

Our protocol is agnostic to which QEC code family and FTQC scheme is utilized, as long as it is capable of a universal set of fault-tolerant logical gates. Specifically, there must be a nonzero threshold p_0 such that, if the processor is subject to local stochastic noise with error parameter $p < p_0$, then for any target error rate and any logical quantum circuit, one can choose the code parameters (e.g., distance) large enough to ensure the ideal logical circuit is simulated up to the required error [87, 89–93].

As we wish to enact the logical n-qubit QRAM operation of Eq. (1), we assume we have chosen some QEC code family that encodes n logical qubits into some number $n' > n$ of physical qubits, along with a scheme for performing fault-tolerant gates. Regardless of our choice, we can identify the following ingredients.

- Encoding: There is an encoding isometry E, which maps any n-qubit physical state $|\psi\rangle$ to its associated encoded n-qubit logical state $|\overline{\psi}\rangle$ (of n' physical qubits). The map E is injective and its image is the codespace of the code. We let $\mathcal{E}$ denote the corresponding quantum channel $E[\cdot]E^\dagger$ that encodes density matrices. Our protocol will also require a fault-tolerant encoding gadget $\mathcal{E}_{\mathrm{FT}}$, which implements $\mathcal{E}$ in the absence of noise, but is constructed in such a way that its noisy implementation $\widetilde{\mathcal{E}}_{\mathrm{FT}}$ is resilient to errors (see Section 4.2).
- QEC: There is a QEC projector $\mathcal{Q}$ that maps states outside of the codespace to states in the codespace, that is, a map that detects and corrects physical errors on encoded states. In FTQC, the projector $\mathcal{Q}$ is implemented with a fault-tolerant QEC gadget, denoted $\mathcal{Q}_{\mathrm{FT}}$, which enacts the map $\mathcal{Q}$ in the absence of noise and a map $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$ in the presence of noise. The gadget $\mathcal{Q}_{\mathrm{FT}}$ is applied after each location in the logical circuit to prevent the buildup and propagation of physical errors; it must satisfy certain formal properties to guarantee the existence of a threshold; see Ref. [90].
- Gates: For each physical gate G, we let $\overline{G}$ denote the associated logical operation on the codespace of the QEC code, and we let $\overline{\mathcal{G}} = \overline{G}[\cdot]\overline{G}^\dagger$ denote the associated map. We let $\overline{\mathcal{G}}_{\mathrm{FT}}$ denote a fault-tolerant gate gadget for G. In the absence of noise, the gadget $\overline{\mathcal{G}}_{\mathrm{FT}}$ implements $\overline{\mathcal{G}}$ when acting on states in the codespace, and in the presence of noise, it implements a map denoted $\widetilde{\mathcal{G}}_{\mathrm{FT}}$, while obeying certain properties related to propagation of physical errors [90]. A scheme for universal FTQC requires the specification of a fault-tolerant gate gadget for a universal set of gates.

We now expand more on the formal properties that FTQC guarantees about these maps. First of all, each subset $S \subset [n']$ of physical qubits is either a correctable or an uncorrectable subset, depending on whether there exists an error channel $(\mathcal{I}_{S^c} \otimes \mathcal{W}_S)$ acting trivially on S^c and nontrivially only on qubits within S that can induce a logical error, in the sense that $\mathcal{Q} \neq \mathcal{Q} \circ (\mathcal{I}_{S^c} \otimes \mathcal{W}_S) \circ \mathcal{Q}$. The assumption that the FTQC scheme has a threshold p_0 against local stochastic noise indicates that whenever $p < p_0$,

$$\sum_{S \text{ uncorrectable}} (1-p)^{n'-|S|} p^{|S|} \leq \Gamma(\mathcal{E})\,, \tag{18}$$

where $\Gamma(\mathcal{E})$ a quantity that depends on p but has the property that, for fixed p, $\Gamma(\mathcal{E})$ can be exponentially driven to zero simply by choosing a larger QEC code (indicated by the encoding map $\mathcal{E}$) from the code family, incurring

at most a polylogarithmic overhead (which is defined as the factor by which the number of physical qubits and physical gates must increase).

Furthermore, abstracting away from physical errors, we can similarly quantify the logical error incurred by performing a fault-tolerant gate as

$$\frac{1}{2}\left\|\overline{\mathcal{G}} \circ \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} - \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{G}}_{\mathrm{FT}} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}}\right\|_{\diamond} \leq \Gamma(\mathcal{E})\,, \tag{19}$$

where again $\Gamma(\mathcal{E})$ is used to denote a quantity that vanishes with increasing code size—the overhead of achieving logical error $\Gamma(\mathcal{E}) \leq \delta$ is at worst polylog$(1/\delta)$. The equation above states that performing the noisy gate $\widetilde{\mathcal{G}}_{\mathrm{FT}}$ sandwiched between noisy rounds of QEC and then projecting onto the codespace is nearly equivalent to simply performing noisy QEC, projecting onto the codespace, and applying the perfect logical gate.

4 Distillation–teleportation protocol: details and error analysis

In this section, we examine each step of our protocol in more detail, propagating errors and proving formal statements about the performance of each step. Whereas in previous sections we labeled inputs and outputs of quantum circuits with kets, here we label them with density matrices, since we will be investigating the impact of noise, which may lead states to lose their purity. It should be understood that the unitary circuit elements are applying the associated unitary channel on the density matrices on which they act.

4.1 Noisy physical resource state preparation with QRAM device

As discussed, we assume we have access to a noisy physical QRAM device that attempts to perform the QRAM operation of Eq. (1) on the state $|+\rangle^{\otimes n}$ at the physical level, in order to produce the resource state $|\Psi(g)\rangle$ of Eq. (9), given as input the dataset g. We depict this relation pictorially via the following circuit, where the red wire indicates that the outgoing quantum register is unencoded and the qubits are physical qubits.

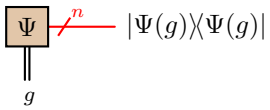

(20)

However, it is unrealistic to assume that the QRAM device can be implemented perfectly at the physical level. Rather, the state it produces is a state $\widetilde{\psi}(g)$ that has some nonzero fidelity with $|\Psi(g)\rangle\langle\Psi(g)|$. We assume that the noise in the device is independent of the dataset, as in Definition 1, so that the QRAM channel decomposes as $\mathcal{N}_2 \circ \mathcal{V}(g) \circ \mathcal{N}_1$. Then, we have

$$\widetilde{\psi}(g) = \mathcal{N}_2\left[V(g)\,\mathcal{N}_1\left[|+\rangle\langle+|^{\otimes n}\right]V(g)^\dagger\right] \tag{21}$$

In the circuit notation, we indicate the existence of noise by outlining the gate and its output in red.

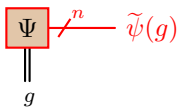

(22)

We characterize the noise strength by the infidelity $1 - F(g)_{\mathrm{phys}}$, where $F(g)_{\mathrm{phys}}$ is the fidelity of the state $\widetilde{\psi}(g)$ with respect to the ideal state $|\Psi(g)\rangle\langle\Psi(g)|$, given by

$$F(g)_{\mathrm{phys}} = \langle\Psi(g)|\widetilde{\psi}(g)|\Psi(g)\rangle\,. \tag{23}$$

4.2 Encoding physical resource states into logical resource states

The noisy resource state $\widetilde{\psi}(g)$ on n physical qubits is unprotected from errors, and once prepared it must immediately be encoded into a QEC code. Ideally, this encoding process would prepare the state $\mathcal{E}[\widetilde{\psi}(g)]$, which is in the codespace of the code, as in the figure below, where the black wire with a black slash-n indicates that the register contains n

logical qubits encoded into some number $n' > n$ of physical qubits.

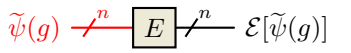

(24)

However, the encoding process may not be fault tolerant, that is, it may introduce additional errors into the logical output state that are proportional to the physical error rate p of the hardware and the number of gates in the encoding circuit that implements $\mathcal{E}$; this logical error cannot be suppressed simply by growing the code size. Hence, we need an encoding procedure $\mathcal{E}_{\mathrm{FT}}$ that is fault-tolerant against the physical noise model, in the sense specified in Proposition 2 below. When we attempt to realize $\mathcal{E}_{\mathrm{FT}}$ on faulty hardware, we instead implement a channel $\widetilde{\mathcal{E}}_{\mathrm{FT}}$. After applying $\widetilde{\mathcal{E}}_{\mathrm{FT}}$, we apply the (noisy) fault-tolerant QEC gadget—implementing the map $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$—on the now-encoded state to correct for physical errors that occurred during $\widetilde{\mathcal{E}}_{\mathrm{FT}}$. The output state is $\widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)]$ (we could alternatively think of $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$ as part of $\widetilde{\mathcal{E}}_{\mathrm{FT}}$ and omit it from the expression[7]). This state is still not necessarily in the codespace, but the action of the noisy-but-fault-tolerant $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$ brings it close to the codespace, in the robust sense required by formal proofs of fault tolerance, for example, in Ref. [90]. The closest codespace state is obtained by applying the QEC projector $\mathcal{Q}$, resulting in

$$\overline{\phi(g)} = \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)] \,. \tag{25}$$

Of course, we cannot apply noiseless $\mathcal{Q}$ on actual hardware, but the probability of logical deviation from $\overline{\phi(g)}$ can be suppressed to an arbitrarily small quantity simply by growing the code size (incurring only logarithmic overheads or even constant overhead if one uses constant-rate codes [89]) provided that the physical error rate is below the threshold. Since the rest of our protocol is performed using fault-tolerant logical gates, we identify $\overline{\phi(g)}$ as the logical output of the noisy encoding, denoted pictorially by outlining the encoding gate in red.

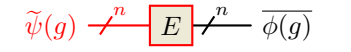

(26)

We define the encoding error as the maximum (over arbitrary n-qubit input ρ) trace distance between the state $\mathcal{E}[\rho]$ obtained from perfect encoding and the state $\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho]$ obtained from noisy encoding, as follows.

Definition 3 (Encoding error)**.** *Consider a QEC code encoding n logical qubits, specified by an encoding isometry $\mathcal{E}$. Let $\mathcal{Q}$ denote a QEC projector for the code $\mathcal{E}$, and let $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$ denote the noisy implementation of a fault-tolerant QEC gadget for $\mathcal{Q}$. Given a fault-tolerant encoding procedure $\mathcal{E}_{\mathrm{FT}}$ and its noisy implementation $\widetilde{\mathcal{E}}_{\mathrm{FT}}$, the encoding error is defined by the following expression*

$$\varepsilon_{\mathrm{enc}} = \sup_{\rho} \frac{1}{2} \| \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho] - \mathcal{E}[\rho] \|_1 \,, \tag{27}$$

where the supremum is taken over all n-qubit physical states ρ. Note that $\varepsilon_{\mathrm{enc}}$ is expected to have a dependence on n, as well as the physical error rate p of the hardware.

The goal of the encoding step is to take a physical state $\widetilde{\psi}(g)$ with some nonzero fidelity $F(g)_{\mathrm{phys}} = \langle \Psi(g) | \widetilde{\psi}(g) | \Psi(g) \rangle$, and map it to an encoded state $\overline{\phi(g)}$ with a smaller but still nonzero fidelity

$$F(g)_{\mathrm{enc}} = \langle \overline{\Psi(g)} | \overline{\phi(g)} | \overline{\Psi(g)} \rangle \,. \tag{28}$$

Using the definition of the encoding error, we can make an additive bound $F(g)_{\mathrm{enc}} \geq F(g)_{\mathrm{phys}} - \varepsilon_{\mathrm{enc}}$, which is sufficient when $F(g)_{\mathrm{phys}}$ is large enough that the right-hand side is greater than zero. However, if $F(g)_{\mathrm{phys}}$ is smaller than $\varepsilon_{\mathrm{enc}}$, then this bound is no longer meaningful. We can instead show a multiplicative bound roughly of the form $F(g)_{\mathrm{enc}} \geq (1 - O(\varepsilon_{\mathrm{enc}})) F(g)_{\mathrm{phys}}$, which is more powerful in the small-fidelity regime. To show this, we have to manually perform a Pauli twirl of the encoding operation, which allows us to guarantee that the twirled encoding channel is stochastic (i.e., it acts as identity with some probability $1 - O(\varepsilon_{\mathrm{enc}})$ and as some other CPTP channel otherwise); see, for example, Ref. [94, Lemma 5.2.4] and Ref. [95, Lemma 3]. Without Pauli twirling, small encoding error alone is not sufficient to rule out the possibility that the encoding channel has coherent error, such as small unitary rotations, which can degrade the fidelity in an additive rather than multiplicative way.

The Pauli twirl involves applying a randomly chosen physical Pauli operator, encoding, and then applying the same Pauli operator but on the logical level. The set of physical Pauli operators can cause the fidelity to degrade

[7] However, note that physical errors in $\widetilde{\mathcal{E}}_{\mathrm{FT}}$ can combine with physical errors in $\widetilde{\mathcal{Q}}_{\mathrm{FT}}$ that create a logical difference between $\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)]$ and $\mathcal{Q} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)]$.

by a factor $(1-p)^n$ (which is on the order of $1-O(\varepsilon_{\text{enc}})$ anyway). We capture this in the following proposition, which states that if we already have an encoding method $\mathcal{E}'_{\text{FT}}$ with small encoding error, we can construct a $\mathcal{E}_{\text{FT}}$ that degrades the fidelity in this multiplicative way. The formal proof is provided in Appendix B.

Proposition 1 (Pauli twirling the encoding channel)**.** *Denote the fidelity of the physical state by $F(g)_{\text{phys}} = \langle\Psi(g)|\widetilde{\psi}(g)|\Psi(g)\rangle$. Suppose the processor is subject to circuit-level stochastic noise with strength p (Definition 2), and et $\mathcal{E}'_{\text{FT}}$ be a fault-tolerant encoding channel with encoding error ε_{enc} (as in Definition 3). Then, there exists another fault-tolerant encoding channel $\mathcal{E}_{\text{FT}}$ (formed by Pauli-twirling $\mathcal{E}'_{\text{FT}}$), for which*

$$\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq (1-p)^n\Big((1-3\varepsilon_{\text{enc}})F(g)_{\text{phys}} - 2\Gamma(\mathcal{E})\Big)\,, \tag{29}$$

where $\overline{\phi(g)}$ is defined from $\mathcal{E}_{\text{FT}}$ as in Eq. (25), and $\Gamma(\mathcal{E})$ is a quantity that vanishes with increasing code size, provided the physical error rate p is below a constant threshold, as discussed in Section 3.3.

Next, we explain how to achieve an encoding with manageable ε_{enc}. It should be noted that for fixed p, it is unavoidable for ε_{enc} to grow at least linearly with the number of logical qubits n, simply because we start with an unencoded state $|\psi\rangle$ on n faulty physical qubits. Remarkably, it is possible to construct a fault-tolerant encoding procedure $\mathcal{E}_{\text{FT}}$ such that in the presence of noise, the logical encoding error ε_{enc} has no direct dependence on the block size and the distance of the code $\mathcal{E}$ that we are encoding into. Such fault-tolerant encoding procedures exist for specific families of quantum codes such as the surface code [49]. Here, to keep our main results as agnostic to the underlying quantum code $\mathcal{E}$ as possible, we opt to use a fault-tolerant encoding procedure that works for any quantum code [59]. This procedure is based on concatenated-code quantum fault tolerance [87], and its fault tolerance for quantum input–quantum output tasks was recently proven under the circuit-level stochastic noise model defined in Definition 2.

The circuit-level stochastic noise model in Definition 2 is a special case of a more general model called local-stochastic noise model [89], which additionally allows for some correlations between gate faults. For our purposes, Proposition 2 below will be using a result of Ref. [59] that is proven under this circuit-level stochastic noise model. However, as is often the case in fault tolerance analysis, we expect the same statement extends to the local-stochastic noise model. Since implementing this extension is beyond the scope of our work, we keep the discussion simple by working with Definition 2 here.

Proposition 2 (Error in fault-tolerant encoding channel for general codes)**.** *Consider a family of quantum error-correcting codes encoding n logical qubits into a codespace, and suppose that this family has a threshold p_0 with respect to local stochastic noise (implying Eq. (18)). Correspondingly, for a particular instance of the family (labeled by its encoding map $\mathcal{E}$), let $\mathcal{Q}$ be the ideal QEC projector, $\mathcal{Q}_{\text{FT}}$ be the fault-tolerant QEC gadget, and $\Gamma(\mathcal{E})$ be the logical error suppression function (see Section 3.3). Then, there exists a fault-tolerant encoding procedure $\mathcal{E}_{\text{FT}}$ of size $|\mathcal{E}|\cdot\text{poly}(k)$, such that, when implemented under the circuit-level stochastic noise model (Definition 2), the encoding error as defined in Definition 3 satisfies*

$$\varepsilon_{\text{enc}} = \sup_{\rho}\frac{1}{2}\|\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\text{FT}}\circ\widetilde{\mathcal{E}}_{\text{FT}}[\rho] - \mathcal{E}[\rho]\|_1 \;\leq \Gamma(\mathcal{E}) + 2\sqrt{cp}n + 2|\mathcal{E}|(cp)^k,$$

where C is an absolute constant and k can take values from a sequence of geometrically increasing integers, provided that the physical error rate p is below some constant threshold.

The formal proof is provided in Appendix B. We remark that the final term in the expression above is similar to the $\Gamma(\mathcal{E})$ term, in the sense that that for any δ, one can choose $k = \text{polylog}(n/\delta)$ and ensure that it is smaller than δ. Thus, the overhead is only polylogarithmic and we neglect its contribution in our analysis elsewhere in the protocol. These two propositions together allow us to show that the encoded state maintains substantial fidelity with the ideal resource state, for use in our protocol.

Corollary 1. *Suppose the quantum processor is subject to circuit-level stochastic noise with error rate p (defined in Definition 2). Let*

$$F_{\min} = (1-np-6n\sqrt{cp})\min_g\langle\Psi(g)|\widetilde{\psi}(g)|\Psi(g)\rangle\,. \tag{30}$$

Then, there exists a fault-tolerant encoding procedure $\mathcal{E}_{\text{FT}}$, such that for all g we have

$$\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq F_{\min}, \tag{31}$$

where $\overline{\phi(g)}$ is defined from $\mathcal{E}_{\text{FT}}$ as in Eq. (25). Moreover, the qubit and gate overhead of applying $\mathcal{E}_{\text{FT}}$ is $\text{poly}(n)$.

Proof. This follows by defining $\mathcal{E}'_{\mathrm{FT}}$ to be the encoding procedure shown to exist in Proposition 2, and then applying Proposition 1 to form $\mathcal{E}_{\mathrm{FT}}$. Note that $(1-p)^n > 1 - np + \delta$ for $\delta = O(n^2p^2)$. When the value of k is taken to be $\mathrm{polylog}(n/\delta)$, and the QEC code size is taken to be $n\,\mathrm{polylog}(n/\delta)$, then the error terms $\Gamma(\mathcal{E})$ and $2|\mathcal{E}|(cp)^k$ can be made $O(\delta)$, and the stated fidelity $F_{\min}$ can be guaranteed. □

4.3 Partial Clifford twirling

The preparation and encoding processes produce a state $\overline{\phi(g)}$, but the distillation process discussed later can only distill $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$ from many copies of $\overline{\phi(g)}$ if the principal eigenvector of $\overline{\phi(g)}$ is $|\overline{\Psi(g)}\rangle$, which we cannot guarantee in general from the noise model we have assumed.

One solution to this challenge is to use twirling [96], also known as randomized compiling. This technique can convert general noise into better-behaved noise by inserting random gates from a certain set into a quantum circuit, and compensating for their effect by modifying other gates in the circuit. We already saw an example of this process in the encoding step (Proposition 1), where the insertion of Pauli gates led the noise to become stochastic. It is well known that twirling by random Clifford gates can lead to stronger results, converting arbitrary (gate-independent) noise into depolarizing noise. For example, Ref. [97] showed how randomized compiling, when actively applied within the physical QRAM device, can help mitigate the impacts of coherent errors on fidelity. However, in our setting—where we ideally consider a fully passive QRAM device and thus can only apply twirling "outside" the device—we cannot use the full Clifford group since, for example, if we conjugate the QRAM unitary $V(g)$ by the H gate, we do not obtain another QRAM unitary. Our twirling set will instead be the subset of the n-qubit Clifford gates generated by Z, X, CZ and CX (CNOT) gates, which do map QRAM unitaries to QRAM unitaries.

4.3.1 Setup and important lemmas

In what follows, we make use of the properties of $\{0,1\}^n$ as an n-dimensional vector space over the field $\mathbb{F}_2$.

Definition 4 (Partial Clifford twirling set)**.** *Suppose we are given A, B, u, and v, where*

- *A is an $n \times n$ invertible matrix over $\mathbb{F}_2$,*
- *B is an upper triangular $n \times n$ matrix (with zeros on the diagonal) with entries in $\mathbb{F}_2$,*
- *$u \in \mathbb{F}_2^n$,*
- *$v \in \mathbb{F}_2^n$.*

Define M_A to be the quantum gate that enacts $M_A : |x\rangle \mapsto |Ax\rangle$ for all $x \in \mathbb{F}_2^n$, which can be composed from $O(n^2)$ CX gates via Gaussian elimination, and define the diagonal unitary

$$Q_B = \prod_{1 \le i < j \le n} \mathrm{CZ}_{ij}^{B_{ij}}\,, \tag{32}$$

where CZ_{ij}^k is the identity gate when $k = 0$ and the CZ gate between qubits i and j when $k = 1$. Then, define the n-qubit quantum gate that corresponds to (A, B, u, v) as

$$C = Z^v Q_B M_A^\dagger X^u, \tag{33}$$

That is, C is a product of $O(n)$ single-qubit Pauli-Z gates determined by the entries of v, $O(n^2)$ CZ gates determined by the entries of B, $O(n)$ Pauli-X gates determined by the entries of X, and $O(n^2)$ CX gates determined by the entries of A (via M_A).

Let the twirling set $\mathbb{T}$ consist of all gates C constructed in this fashion from some choice of A, B, u, v. When we say to generate a random gate from $\mathbb{T}$, we mean to generate a uniformly random A, B, u, and v and choose the corresponding $C \in \mathbb{T}$.

We call this partial Clifford twirling because the set $\mathbb{T}$ is a subset of the n-qubit Clifford group. In particular, the set $\mathbb{T}$ is generated by $X, Z, \mathrm{CX}, \mathrm{CZ}$, and the Clifford group is obtained by adding the Hadamard H and phase gate S to the generating set.

Proposition 3. *Let $g : \mathbb{F}_2^n \to \mathbb{F}_2$ be a data table (Boolean function) and let $C \in \mathbb{T}$ be a twirling gate corresponding to choice (A, B, u, v). Let M_A and Q_B be defined as in Definition 4. Then, we have*

$$|\Psi(g)\rangle = V(g)|+\rangle^{\otimes n} = C\,V(g_C)|+\rangle^{\otimes n} = C|\Psi(g_C)\rangle \tag{34}$$

where for any $x \in \mathbb{F}_2^n$

$$g_C(x) = g(Ax \oplus u) \oplus [x \cdot v] \oplus \left[x^\top B x\right] \tag{35}$$

Proof. We note that $X^u M_A |+\rangle^{\otimes n} = |+\rangle^{\otimes n}$ since X^u and M_A are made from X and CX gates. The gates X^u and M_A simply permute the computational basis states, viewed as vectors in $\mathbb{F}_2^n$, by an affine transformation, and we can further note that $M_A^\dagger X^u V(g) X^u M_A = V(g')$ where $g'(x) = g(Ax \oplus u)$. Finally, the operation $Z^v Q_B$ is diagonal, and equal to $V(h)$ where $h(x) = [x \cdot v] \oplus \left[x^\top B x\right]$. We have the composition rule $V(h)V(g') = V(g' \oplus h) = V(g_C)$. The statement follows from these facts as

$$\begin{aligned} C\,V(g)|+\rangle^{\otimes n} &= Z^v Q_B M_A^\dagger X^u V(g)|+\rangle^{\otimes n} = Z^v Q_B M_A^\dagger X^u V(g) X^u M_A |+\rangle^{\otimes n} = Z^v Q_B V(g')|+\rangle^{\otimes n} \\ &= V(h)V(g')|+\rangle^{\otimes n} = V(g_C)|+\rangle^{\otimes n} \end{aligned} \tag{36}$$

□

Next, we will examine how random choice of $C \in \mathbb{T}$ spreads Pauli operators. We will examine the set of signed Pauli operators, and specific subsets of it.

Definition 5 (Signed Pauli set and noteworthy subsets)**.** *Define* $\mathbb{P}$ *to be the set of* 2^{2n+1} *signed Pauli strings written in the canonical form* $\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a$*, where* $s \in \mathbb{F}_2$*,* $a, b \in \mathbb{F}_2^n$*. The factor of* i *ensures that each operator in* $\mathbb{P}$ *is Hermitian. We partition* $\mathbb{P}$ *into several nonoverlapping subsets*

$$\mathbb{P}_0 = \{\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a \in \mathbb{P} : a = b = 0^n, s = 0\} = \{\mathbb{I}\} \tag{37}$$

$$\mathbb{P}_1 = \{\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a \in \mathbb{P} : a = b = 0^n, s = 1\} = \{-\mathbb{I}\} \tag{38}$$

$$\mathbb{P}_Z = \{\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a \in \mathbb{P} : a \neq 0^n, b = 0^n\} \tag{39}$$

$$\mathbb{P}_{\text{even}} = \{\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a \in \mathbb{P} : b \neq 0^n, a \cdot b = 0\} \tag{40}$$

$$\mathbb{P}_{\text{odd}} = \{\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a \in \mathbb{P} : a \cdot b = 1\} \tag{41}$$

We argue that the Pauli operators are spread uniformly over the subset of $\mathbb{P}$ to which it belongs. The full proof of the following proposition is provided in Appendix C.1.

Proposition 4 (Twirling spreads Paulis uniformly)**.** *Let* $C \sim \mathbb{T}$ *denote choosing* C *randomly from* $\mathbb{T}$ *as described in Definition 4. Given a fixed* $P \in \mathbb{P}$*, for any* $C \in \mathbb{T}$*,* $CPC^\dagger \in \mathbb{P}$ *since* C *is Clifford. Furthermore, let* $Q \in \mathbb{P}$ *be a random variable formed by choosing* $C \sim \mathbb{T}$ *and defining* $Q = CPC^\dagger$*. Then, the distribution over* Q *is the uniform distribution over the subset of* $\mathbb{P}$ *(i.e.,* $\mathbb{P}_0$*,* $\mathbb{P}_1$*,* $\mathbb{P}_Z$*,* $\mathbb{P}_{\text{even}}$*, or* $\mathbb{P}_{\text{odd}}$*) to which* P *belongs.*

Proof idea. Consider a Pauli $P \in \mathbb{P}$ written in canonical form as $P = \mathrm{i}^{a\cdot b}(-1)^s X^b Z^a$, and a Clifford $C = Z^v Q_B M_A^\dagger X^u \in \mathbb{T}$. We explicitly compute the Pauli $CPC^\dagger = \mathrm{i}^{a'\cdot b'}(-1)^{s'} X^{b'} Z^{a'}$ and give formulas for s', a', b' in terms of s, a, b, v, B, A, u. Then, we use these formulas to verify that for $t \in \{0, 1, Z, \text{odd}, \text{even}\}$ if $P \in \mathbb{P}_t$ and v, B, A, u are chosen uniformly at random, then s', a', b' are uniformly random over all values consistent with the definition of $\mathbb{P}_t$. □

4.3.2 Modification to circuit

We now explain precisely how our protocol changes when we implement partial Clifford twirling. We want to implement $V(g)$. Each time we query the physical QRAM device, we generate a C uniformly at random from the Clifford subset $\mathbb{T}$, as defined in Definition 4. We update the function g to be g_C using Eq. (35). In particular, the value at each address x may need to be updated, but each one can be computed with a simple poly(n)-time classical computation and a single query to learn $g(y)$ for a particular y. We use the QRAM device to produce the noisy physical resource state $\widetilde{\psi}(g_C)$, and then we encode that resource state, yielding $\overline{\phi(g_C)}$, as in Eq. (25). Only then do we fault tolerantly apply the gate $\overline{C}$, which consists of $O(n^2)$ logical Clifford gates. This full procedure is depicted in the following circuit.

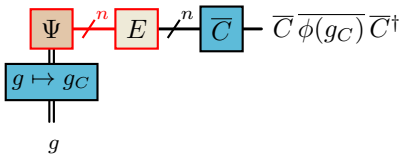

(42)

Each time the QRAM device is called, an independent random C is chosen. Thus, the output state may be modeled as the mixture

$$\overline{\phi(g)}_{\text{twirl}} = \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \overline{C}\,\overline{\phi(g_C)}\,\overline{C}^\dagger\,. \tag{43}$$

We have not included the twirling step in the main circuit of Figure 2; partly because it would clutter the figure, and partly because we feel that the twirling step may not be necessary in practice if the QRAM device can be constructed in such a way that $\overline{\phi(g)}$ already has $|\overline{\Psi(g)}\rangle$ as its principal eigenvector.

4.3.3 Twirling ensures the principal eigenvector is correct

Now, we are ready to present the main finding of this section. The full proof is provided in Appendix C.2.

Proposition 5 (Correct top eigenvector)**.** *Suppose that for every g, the state $\overline{\phi(g)}$ defined in Eq. (25) satisfies $\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq F_{\min}$, and suppose that the faulty QRAM device is subject to dataset-independent noise (Definition 1). Let $C\sim\mathbb{T}$ denote drawing C randomly from the twirling set as described in Definition 4, and let $\mathbb{E}$ denote expectation value. For each g, let $\overline{\phi(g)}_{\text{twirl}}$ be defined as in Eq. (43). Then $\overline{\phi(g)}_{\text{twirl}}$ satisfies the eigenvalue equation*

$$\overline{\phi(g)}_{\text{twirl}}|\overline{\Psi(g)}\rangle = \lambda_{\text{twirl}}|\overline{\Psi(g)}\rangle\,, \tag{44}$$

with $\lambda_{\text{twirl}} \geq F_{\min}$. Furthermore, all other eigenvalues of $\overline{\phi(g)}_{\text{twirl}}$ are no larger than 2^{-n+1}.

Proof idea. Due to the dataset-independent assumption, the noise in the physical QRAM device and the encoding step can be consolidated into a noise matrix $\chi_{P,P'}$ (where $P, P' \in \mathbb{P}$) for which it is always possible to write

$$\overline{\phi(g_C)} = \sum_{P,P'\in\mathbb{P}} \chi_{P,P'}\,\overline{P}|\overline{\Psi(g_C)}\rangle\langle\overline{\Psi(g_C)}|\overline{P}'\,. \tag{45}$$

Noting $|\overline{\Psi(g_C)}\rangle = \overline{C}^\dagger|\overline{\Psi(g)}\rangle$ (Proposition 3), we can then say

$$\overline{\phi(g)}_{\text{twirl}} = \sum_{P,P'\in\mathbb{P}} \chi_{P,P'} \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \overline{C}\,\overline{P}\,\overline{C}^\dagger|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{C}\,\overline{P}'\,\overline{C}^\dagger\,. \tag{46}$$

We now use the fact that randomly choosing C leads $CPC^\dagger$ to uniformly cover a large subset of of $\mathbb{P}$ (Proposition 4). The offdiagonal terms with $P \neq \pm P'$ vanish due to the uniformly random sign. This conclusion did not require the full size of $\mathbb{T}$; it is a consequence of the fact that $\mathbb{T}$ contains the Pauli group, and Pauli-twirling leads the effective channel to become a Pauli channel. If $\mathbb{T}$ were the entire Clifford group, then $CPC^\dagger$ would be a uniformly random Pauli, and we would immediately be able to use the 1-design property of the Pauli set to say that, for any g and for any case where $P \neq \pm\mathbb{I}$, the quantity $\mathbb{E}_{C\sim\mathbb{T}}\,\overline{C}\,\overline{P}\,\overline{C}^\dagger|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{C}\,\overline{P}\,\overline{C}^\dagger$ is equal to the maximally mixed state. This would then immediately imply the statement we seek, and also that all other eigenvalues λ_{other} satsify $\lambda_{\text{other}} \leq 2^{-n}$. Generally, this would be a manifestation of the fact that Clifford twirling transforms any noise channel into a depolarizing channel. However, as $\mathbb{T}$ is not the full Clifford group, we have to do more work; some subsets of $\mathbb{P}$ may be underweighted or overweighted. Nevertheless, we show that there is enough uniformity to recover a similar result, albeit with the bound on λ_{other} suffering a factor-of-2 overhead. □

4.4 Distillation of logical resource states

The encoding step, combined with twirling, produces the logical encoded states $\overline{\phi(g)}_{\text{twirl}}$, which are guaranteed to have $|\overline{\Psi(g)}\rangle$ as their principal eigenvector (Proposition 5). The remaining steps of the protocol act directly on the encoded states (using fault-tolerant gadgets for a universal set of gates): all physical errors created during the distillation and teleportation portions of the protocol can be prevented from turning into logical errors by growing the code distance. Thus, we describe the distillation and teleportation protocol by their logical quantum circuits, and in the description of our protocol, we keep the overline notation to remind the reader that these operations are meant to be performed fault tolerantly on the encoded Hilbert space.

However, the distillation ideas presented here apply generally, regardless of whether/how the states and operations are encoded with QEC. Later in the section, including some of the proposition statements, we drop the overlines, since the statements could be of independent interest.

The procedures we consider for distilling the logical QRAM resource state are state agnostic. Namely, given many copies of an arbitrary state $\overline{\rho}_{\text{in}}$, the distillation protocol prepares (up to small trace distance error) the pure state $|\overline{\Xi}\rangle\langle\overline{\Xi}|$, where $|\overline{\Xi}\rangle$ is the principal eigenvector of $\overline{\rho}_{\text{in}}$. In other words, our distillation procedure is equivalent to the task of *quantum purity amplification* [68].

Pictorially, the distillation step accomplishes the following operation within our protocol. For simplicity, the circuit below depicts all copies being prepared at the beginning and processed at once; in practice, it is possible to prepare the states in a streaming fashion with less space requirement, as we discuss later.

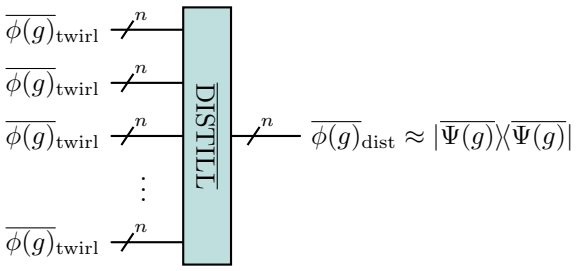

(47)

Now we present the main result of this section, as applied to our protocol, utilizing the general techniques discussed later in the section.

Proposition 6. *Suppose that the QRAM device is subject to dataset-independent noise, as in Definition 1, and that for every g, the states $\overline{\phi(g)}$ produced on input g (see Eq. (25)) satisfy $\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq F_{\min}$, with $F_{\min} \geq 2^{-n+2}$. Then, for any error parameter $\varepsilon_{\text{dist}}$, by applying a carefully crafted sequence of subsequent (fractional) swap operations on $O(\frac{1-F_{\min}}{F_{\min}^2}(\frac{1}{\varepsilon_{\text{dist}}} + \frac{1}{F_{\min}}))$ copies of $\overline{\phi(g)}_{\text{twirl}}$ (from Eq. (43)) we can distill a state $\overline{\phi(g)}_{\text{dist}}$ that satisfies*

$$\frac{1}{2}\||\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\text{dist}}\|_1 \leq \varepsilon_{\text{dist}} \,. \tag{48}$$

The protocol requires $O(1)$ single-qubit gates and $O(n)$ controlled swap operations, per copy consumed.

Proof. This is based on Proposition 5, which shows that $\overline{\phi(g)}_{\text{twirl}}$ has the correct top eigenvector, combined with some state-agnostic quantum purity amplification protocol to distill the top eigenvector. When $F_{\min}$ is close to 1 the iterated swap test achieves this with a close-to-optimal $\approx (1 - F_{\min})/\varepsilon_{\text{dist}}$ number of copies of $\overline{\phi(g)}_{\text{twirl}}$ and the same order of swap tests, that is, circuit (50)—see Proposition 7 and Lemma 1. However, for smaller values of $F_{\min}$, the iterated swap test incurs an $\exp(\Theta(1/F_{\min}))$ overhead, see the discussion at the end of Section 4.4.1.

In the general case, we can exploit the fact that the second largest eigenvalue of $\overline{\phi(g)}_{\text{twirl}}$ is upper bounded by $2^{-n+1} \leq F_{\min}/2$ due to Proposition 5, and use a protocol based on quantum principal component analysis (Proposition 10), achieving the stated copy complexity utilizing a matching number of swap-test-like gadgets from Fig. 4.

Both Proposition 10 and Proposition 7 are described as producing a state for which the largest eigenvalue is close to 1. We can turn this into a trace distance bound via Lemma 1. The stated gate complexity follows from the observation that Fig. 4 has 3 controlled swap operations (targetting $(n+1)$-qubit registers) and 4 single-qubit gates. □

To convert to trace distance as in Eq. (48), our analysis uses the fact that a mixed state is as close to its principal eigenvector as its principal eigenvalue is to 1, captured in the following lemma.

Lemma 1. *A quantum state $\overline{\rho}_{\text{out}}$, whose principal eigenvector is $|\overline{\Xi}\rangle$ with eigenvalue $1-\eta$, satisfies*

$$\frac{1}{2}\|\overline{\rho}_{\text{out}} - |\overline{\Xi}\rangle\langle\overline{\Xi}|\|_1 = \eta. \tag{49}$$

Proof. Let $\lambda_1, \ldots, \lambda_d$ be the eigenvalues of the state $\overline{\rho}_{\text{out}}$, with $\lambda_1 = 1 - \eta$ the principal eigenvector. They satisfy $\sum_{j=1}^{d} \lambda_j = 1$, so $\sum_{j=2}^{d} \lambda_j = \eta$. The operator $\overline{\rho}_{\text{out}} - |\overline{\Xi}\rangle\langle\overline{\Xi}|$ has the same eigenvectors as $\overline{\rho}_{\text{out}}$, and the eigenvalues are $-\eta, \lambda_2, \ldots, \lambda_d$. The sum of the singular values is thus equal to 2η, which completes the proof. □

4.4.1 Distillation with the iterated swap test

In this subsection, we drop the overlines in our notation, and speak generally about the task of quantum purity amplification. We first consider the iterated swap test, a quantum purity amplification procedure studied in detail in

Ref. [67] (a similar procedure was discussed in Ref. [70]), although there the analysis assumed that the input states ρ_{in} are a mixture of a rank-1 pure state and the maximally mixed state, that is, of the form $\rho_{\text{dep}} = (1-\delta)|\Xi\rangle\langle\Xi| + \frac{\delta}{d}\mathbb{I}$, with d the Hilbert space dimension and $\mathbb{I}/d$ the maximally mixed state. We do not make this assumption here. See also Ref. [69] for an analysis of the iterated swap test without this assumption.

The basic ingredient of this distillation procedure is the swap test (which in our application would be performed fault tolerantly using fault-tolerant gadgets for controlled swap, Hadamard, and measurement).

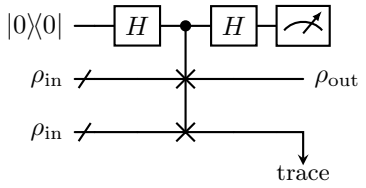

(50)

We say that the swap test passes if the first qubit is measured in $|0\rangle$ at the end of the circuit. Acting on two copies of the trace-1 state ρ_{in}, the probability of the swap test passing is given by

$$\Pr[\text{swap test passes}] = \frac{1 + \text{Tr}\left(\rho_{\text{in}}^2\right)}{2}, \tag{51}$$

and the state one obtains conditioned on the swap test passing is

$$\rho_{\text{out}} = \frac{\rho_{\text{in}} + \rho_{\text{in}}^2}{1 + \text{Tr}(\rho_{\text{in}}^2)}. \tag{52}$$

If the principal eigenvalue of ρ_{in} is $1 - \eta_{\text{in}}$, then we can bound

$$\Pr[\text{swap test passes}] \geq \frac{1 + (1-\eta_{\text{in}})^2}{2} \geq 1 - \eta_{\text{in}}. \tag{53}$$

It is easy to verify that ρ_{out} and ρ_{in} commute, and they also have the same eigensubspaces since $[0,1] \ni x \to x + x^2$ is strictly monotone. Moreover, denoting the principal eigenvalue of ρ_{out} by $1 - \eta_{\text{out}}$, we have

$$\begin{aligned}
\eta_{\text{out}} &= 1 - \frac{(1-\eta_{\text{in}}) + (1-\eta_{\text{in}})^2}{1 + \text{Tr}(\rho_{\text{in}}^2)} \\
&\leq 1 - \frac{(1-\eta_{\text{in}}) + (1-\eta_{\text{in}})^2}{1 + (1-\eta_{\text{in}})^2 + \eta_{\text{in}}^2} \\
&= \frac{\eta_{\text{in}}}{2}\left(\frac{1+\eta_{\text{in}}}{1-\eta_{\text{in}}+\eta_{\text{in}}^2}\right),
\end{aligned} \tag{54}$$

which is about $\eta_{\text{in}}/2$ for $\eta_{\text{in}} \ll 1$.

The idea of the distillation protocol studied in Refs. [67, 69, 70] is to iterate the swap test by feeding two copies of ρ_{out} that both passed the swap test back into the swap test as ρ_{in}, and repeating. Each time we successfully pass the swap test, the output state has higher purity and also the probability that the swap test passes at the next level is closer to 1. By continuing this process for sufficiently many iterations and consuming sufficiently many copies of ρ_{in}, we can produce a state with purity arbitrarily close to 1.

For small η_{in} (i.e., the regime of high input fidelity), the number of copies needed scales as $O(\eta_{\text{in}}/\varepsilon_{\text{dist}})$. Circuit (47) depicts creating all copies of the input state at the beginning of the protocol. For the swap test approach, if swap tests on disjoint pairs can be performed in parallel, the overall depth of the protocol could be $O(\log(\eta_{\text{in}}/\varepsilon_{\text{dist}}))$. However, as described in Ref. [67], the swap test approach can also be implemented in a streaming fashion, allowing one to reduce the space requirements to $O(\log(\eta_{\text{in}}/\varepsilon_{\text{dist}}))$ at the expense of requiring $O(\eta_{\text{in}}/\varepsilon_{\text{dist}})$ depth. We now give the formal statement of performance and costs for this type of distillation when $\eta_{\text{in}} < 1/4$, for general input states.

Proposition 7. *Consider a qudit in a mixed state ρ_{in} with its principal eigenvector $|\Xi\rangle$ having eigenvalue $1 - \eta_{\text{in}}$. After k successful iterations of the swap test procedure we obtain a state ρ_k such that $[\rho_{\text{in}}, \rho_k] = 0$ and if $\eta_{\text{in}} < 1/4$, then $\langle\Xi|\rho_k|\Xi\rangle \geq 1 - \frac{2^{-k}\eta_{\text{in}}}{1-4\eta_{\text{in}}}$. The expected number of copies of ρ_{in} consumed is upper bounded by $2^k/\sqrt{1-4\eta_{\text{in}}}$, and*

the expected number of required individual swap tests is upper bounded by $(2^k-1)/\sqrt{1-4\eta_{\text{in}}}$*, moreover the protocol needs to store at most* 1 *qubit and* $k+1$ *qudits for preparing* ρ_k*.*

Proof. Let $\rho_0 = \rho_{\text{in}}$ and $\eta_0 = \eta_{\text{in}}$. The fact that $[\rho_0, \rho_k] = 0$ follows directly from Eq. (52). For the stated space-efficient implementation, observe that until ρ_k is prepared it suffices to store at most one copy of each of $\rho_0, \rho_1, \ldots, \rho_{k-1}$, except for a single ρ_i where a swap test is performed, see Ref. [67, Algorithm 3].

Our proof is inspired by the calculations of Ref. [67]. Let us define $\eta_i := 1 - \langle \Xi | \rho_i | \Xi \rangle$. We can verify by induction that $\eta_i \leq \frac{2^{-i}\eta_0}{1-4\eta_0+2^{2-i}\eta_0} \left(\leq \frac{2^{-i}\eta_0}{1-4\eta_0} \right)$. This trivially holds for $i = 0$ and the induction step can be verified as follows:

$$\begin{aligned}
\eta_{i+1} &\leq \frac{\eta_i}{2}\left(\frac{1+\eta_i}{1-\eta_i+\eta_i^2}\right) && \text{(by (54))} \\
&\leq \frac{\eta_i}{2-4\eta_i} && \text{(since } (1+\eta_i)(1-2\eta_i) \leq 1-\eta_i+\eta_i^2\text{)} \\
&\leq \frac{2^{-i}\eta_0}{2(1-4\eta_0+2^{2-i}\eta_0)-2^{2-i}\eta_0} && \text{(by monotonicity of } \tfrac{x}{2-4x} \text{ and the induction hypothesis)} \\
&= \frac{2^{-1-i}\eta_0}{1-4\eta_0+2^{1-i}\eta_0}. && (55)
\end{aligned}$$

Let us denote by $p_i^{\text{succ}} = \frac{1+\text{Tr}\left(\rho_{i-1}^2\right)}{2}$ the success probability Eq. (51) of the swap test on ρ_{i-1}. The expected number c_i of copies of ρ_0 needed for preparing ρ_i is $c_i = 2^i \prod_{j=1}^{i} \frac{1}{p_j^{\text{succ}}}$, which is easy to verify by induction. The $i = 0$ case is trivial, and the induction step follows from the observation that to obtain ρ_{i+1} we need to repeat the swap test an expected number of $\frac{1}{p_{i+1}^{\text{succ}}}$ many times on two copies of ρ_i, i.e., $c_{i+1} = 2\frac{c_i}{p_{i+1}^{\text{succ}}}$.

We can bound c_i by deriving the bound $\prod_{j=1}^{i} p_j^{\text{succ}} \geq \sqrt{1-4\eta_0}$ as follows

$$\begin{aligned}
\left(\prod_{j=1}^{i} p_j^{\text{succ}}\right)^2 &\geq \prod_{j=1}^{i}(1-\eta_{j-1})^2 && \text{(by Eq. (53))} \\
&\geq \prod_{j=1}^{i}(1-2\eta_{j-1}) \geq \prod_{j=1}^{i}\left(1-\frac{2^{2-j}\eta_0}{1-4\eta_0+2^{3-j}\eta_0}\right) && \text{(by Eq. (55))} \\
&= 1-4\eta_0+2^{2-i}\eta_0. \qquad \left(\text{by induction: } (1-\tfrac{2^{1-i}\eta_0}{1-4\eta_0+2^{2-i}\eta_0})(1-4\eta_0+2^{2-i}\eta_0) = (1-4\eta_0+2^{1-i}\eta_0)\right)
\end{aligned}$$

The expected number s_i of swap tests used for preparing ρ_i can also be bounded by $(2^i-1)\prod_{j=1}^{i}\frac{1}{p_j^{\text{succ}}}$ via induction:

$$s_{i+1} = \frac{1+2s_i}{p_{i+1}^{\text{succ}}} \leq \frac{1+2\cdot(2^i-1)\prod_{j=1}^{i}\frac{1}{p_j^{\text{succ}}}}{p_{i+1}^{\text{succ}}} = (2^{i+1}-1)\prod_{j=1}^{i+1}\frac{1}{p_j^{\text{succ}}} + \frac{1-\prod_{j=1}^{i}\frac{1}{p_j^{\text{succ}}}}{p_{i+1}^{\text{succ}}} \leq (2^{i+1}-1)/\sqrt{1-4\eta_0}. \qquad \square$$

Proposition 7 only applies when $\eta_{\text{in}} < 1/4$, but the iterated swap test can still be successful even when the input fidelity is lower. We now consider what happens when $\eta_{\text{in}} \gg 0$. Let $\gamma_{\text{in}} := 1 - \eta_{\text{in}} = \langle \Xi | \rho_{\text{in}} | \Xi \rangle = \lambda_1(\rho_{\text{in}})$ denote the principal eigenvalue of ρ_{in}, where the notation $\lambda_i(\sigma)$ denotes the i-th largest eigenvalue of σ. Let us

assume that $\frac{\lambda_2(\rho_{\text{in}})}{\lambda_1(\rho_{\text{in}})} \leq \alpha$ for some value $\alpha < 1$, then we get the following guarantee on $\gamma_{\text{out}} := \langle \Xi | \rho_{\text{out}} | \Xi \rangle$

$$\begin{aligned}
\gamma_{\text{out}} &= \frac{\gamma_{\text{in}} + \gamma_{\text{in}}^2}{1 + \text{Tr}(\rho_{\text{in}}^2)} \\
&\geq \frac{\gamma_{\text{in}} + \gamma_{\text{in}}^2}{1 + \gamma_{\text{in}}^2 + \alpha\gamma_{\text{in}}(1 - \gamma_{\text{in}})} \\
&= \frac{\gamma_{\text{in}}(1 + \gamma_{\text{in}})}{1 + \gamma_{\text{in}} + (\alpha - 1)\gamma_{\text{in}}(1 - \gamma_{\text{in}})} \\
&= \frac{\gamma_{\text{in}}}{1 - (1 - \alpha)\gamma_{\text{in}}\frac{1-\gamma_{\text{in}}}{1+\gamma_{\text{in}}}} =: p(\gamma_{\text{in}})
\end{aligned} \tag{56}$$

which is always greater than γ_{in} when $\gamma_{\text{in}} \in (0, 1)$. Moreover $\frac{\lambda_2(\rho_{\text{out}})}{\lambda_1(\rho_{\text{out}})} = \frac{\lambda_2(\rho_{\text{in}})}{\lambda_1(\rho_{\text{in}})} \cdot \frac{1+\lambda_2(\rho_{\text{in}})}{1+\lambda_1(\rho_{\text{in}})} \leq \alpha$. Finally, by computing the derivative of Eq. (56) in γ_{in}, we get

$$\frac{1 + \gamma_{\text{in}}(2 - \gamma_{\text{in}} + 2\alpha\gamma_{\text{in}})}{(1 + \gamma_{\text{in}}(\alpha + (1 - \alpha)\gamma_{\text{in}}))^2} > 0,$$

which means that $p(\gamma_{\text{in}})$ in Eq. (56) is monotonically increasing in γ_{in}; therefore, if we replace γ_{in} with a lower bound on γ_{in} we still get a valid lower bound on γ_{out}. Let's assume that we have an "easy" scenario, where $\alpha \leq 10^{-3}$; a direct calculation shows that if $\gamma_{\text{in}} \geq 0.2$, then $p(\gamma_{\text{in}}) > 0.23$, $p^{\circ 2}(\gamma_{\text{in}}) = p(p(\gamma_{\text{in}})) > 0.26, \ldots, p^{\circ 9}(\gamma_{\text{in}}) > \frac{5}{6}$. By using Eq. (53), similarly to the proof of Proposition 7 we can see that the expected number of copies for successfully completing the 9 iterations of the swap test is at most

$$2^9 \prod_{j=0}^{8} \frac{2}{1 + (p^{\circ j}(\gamma_{\text{in}}))^2} = 2^{18} \prod_{j=0}^{8} \frac{1}{1 + (p^{\circ j}(\gamma_{\text{in}}))^2} < \frac{2^{18}}{22.6} < 11600. \tag{57}$$

Therefore, one can see that in the $\gamma_{\text{in}} = 1 - \eta_{\text{in}} \leq \frac{3}{4}$ regime the iterated swap test still works, but its efficiency degrades rapidly [69, Theorems 15 & 30]. In fact, even if we assume that all non-principal eigenvalues are the same, the protocol incurs an exponential cost in $1/\gamma_{\text{in}}$ because the initial swap tests make only a small increase in γ_{out} while they succeed with probability about $\frac{1}{2}$; see, for example, Ref. [67, Theorem 9]. Nevertheless, when $\gamma_{\text{in}} = 1 - \eta_{\text{in}}$ is lower bounded by a constant we get the desired asymptotically optimal complexity (see Section 4.4.2), by first magnifying γ_{in} to at least $\frac{4}{5}$ as in Eq. (56)–(57), and then applying Proposition 7, stated formally as follows.

Proposition 8. *In the setting of Proposition 7, suppose that ρ_{in} has principal eigenvalue $1 - \eta_{\text{in}}$, where $1 - \eta_{\text{in}}$ is greater than $\Omega(1)$. Furthermore, suppose that all other eigenvalues of ρ_{in} are bounded above by $\alpha(1 - \eta_{\text{in}})$ for some constant $\alpha < 1$. Then the expected number of copies consumed and the expected number of swap tests is the same as stated in Proposition 7, up to a multiplicative $O(1)$ constant.*

Proof. This follows from a generalization of the example above to arbitrary $\alpha < 1$. □

4.4.2 An asymptotically optimal distillation protocol with simultaneous use of all copies

Ref. [68] describes a state-agnostic quantum purity amplification protocol based on the Schur transform, which processes all copies in parallel. The authors prove [68, Theorem II.3] that for a generic quantum state ρ_{in} with principal eigenvector $|\Xi\rangle$, their protocol's sample complexity for outputting a quantum state ρ_{out} such that $\langle \Xi | \rho_{\text{out}} | \Xi \rangle \geq 1 - \varepsilon_{\text{dist}}$ is asymptotically optimal in the $\varepsilon_{\text{dist}} \to 0$ limit[8] and their protocol has sample complexity

$$\underset{\varepsilon_{\text{dist}} \to 0}{\sim} \frac{1}{\varepsilon_{\text{dist}}} \sum_{i=2}^{d} \frac{\lambda_i(\rho_{\text{in}})}{(\lambda_1(\rho_{\text{in}}) - \lambda_i(\rho_{\text{in}}))^2} + O(1). \tag{58}$$

When $\lambda_1(\rho_{\text{in}}) = 1 - \eta_{\text{in}}$, the above expression is maximized by $\lambda_2(\rho_{\text{in}}) = \eta_{\text{in}}$, resulting in complexity

$$\underset{\varepsilon_{\text{dist}} \to 0}{\sim} \frac{1}{\varepsilon_{\text{dist}}} \frac{\eta_{\text{in}}}{(1 - 2\eta_{\text{in}})^2} + O(1),$$

[8] This means that for any fixed spectrum $S = (\lambda_1, \lambda_2, \ldots, \lambda_d)$ the optimal sample complexity is $\frac{1}{\varepsilon_{\text{dist}}} \sum_{i=2}^{d} \frac{\lambda_i}{(\lambda_1 - \lambda_i)^2} + O(1)$ given that ρ_{in} has spectrum S. However, this does not say much about what happens for, say, constant $\varepsilon_{\text{dist}} \approx 1$, see Ref. [68, Appendix D].

which is rather close to the complexity achieved by Proposition 7.

If we only assume that $\lambda_2(\rho_{\text{in}}) \le \alpha\lambda_1(\rho_{\text{in}})$, then the expression in Eq. (58) is maximized when all nonzero, non-principal eigenvalues equal $\alpha\lambda_1(\rho_{\text{in}})$, giving rise to the complexity expression

$$\underset{\varepsilon_{\text{dist}}\to 0}{\sim} \quad \frac{1}{\varepsilon_{\text{dist}}} \frac{1-\gamma_{\text{in}}}{(1-\alpha)^2\gamma_{\text{in}}^2} + O(1). \tag{59}$$

As noted in Ref. [68], this is exponentially better in the $\gamma_{\text{in}} \to 0$ regime (i.e., low input fidelity) than the iterated swap test protocol described in Section 4.4.1. However, a major drawback of the corresponding protocol of Ref. [68] is that it requires storing and processing all copies in parallel, resulting in a large space complexity.

The authors of Ref. [69] note that there is no known protocol that can be applied in a streaming fashion but gets close to the complexity of Eq. (59) in the $\gamma_{\text{in}} \ll 1$ regime. In the following subsection we derive such a protocol that uses only two qudits of memory while matching the above asymptotically optimal sample complexity.

4.4.3 Improved distillation in the regime of small input fidelity via quantum PCA

Now we show that a gate-efficient procedure inspired by quantum principal component analysis (PCA) [71, 72] requires only two qudits plus three qubits of storage to output the top eigenstate with fidelity at least $1-\varepsilon_{\text{dist}}$ using

$$O\left(\left(\frac{1}{\varepsilon_{\text{dist}}} + \frac{1}{\gamma_{\text{in}}}\right)\frac{1-\gamma_{\text{in}}}{(1-\alpha)^2\gamma_{\text{in}}^2}\right)$$

copies of ρ_{in} in expectation, which matches the optimal asymptotic complexity of Eq. (59) up to a constant factor.

Intuitively speaking, the additional $1/\gamma_{\text{in}}$ term next to $1/\varepsilon_{\text{dist}}$ comes from the fact that we need to find the top eigenstate within the states stored in memory. Since the protocol of Ref. [68] stores and processes all of the required copies in parallel, a tighter analysis might reveal that it performs better in the $\varepsilon_{\text{dist}} \gg \gamma_{\text{in}}$ regime. However, once we pay the $1/\gamma_{\text{in}}$ price of postselection, we can very efficiently distill further, so the overhead does not multiply with the high-precision-induced $1/\varepsilon_{\text{dist}}$ cost. We expect that in the single-pass constant-storage setting, our protocol is essentially optimal, but we leave this problem of optimality as an open question. Finally, we speculate that one may be able to improve this protocol's sample complexity in the following way: if an attempt of locating the top eigenstate failed, one may reuse the earlier copies that were used for density matrix exponentation in earlier rounds.

Density matrix exponentiation. Quantum PCA [71, 72] is based on the core primitive of density matrix exponentiation using a fractional swap operation $\exp(-\mathrm{i}\mathbb{S}t) = \cos(t)\mathbb{I} - \mathrm{i}\sin(t)\mathbb{S}$, where $\mathbb{I}$ denotes the identity operation on a two-qudit system, and $\mathbb{S}$ denotes the swap operation of two qudits. Suppose that we have a density operator ς on systems A and S_1, and we get a copy of ϱ on system S_2 matching the dimension of S_1. The Lloyd–Mohseni–Rebentrost (LMR) density matrix exponentiation primitive applies a fractional swap of S_1 and S_2, then discards S_2, and the resulting state can be described as follows [72]:

$$\begin{aligned}
&\operatorname{Tr}_{S_2}\left[\left(\mathbb{I}_A \otimes \exp(-\mathrm{i}\mathbb{S}t)\right)\left(\varsigma \otimes \varrho\right)\left(\mathbb{I}_A \otimes \exp(\mathrm{i}\mathbb{S}t)\right)\right] \qquad (60)\\
&= \operatorname{Tr}_{S_2}[\cos^2(t)\varsigma\otimes\varrho + \mathrm{i}\cos(t)\sin(t)(\varsigma\otimes\varrho)(\mathbb{I}_A\otimes\mathbb{S}) - \mathrm{i}\cos(t)\sin(t)(\mathbb{I}_A\otimes\mathbb{S})(\varsigma\otimes\varrho) + \sin^2(t)(\mathbb{I}\otimes\mathbb{S})(\varsigma\otimes\varrho)(\mathbb{I}_A\otimes\mathbb{S})]\\
&= \cos^2(t)\varsigma + \mathrm{i}\cos(t)\sin(t)\varsigma(\mathbb{I}_A\otimes\varrho) - \mathrm{i}\cos(t)\sin(t)(\mathbb{I}_A\otimes\varrho)\varsigma + \sin^2(t)\operatorname{Tr}_{S_1}[\varsigma]\otimes\varrho. \qquad \text{(see Fig. 3)}
\end{aligned}$$

This can be viewed as density matrix exponentiation since it represents the map $\varsigma \mapsto (\mathbb{I}_A \otimes \mathrm{e}^{-\mathrm{i}\varrho t})\varsigma(\mathbb{I}_A \otimes \mathrm{e}^{\mathrm{i}\varrho t})$ up to corrections of order $O(t^2)$. The procedure consumes one copy of ϱ in order to approximately implement the unitary evolution generated by the Hermitian operator ϱ for a short time t. One can also approximately implement controlled density matrix exponentiation $\varsigma \mapsto \left(\mathbb{I}_A \otimes (|0\rangle\langle 0| \otimes \mathbb{I} + |1\rangle\langle 1| \otimes \mathrm{e}^{-\mathrm{i}\varrho' t})\right)\varsigma\left(\mathbb{I}_A \otimes (|0\rangle\langle 0| \otimes \mathbb{I} + |1\rangle\langle 1| \otimes \mathrm{e}^{\mathrm{i}\varrho' t})\right)$ by choosing $\varrho = |1\rangle\langle 1| \otimes \varrho'$; see Ref. [72, Appendix C]. We will exploit this trick in our protocol below.

For an efficient implementation of the fractional swap operation, note that the unitary $(\mathrm{e}^{\mathrm{i}\theta_+ Y} \otimes \mathbb{I})\mathrm{C}\mathbb{S}(\mathrm{e}^{-\mathrm{i}\theta_- X} \otimes \mathbb{I})$ is a block-encoding of the operator $\exp(-\mathrm{i}\mathbb{S}t)/2$, with $\theta_\pm$ defined below and $\mathrm{C}\mathbb{S}$ the controlled swap gate with the first register acting as the control; thus, the fractional swap gate can be implemented using 3 $\mathrm{C}\mathbb{S}$ gates and 4

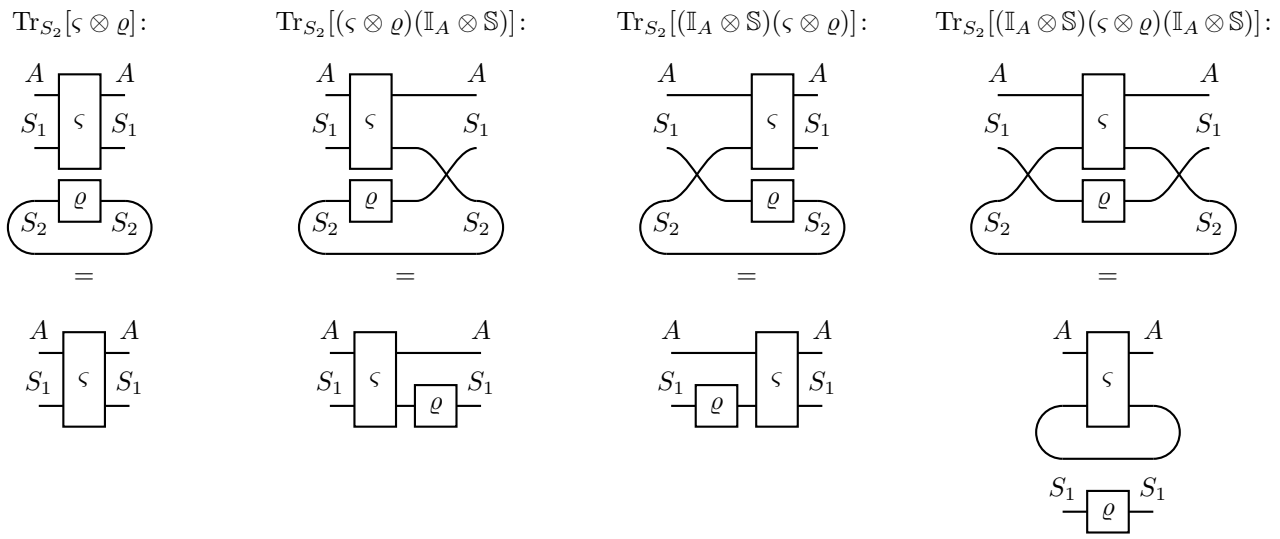

Figure 3: Diagrammatic representation [98], and simplification of the four terms in Eq. (60). For a derivation by direct computation consider that $\mathrm{Tr}_{S_2}[(\varsigma \otimes \varrho)(\mathbb{I}_A \otimes \mathbb{S})] = \sum_i (\mathbb{I}_{AS_1} \otimes \langle i|)(\varsigma \otimes \varrho)(\mathbb{I}_A \otimes \mathbb{S})(\mathbb{I}_{AS_1} \otimes |i\rangle)$ which is $\sum_i (\varsigma(\mathbb{I}_A \otimes |i\rangle)) \otimes (\langle i|\varrho) = \sum_i \varsigma(\mathbb{I}_A \otimes |i\rangle\langle i|\varrho) = \varsigma(\mathbb{I}_A \otimes \varrho)$.

single-qubit gates

$$|0\rangle\langle 0| \otimes (\cos(t)\mathbb{I} - \mathrm{i}\sin(t)\mathbb{S}) + |1\rangle\langle 1| \otimes (\mathrm{i}\sin(t)\mathbb{I} + \cos(t)\mathbb{S})$$
$$= -(\mathrm{e}^{\mathrm{i}\theta_+ Y} \otimes \mathbb{I})\,\mathrm{C}\mathbb{S}\,(\mathrm{e}^{-\mathrm{i}\theta_- X} Z \mathrm{e}^{\mathrm{i}\theta_- X} \otimes \mathbb{I})\,\mathrm{C}\mathbb{S}\,(\mathrm{e}^{-\mathrm{i}\theta_+ Y} Z \mathrm{e}^{\mathrm{i}\theta_+ Y} \otimes \mathbb{I})\,\mathrm{C}\mathbb{S}\,(\mathrm{e}^{-\mathrm{i}\theta_- X} \otimes \mathbb{I}),$$
$$\text{where} \quad \theta_\pm = \frac{\arccos\left(\frac{\cos(t)-\sin(t)}{2}\right) \pm \arccos\left(\frac{\cos(t)+\sin(t)}{2}\right)}{2},$$

which corresponds to one iteration of oblivious amplitude amplification. As elsewhere in the paper, X, Y, and Z refer to the single-qubit Pauli operators.

A simple (suboptimal) protocol for the case $\lambda_2(\rho_{\mathrm{in}}) \ll \lambda_1(\rho_{\mathrm{in}})\sqrt{\lambda_1(\rho_{\mathrm{in}})\varepsilon_{\mathrm{dist}}}$. We now show how to use a simple version of Kitaev's phase estimation [99] to distill the top eigenstate of ρ_{in} when we are promised that $\lambda_1(\rho_{\mathrm{in}}) \in [\gamma, 3\gamma]$, $\varepsilon_{\mathrm{dist}} \leq (1-\gamma)$ and $\lambda_2(\rho_{\mathrm{in}}) \ll \gamma\sqrt{\gamma\varepsilon_{\mathrm{dist}}/(1-\gamma)}$, for a known value of γ. Specifically, the protocol aims to implement the following circuit involving controlled density matrix exponentiation, which may be viewed as phase estimation to one bit of precision.

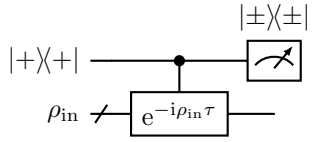

(61)

With the right choice of τ, if density matrix exponentiation is performed in small enough steps t, then there is a substantial chance of measuring the first register in $|-\rangle\langle -|$, and when this occurs, the non-principal eigenstates of the input state ρ_{in} are appropriately suppressed in the output.

The controlled density matrix exponentiation in circuit (61) is approximated with $r = \frac{\tau}{t}$ applications of the LMR procedure, giving the circuit in Fig. 4. Namely, consider the effect of the LMR protocol of Eq. (60) in the special case when the system A is not present, and the protocol is repeatedly applied on the initial state $\varsigma = \rho^{(0)} := |+\rangle\langle +| \otimes \rho_{\mathrm{in}}$ on system S_1, using copies of the amended mixed state $\varrho = |1\rangle\langle 1| \otimes \rho_{\mathrm{in}}$ on system S_2. Denote the state (on system S_1) after r iterations of the LMR protocol by $\rho^{(r)}$, which can be written as

$$\rho^{(r)} = \sum_j \sigma_j^{(r)} \otimes \lambda_j |\psi_j\rangle\langle\psi_j|, \tag{62}$$

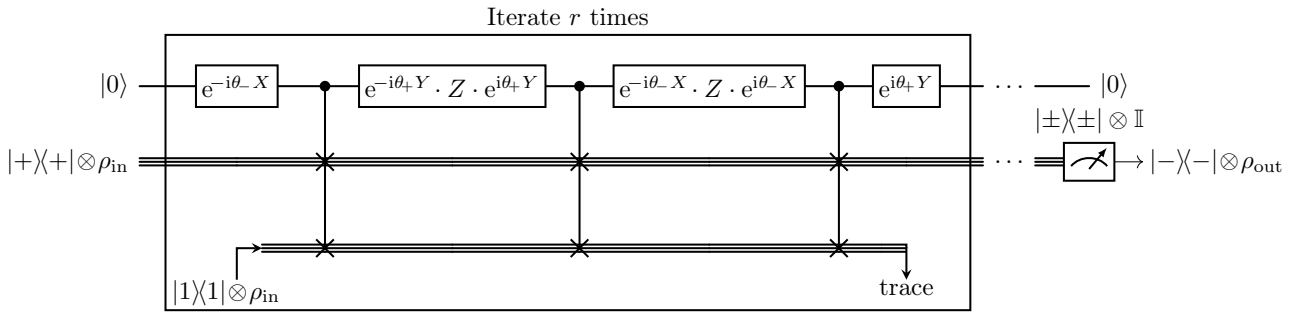

Figure 4: A simple procedure based on density matrix exponentiation for extracting the top eigenstate of an unknown density operator ρ_{in}. The procedure approximately implements one step of Kitaev's phase estimation of circuit (61). The parameters $\theta_\pm = \frac{1}{2}\arccos\left(\frac{\cos(t)-\sin(t)}{2}\right) \pm \frac{1}{2}\arccos\left(\frac{\cos(t)+\sin(t)}{2}\right)$ determine the length t of the approximated density matrix evolution-time segment per iteration. Each iteration consumes a fresh copy of ρ_{in} and the first ancilla qubit returns to state $|0\rangle$ after each iteration. After all iterations are completed, the second ancilla qubit is measured in the $|\pm\rangle$ basis and we only accept the $|-\rangle$ outcome.

where $\rho_{\text{in}} = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$ is the eigendecomposition of ρ_{in} and $\sigma_j^{(r)}$ is a normalized single-qubit density operator, for example, $\sigma_j^{(0)} = |+\rangle\langle+|$. According to Eq. (60), we find that

$$\rho^{(r+1)} = \cos^2(t)\rho^{(r)} + \mathrm{i}\cos(t)\sin(t)\rho^{(r)}(|1\rangle\langle 1|\otimes\rho_{\text{in}}) - \mathrm{i}\cos(t)\sin(t)(|1\rangle\langle 1|\otimes\rho_{\text{in}})\rho^{(r)} + \sin^2(t)(|1\rangle\langle 1|\otimes\rho_{\text{in}}),$$

in particular $\rho^{(r+1)} = \sum_j \sigma_j^{(r+1)} \otimes \lambda_j |\psi_j\rangle\langle\psi_j|$ where

$$\sigma_j^{(r+1)} = \cos^2(t)\sigma_j^{(r)} + \mathrm{i}\cos(t)\sin(t)\sigma_j^{(r)}(\lambda_j|1\rangle\langle 1|) - \mathrm{i}\cos(t)\sin(t)(\lambda_j|1\rangle\langle 1|)\sigma_j^{(r)} + \sin^2(t)\,\mathrm{Tr}\left[\sigma_j^{(r)}\right]|1\rangle\langle 1|. \tag{63}$$

This means that $\sigma_j^{(r+1)} = \Phi_j[\sigma_j^{(r)}]$ for a quantum channel Φ_j with Choi matrix

$$\begin{pmatrix} \cos^2(t) & 0 & 0 & \cos^2(t) + \mathrm{i}\lambda_j\sin(t)\cos(t) \\ 0 & \sin^2(t) & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \cos^2(t) - \mathrm{i}\lambda_j\sin(t)\cos(t) & 0 & 0 & 1 \end{pmatrix}, \tag{64}$$

which is indeed positive semidefinite for all $\lambda_j \in [0,1]$. The vectorization of the superoperator $\Phi_j[\cdot]$ is

$$\begin{pmatrix} \cos^2(t) & 0 & 0 & 0 \\ 0 & \cos^2(t) + \mathrm{i}\sin(t)\cos(t)\lambda_j & 0 & 0 \\ 0 & 0 & \cos^2(t) - \mathrm{i}\sin(t)\cos(t)\lambda_j & 0 \\ \sin^2(t) & 0 & 0 & 1 \end{pmatrix}.$$

The difference from the desired time-evolution superoperator $\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|}$ is

$$\begin{pmatrix} \sin^2(t) & 0 & 0 & 0 \\ 0 & \mathrm{e}^{\mathrm{i}\lambda_j t} - \cos^2(t) - \mathrm{i}\cos(t)\sin(t)\lambda_j & 0 & 0 \\ 0 & 0 & \mathrm{e}^{-\mathrm{i}\lambda_j t} - \cos^2(t) + \mathrm{i}\cos(t)\sin(t)\lambda_j & 0 \\ -\sin^2(t) & 0 & 0 & 0 \end{pmatrix}. \tag{65}$$

For all $t \in [0, \pi/2]$ the operator norm of the matrix in Eq. (65) is $\sqrt{2}\sin^2(t)$, since

$$\begin{aligned} |\mathrm{e}^{\mathrm{i}\lambda_j t} - \cos^2(t) - \mathrm{i}\cos(t)\sin(t)\lambda_j| &= \sqrt{(\cos(\lambda_j t) - \cos^2(t))^2 + (\sin(\lambda_j t) - \cos(t)\sin(t)\lambda_j)^2} && (66) \\ &\leq \sqrt{(1-\cos^2(t))^2} && \text{(monotonically decreasing in } \lambda_j) \\ &= \sin^2(t)\,. && (67) \end{aligned}$$

In the above inequality we used that Eq. (66) is monotonically decreasing for all $\lambda_j \in [0,1]$ when $t \in [0, \pi/2]$. Indeed, the derivative of the square of Eq. (66) in λ_j can be computed as follows:

$$2\cos(t)\left(\underbrace{(t\cos(t) - \sin(t))}_{=\int_0^t -t\sin(t)\mathrm{d}t \leq 0}\sin(t\lambda_j) + \lambda_j \sin(t)\underbrace{\left(\sin(t)\cos(t) - t\underbrace{\cos(t\lambda_j)}_{\geq \cos(t)}\right)}_{\leq 0}\right) \leq 0\,.$$

Let us now consider superoperator norms induced by Schatten p-norms [100] for some $p \in [1, \infty]$. By norm conversion and the observation that the operator norm of Eq. (65) is $\sqrt{2}\sin^2(t)$ due to Eq. (67), we directly get that

$$\|\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \Phi_j[\cdot]\|_{1\text{–}1} \leq \sqrt{2}\|\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \Phi_j[\cdot]\|_{2\text{–}2} = 2\sin^2(t) \leq 2t^2. \tag{68}$$

Due to the contractiveness of quantum channels in the Schatten 1-norm we get that after r iterations

$$\|\sigma_j^{(r)} - \mathrm{e}^{-\mathrm{i}r\lambda_j t|1\rangle\langle 1|}[\sigma_j^{(0)}]\mathrm{e}^{\mathrm{i}r\lambda_j t|1\rangle\langle 1|}\|_1 \leq 2rt^2. \tag{69}$$

Proposition 9. *Suppose we are given real numbers $\gamma, \varepsilon_{\mathrm{dist}} \in (0,1)$, and copies of a qudit density operator ρ_{in} such that $\lambda_1(\rho_{\mathrm{in}}) \in [\gamma, 3\gamma]$, $\varepsilon_{\mathrm{dist}} \leq 1-\gamma$ and $\lambda_2(\rho_{\mathrm{in}}) \leq \gamma\sqrt{\frac{8\gamma\varepsilon_{\mathrm{dist}}}{3\pi^2(1-\gamma)}}$. Choosing $r = \left\lceil \frac{3\pi^2(1-\gamma)}{2\gamma^3\varepsilon_{\mathrm{dist}}} \right\rceil$ and $t = \frac{\pi}{2r\gamma}$ the protocol of Fig. 4 succeeds with probability at least $\frac{\gamma}{3}$, and upon success produces a state $\rho_-^{(r)}$ such that $[\rho_{\mathrm{in}}, \rho_-^{(r)}] = 0$, and $\langle\Xi|\rho_-^{(r)}|\Xi\rangle \geq 1 - \varepsilon_{\mathrm{dist}}$, where $|\Xi\rangle$ is the principal eigenvector of ρ_{in}. The protocol uses an expected number of copies of ρ_{in} which is at most $\frac{3}{\gamma}(r+1)$, and the same order of controlled qudit swap and single-qubit gates. The space complexity is two qudits and three qubits of storage.*

Proof. The expected complexity bound follows from Fig. 4 and the success probability bound $\geq \frac{\gamma}{3}$, so it suffices to prove this latter bound. By Eq. (69), we get that projecting down $\rho^{(r)}$ to the $|-\rangle\langle -|$ ancilla state we get the subnormalized state

$$\begin{aligned} (\langle -| \otimes \mathbb{I})\rho^{(r)}(|-\rangle \otimes \mathbb{I}) &= \sum_j \langle -|\sigma_j^{(r)}|-\rangle \otimes \lambda_j |\psi_j\rangle\langle\psi_j| && (70) \\ &= \sum_j \lambda_j |\psi_j\rangle\langle\psi_j| \Big(\underbrace{|\langle -|\mathrm{e}^{-\mathrm{i}r\lambda_j t|1\rangle\langle 1|}|+\rangle|^2}_{= \frac{|1-\mathrm{e}^{-\mathrm{i}r\lambda_j t}|^2}{4} \leq \frac{|r\lambda_j t|^2}{4}} + \underbrace{\chi_j}_{|\cdot| \leq rt^2} \Big). && (71) \end{aligned}$$

Let $\rho_-^{(r)} := (\langle -|\otimes\mathbb{I})\rho^{(r)}(|-\rangle\otimes\mathbb{I})/\operatorname{Tr}\big[(\langle -|\otimes\mathbb{I})\rho^{(r)}(|-\rangle\otimes\mathbb{I})\big]$ be the state we get by postselecting on the $|-\rangle$ outcome of the ancilla measurement in the $|\pm\rangle$ basis. From Eq. (63) it is evident that $\rho_-^{(r)}$ commutes with ρ_{in}. Since $\lambda_1(\rho_{\text{in}}) \in [\gamma, 3\gamma]$ and $rt = \frac{\pi}{2\gamma}$ we get that the probability that we measure the ancilla qubit in the $|-\rangle$ state is

$$\operatorname{Tr}\Big[(\langle -|\otimes\mathbb{I})\rho^{(r)}(|-\rangle\otimes\mathbb{I})\Big] \geq \lambda_1\langle -|\sigma_1^{(r)}|-\rangle \geq \gamma\langle -|\sigma_1^{(r)}|-\rangle \tag{72}$$

$$\text{where } \langle -|\sigma_1^{(r)}|-\rangle \geq \Big(\frac{|1-\mathrm{e}^{-\mathrm{i}r\lambda_1 t}|^2}{4} - rt^2\Big) \geq \Big(\frac{1}{2} - \frac{\pi^2}{4\gamma^2 r}\Big). \qquad (\text{since } \lambda_1 \in [\gamma,3\gamma] \text{ and } rt = \tfrac{\pi}{2\gamma})$$

By our choice of r we get that $\frac{\pi^2}{4\gamma^2 r} \leq \frac{\gamma\varepsilon_{\text{dist}}}{6(1-\gamma)} \leq \frac{1}{6}$ and (recalling the principal eigenvector $|\Xi\rangle$ is denoted $|\psi_1\rangle$)

$$\begin{aligned}
\langle\psi_1|\rho_-^{(r)}|\psi_1\rangle &= \frac{\lambda_1\langle -|\sigma_1^{(r)}|-\rangle}{\lambda_1\langle -|\sigma_1^{(r)}|-\rangle + \sum_{j>1}\lambda_j \underbrace{\langle -|\sigma_j^{(r)}|-\rangle}_{\leq \frac{|r\lambda_j t|^2}{4} + rt^2}} \\
&\geq \frac{\lambda_1\langle -|\sigma_1^{(r)}|-\rangle}{\lambda_1\underbrace{\langle -|\sigma_1^{(r)}|-\rangle}_{\geq \frac{1}{3} \text{ by (72)}} + \underbrace{(1-\lambda_1)}_{\leq 1-\gamma}\underbrace{\Big(\Big|\frac{\pi\lambda_2}{4\gamma}\Big|^2 + \frac{\pi^2}{4\gamma^2 r}\Big)}_{\leq \frac{\gamma\varepsilon_{\text{dist}}}{3(1-\gamma)}}} \qquad (\text{since } rt = \tfrac{\pi}{2\gamma},\ \lambda_2^2 \leq \tfrac{8\gamma^3\varepsilon_{\text{dist}}}{3\pi^2(1-\gamma)}) \\
&\geq \frac{\lambda_1\langle -|\sigma_1^{(r)}|-\rangle}{(1+\varepsilon_{\text{dist}})\lambda_1\langle -|\sigma_1^{(r)}|-\rangle} = \frac{1}{1+\varepsilon_{\text{dist}}} \geq 1-\varepsilon_{\text{dist}}. \qquad \square
\end{aligned}$$

Note that if we only know that $\lambda_1 \geq \gamma$ and $\lambda_2 \leq \gamma\sqrt{\frac{8\gamma\varepsilon_{\text{dist}}}{3\pi^2(1-\gamma)}}$, but don't know whether $\lambda_1 \leq 3\gamma$, we can still perform the distillation using the above protocol through combining it with standard techniques, such as exponential search to guess the right order of $\frac{\lambda_1}{\gamma}$. The resulting protocol shall even have the same asymptotic complexity up to constant factors.

An improved protocol for the general case. We can improve the overhead in the previous protocol by recursively filtering the smaller eigenvalues in a similar fashion to Ref. [101]. As we show, it suffices to filter a constant fraction of the unwanted eigenstates in each iteration. This relieves the burden of error magnification due to the postselection on a small probability $\approx \lambda_1$ event hindering the previous simple variant described in the proof of Proposition 9. For this purpose, we can use Kitaev's phase estimation [99] combined by standard error reduction techniques, which requires a total simulation time of $\Theta(\frac{\log(1/\epsilon)}{\delta})$ in controlled density matrix exponentiation to output, with probability at least $1-\epsilon$, a phase estimate that is less than δ off [102]. For practical considerations one might also consider using improved iterative phase estimation variants [103] or eigenstate filtering techniques via quantum singular value transformation or quantum signal processing [104, 105].

The controlled Hamiltonian simulation can be performed using the same circuit as before (Fig. 4), however the subsequent applications of the protocol require a slightly adapted analysis because we need to track the state of multiple qubits and/or the entire measurement history.

Proposition 10. *Suppose we are given real numbers $\gamma, \varepsilon_{\text{dist}}, \alpha \in (0,1)$, and can request copies of a qudit density operator ρ_{in} such that $\lambda_1(\rho_{\text{in}}) \geq \gamma$, $\varepsilon_{\text{dist}} < (1-\gamma)$ and $\lambda_2(\rho_{\text{in}}) \leq \alpha\gamma$. By iteratively applying the circuit of Fig. 4 with appropriate choices of r and t we can prepare a state ρ_{out} such that $[\rho_{\text{in}}, \rho_{\text{out}}] = 0$, and $\langle\Xi|\rho_{\text{out}}|\Xi\rangle \geq 1-\varepsilon_{\text{dist}}$, where $|\Xi\rangle$ is the principal eigenvector of ρ_{in}. The protocol uses an expected number of copies of ρ_{in} which is at most $O\Big(\frac{1-\gamma}{(1-\alpha)^2\gamma^2}\Big(\frac{1}{\varepsilon_{\text{dist}}} + \frac{1}{\gamma}\Big)\Big)$, and the same order of controlled qudit swap gates and single-qubit gates. The space complexity is two qudits and three qubits of storage.*

The core of our analysis is to understand what happens when we apply the LMR controlled density matrix exponentiation protocol to a mixed state whose reduced density matrix commutes with $\rho_{\text{in}} = \sum_j \lambda_j|\psi_j\rangle\langle\psi_j|$, where $|\Xi\rangle = |\psi_1\rangle$.

Lemma 2. *Let A label a system of arbitrary finite dimension, and let C label a two-dimensional (qubit) system. Suppose that we have a quantum algorithm that receives as input a normalized state $\rho^{(\text{start})} := \sum_j \sigma_j^{(\text{start})} \otimes |\psi_j\rangle\langle\psi_j|$, where $\sigma_j^{(\text{start})}$ are subnormalized quantum states on the AC register and $|\psi_j\rangle\langle\psi_j|$ are the eigenstates of ρ_{in}. If the*

algorithm only interacts with the final register through controlled Hamiltonian simulation $\mathbb{I}_A \otimes |0\rangle\langle 0|_C \otimes \mathbb{I} + \mathbb{I}_A \otimes |1\rangle\langle 1|_C \otimes \mathrm{e}^{-\mathrm{i}\rho_{\mathrm{in}}\tau_k}$*, then the output state can be written as* $\rho^{(\mathrm{end})} := \sum_j \sigma_j^{(\mathrm{end})} \otimes |\psi_j\rangle\langle\psi_j|$*. Moreover, if the total simulation time is* $T = \sum_k \tau_k$ *and we approximate each such controlled Hamiltonian simulation step by the LMR protocol with step size at most* $t \leq \frac{\pi}{2}$*, then the output state can be written as* $\tilde{\rho}^{(\mathrm{end})} := \sum_j \tilde{\sigma}_j^{(\mathrm{end})} \otimes |\psi_j\rangle\langle\psi_j|$*, where it holds that* $\|\tilde{\sigma}_j^{(\mathrm{end})} - \sigma_j^{(\mathrm{end})}\|_1 \leq 3Tt \max(\mathrm{Tr}\left[\sigma_j^{(\mathrm{start})}\right], \lambda_j)$*.*

Proof. Operations that do not touch the final register clearly preserve the diagonal form $\sum_j \sigma_j \otimes |\psi_j\rangle\langle\psi_j|$ of quantum states. Controlled Hamiltonian simulation can be equivalently described as an operation controlled by eigenstates on the final register, that is, $|0\rangle\langle 0|_C \otimes \mathbb{I} + |1\rangle\langle 1|_C \otimes \mathrm{e}^{-\mathrm{i}\rho_{\mathrm{in}}\tau} = \sum_j \mathrm{e}^{\lambda_j |1\rangle\langle 1|\tau} \otimes |\psi_j\rangle\langle\psi_j|$, which therefore also preserves the diagonal form, proving that we can write $\rho^{(\mathrm{end})} = \sum_j \sigma_j^{(\mathrm{end})} \otimes |\psi_j\rangle\langle\psi_j|$.

Now consider what happens when we apply a density matrix exponentiation step of Eq. (60) on $\varsigma = \rho^{(r)} := \sum_j \tilde{\sigma}_j^{(r)} \otimes |\psi_j\rangle\langle\psi_j|$ using mixed state $\varrho = |1\rangle\langle 1|_C \otimes \rho_{\mathrm{in}}$. According to Eq. (60) we get that

$$\begin{aligned}\rho^{(r+1)} = \cos^2(t)\rho^{(r)} + \mathrm{i}\cos(t)\sin(t)\rho^{(r)}(\mathbb{I}_A \otimes |1\rangle\langle 1|_C \otimes \rho_{\mathrm{in}}) \\ - \mathrm{i}\cos(t)\sin(t)(\mathbb{I}_A \otimes |1\rangle\langle 1|_C \otimes \rho_{\mathrm{in}})\rho^{(r)} + \sin^2(t)(\tilde{\sigma}^{(r)} \otimes |1\rangle\langle 1|_C \otimes \rho_{\mathrm{in}})\,,\end{aligned}$$

where $\tilde{\sigma}^{(r)}$ is a normalized state on the A register defined by

$$\tilde{\sigma}^{(r)} = \mathrm{Tr}_C\left[\sum_j \tilde{\sigma}_j^{(r)}\right] = \sum_j (\mathbb{I}_A \otimes \langle 0|_C)\tilde{\sigma}_j^{(r)}(\mathbb{I}_A \otimes |0\rangle_C) + (\mathbb{I}_A \otimes \langle 1|_C)\tilde{\sigma}_j^{(r)}(\mathbb{I}_A \otimes |1\rangle_C)\,.$$

From this we can see that $\rho^{(r+1)} := \sum_j \tilde{\sigma}_j^{(r+1)} \otimes |\psi_j\rangle\langle\psi_j|$ where

$$\tilde{\sigma}_j^{(r+1)} = \underbrace{\cos^2(t)\tilde{\sigma}_j^{(r)} + \mathrm{i}\cos(t)\sin(t)\lambda_j\tilde{\sigma}_j^{(r)}(\mathbb{I}_A \otimes |1\rangle\langle 1|_C) - \mathrm{i}\cos(t)\sin(t)\lambda_j(\mathbb{I}_A \otimes |1\rangle\langle 1|_C)\tilde{\sigma}_j^{(r)}}_{\widetilde{\Phi}_j[\tilde{\sigma}_j^{(r)}]:=} + \sin^2(t)\lambda_j(\tilde{\sigma}^{(r)} \otimes |1\rangle\langle 1|_C). \tag{73}$$

Since $\rho^{(r+1)} \succeq 0$ we immediately get that $\tilde{\sigma}_j^{(r+1)} \succeq 0$. We also get that

$$\|\tilde{\sigma}_j^{(r+1)}\|_1 = \mathrm{Tr}\left[\tilde{\sigma}_j^{(r+1)}\right] = \cos^2(t)\,\mathrm{Tr}\left[\tilde{\sigma}_j^{(r)}\right] + \sin^2(t)\lambda_j \leq \max\left(\mathrm{Tr}\left[\tilde{\sigma}_j^{(r)}\right], \lambda_j\right). \tag{74}$$

If we start the procedure with $\rho^{(0)}$, consequently by induction we get that after r iterations we have

$$\mathrm{Tr}\left[\tilde{\sigma}_j^{(r)}\right] \leq \max\left(\mathrm{Tr}\left[\tilde{\sigma}_j^{(0)}\right], \lambda_j\right).$$

Note that the channel $\widetilde{\Phi}_j$ defined in Eq. (73) may be written as $\widetilde{\Phi}_j = \mathcal{I}_A \otimes \widetilde{\Phi}_j^{(C)}$, where $\mathcal{I}_A$ is the identity channel on system A and

$$\widetilde{\Phi}_j^{(C)}[\cdot] = \cos^2(t)[\cdot] + \mathrm{i}\cos(t)\sin(t)\lambda_j[\cdot]|1\rangle\langle 1| - \mathrm{i}\cos(t)\sin(t)\lambda_j|1\rangle\langle 1|[\cdot]\,. \tag{75}$$

Let us now consider the difference of $\widetilde{\Phi}_j[\cdot]$ and $\mathcal{I}_A[\cdot] \otimes \mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|}$.

$$\begin{aligned}\|\mathcal{I}_A[\cdot] \otimes \mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \widetilde{\Phi}_j[\cdot]\|_{1\text{–}1} &\leq \|\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \widetilde{\Phi}_j^{(C)}[\cdot]\|_\diamond \\ &\leq 2\|\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \widetilde{\Phi}_j^{(C)}[\cdot]\|_{2\text{–}2} \\ &= 2\sin^2(t),\end{aligned} \tag{76}$$

where the first inequality follows from the definition of the diamond norm, the second inequality follows from norm conversion from the diamond norm to the 2–2 norm for superoperators on 2-dimensional systems [106, Appendix C], and the last equality follows by applying Eq. (67) to evaluate the operator norm of the matrix representation

of the superoperator $\mathrm{e}^{-\mathrm{i}\lambda_j t|1\rangle\langle 1|}[\cdot]\mathrm{e}^{\mathrm{i}\lambda_j t|1\rangle\langle 1|} - \widetilde{\Phi}_j^{(C)}[\cdot]$, given by

$$\begin{pmatrix} \sin^2(t) & 0 & 0 & 0 \\ 0 & \mathrm{e}^{\mathrm{i}\lambda_j t} - \cos^2(t) - \mathrm{i}\cos(t)\sin(t)\lambda_j & 0 & 0 \\ 0 & 0 & \mathrm{e}^{-\mathrm{i}\lambda_j t} - \cos^2(t) + \mathrm{i}\cos(t)\sin(t)\lambda_j & 0 \\ 0 & 0 & 0 & \sin^2(t) \end{pmatrix}. \tag{77}$$

It follows (dropping the subscripts A and C for brevity) that

$$\|\tilde{\sigma}_j^{(r)} - \mathrm{e}^{-\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(0)}\mathrm{e}^{\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 \leq 3r\sin^2(t)\max\Big(\mathrm{Tr}\Big[\tilde{\sigma}_j^{(0)}\Big], \lambda_j\Big).$$

Indeed, this trivially holds for $r = 0$, and we can prove it by induction for $r \geq 1$ as follows

$$\begin{aligned}
&\|\tilde{\sigma}_j^{(r+1)} - \mathrm{e}^{-\mathrm{i}(r+1)\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(0)}\mathrm{e}^{\mathrm{i}(r+1)\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 \\
&\leq \|\tilde{\sigma}_j^{(r+1)} - \widetilde{\Phi}_j[\tilde{\sigma}_j^{(r)}]\|_1 + \|\widetilde{\Phi}_j[\tilde{\sigma}_j^{(r)}] - \mathrm{e}^{-\mathrm{i}(r+1)\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(0)}\mathrm{e}^{\mathrm{i}(r+1)\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 \\
&\leq \sin^2(t)\lambda_j + \|\widetilde{\Phi}_j[\tilde{\sigma}_j^{(r)}] - \mathrm{e}^{-\mathrm{i}\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(r)}\mathrm{e}^{\mathrm{i}\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 && \text{(by Eq. (73))} \\
&\qquad + \|\mathrm{e}^{-\mathrm{i}\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}[\tilde{\sigma}_j^{(r)} - \mathrm{e}^{-\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(0)}\mathrm{e}^{\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}]\mathrm{e}^{\mathrm{i}\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 \\
&\leq \sin^2(t)\lambda_j + 2\sin^2(t)\|\tilde{\sigma}_j^{(r)}\|_1 && \text{(by Eq. (76))} \\
&\qquad + \|\tilde{\sigma}_j^{(r)} - \mathrm{e}^{-\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\tilde{\sigma}_j^{(0)}\mathrm{e}^{\mathrm{i}r\lambda_j t(\mathbb{I}\otimes|1\rangle\langle 1|)}\|_1 \\
&\leq 3(r+1)\sin^2(t)\max\Big(\mathrm{Tr}\Big[\tilde{\sigma}_j^{(0)}\Big], \lambda_j\Big). && \text{(by Eq. (74) and induction)}
\end{aligned}$$

Thus, when approximating controlled Hamiltonian simulation $\mathbb{I}_A \otimes \mathrm{e}^{-\mathrm{i}\tau_k(|1\rangle\langle 1|\otimes\rho_{\mathrm{in}})}$ with $t' \leq t$ step size we need $r = \frac{\tau_k}{t'}$ steps, inducing an overall error in Schatten 1-norm upper bounded by $3\frac{\tau_k}{t'}t'^2\max\Big(\mathrm{Tr}\Big[\tilde{\sigma}_j^{(0)}\Big], \lambda_j\Big) \leq 3\tau_k t\max\Big(\mathrm{Tr}\Big[\tilde{\sigma}_j^{(0)}\Big], \lambda_j\Big)$. Other steps that do not depend on the final register can be described by quantum channels, which are contractive with respect to the 1-norm. We can conclude that $\|\tilde{\sigma}_j^{(\mathrm{end})} - \sigma_j^{(\mathrm{end})}\|_1 \leq 3\sum_k \tau_k t\max(\mathrm{Tr}\Big[\sigma_j^{(\mathrm{start})}\Big], \lambda_j)$, as claimed. □

Proof of Proposition 10. The protocol will proceed through a sequence of ℓ iterations, each of which aims to make progress on amplifying the principal eigenstate, but has some probability of failure. If a round fails, the procedure is restarted from the beginning. First, we establish some properties of the starting and ending state for an individual iteration.

Suppose that at the beginning of an iteration, we have a quantum state $\rho^{(\mathrm{start})} := \sum_j p_j^{(\mathrm{start})}|\psi_j\rangle\langle\psi_j|$. Performing Kitaev's phase estimation with precision δ and success probability at least $1-\epsilon$ requires $T = \Theta(\frac{\log(1/\epsilon)}{\delta})$ controlled Hamiltonian simulation time [102]. Kitaev's phase estimation is accomplished using only 1 ancilla qubit (denoted system C), through a sequence of circuits like circuit (61) for different values of $\tau = r \cdot t$, which result in a single-bit measurement outcome stored in classical memory. We may equivalently view the system A as initially holding a set of fresh ancilla qubits in the state $|+\rangle$, and for each application of circuit (61), one of these ancilla qubits is swapped into register C and then swapped back after being measured. The swapping in of these fresh qubits and the postprocessing of measurement outcomes is entirely classical and does not incur any additional quantum gates. Viewed this way, it is clear that the full (boosted) Kitaev phase estimation procedure only interacts with the register holding $\rho^{(\mathrm{start})}$ through controlled Hamiltonian simulation and thus, when the controlled Hamiltonian simulation is implemented via the LMR density matrix exponentiaton protocol, Lemma 2 applies. If the stepsize is $t = \frac{\zeta}{3T} < \pi/2$, then a total of $\Theta\Big(\frac{\log^2(1/\epsilon)}{\delta^2\zeta}\Big)$ copies of ρ_{in} are consumed via the protocol of Fig. 4. The gate complexity is 4 single-qubit gates and 3 controlled qudit swap gates per copy consumed (with no additional gate complexity associated to the input state); this tells us that it suffices to bound the number of consumed copies. We will set the precision of the phase estimation to $\delta = \frac{(1-\alpha)\gamma}{2}$, and say that the round succeeds if the energy estimate register (a classical register) is greater than $\frac{1+\alpha}{2}\gamma$ (in order to distinguish $\lambda_1 \geq \gamma$ from $\lambda_2 \leq \alpha\gamma$ with

failure probability at most ϵ). Let $\rho^{(\mathrm{proj})} = \sum_j p_j^{(\mathrm{proj})} |\psi_j\rangle\langle\psi_j|$ be the subnormalized output state of the LMR-based algorithm when projected down to the energy estimate register being greater than $\frac{1+\alpha}{2}\gamma$ and tracing out all ancilla registers. If $p_1^{(\mathrm{start})} \geq \lambda_1/2$, then according to Lemma 2, we get that

$$p_1^{(\mathrm{proj})} \geq p_1^{(\mathrm{start})}(1-\epsilon-3Tt) \geq p_1^{(\mathrm{start})}(1-\epsilon-\zeta)\,, \tag{78}$$

$$\text{and for } j>1, \qquad p_j^{(\mathrm{proj})} \in [0, \epsilon p_j^{(\mathrm{start})} + 3Tt\max(p_j^{(\mathrm{start})}, \lambda_j)/2)] \subseteq [0, (\epsilon+\frac{\zeta}{2})p_j^{(\mathrm{start})} + \frac{\zeta}{2}\lambda_j]. \tag{79}$$

The probability that the round succeeds is denoted by

$$p^{(\mathrm{succ})} = \sum_j p_j^{(\mathrm{proj})} \tag{80}$$

and the output state conditioned on success is denoted by

$$\rho^{(\mathrm{end})} = \frac{\rho^{(\mathrm{proj})}}{p^{(\mathrm{succ})}} := \sum_j p_j^{(\mathrm{end})} |\psi_j\rangle\langle\psi_j| \tag{81}$$

From Eq. (78), this state satisfies

$$p_1^{(\mathrm{end})} = \frac{p_1^{(\mathrm{proj})}}{p^{(\mathrm{succ})}} \geq \frac{1-\epsilon-\zeta}{p^{(\mathrm{succ})}} p_1^{(\mathrm{start})}, \tag{82}$$

This equation shows how performing (approximate) phase estimation and postselecting on energy estimates above $\frac{1+\alpha}{2}\gamma$ can lead to magnification of the overlap p_1 with the principal eigenvector, depending on parameters ϵ and ζ.

The full protocol iterates this process over ℓ rounds, indexed by $i = 1, 2, \ldots, \ell$. Let $\rho^{(\mathrm{start},i)}$ denote the starting state for iteration i and $\rho^{(\mathrm{end},i)}$ the ending state, which is taken to be the starting state $\rho^{(\mathrm{start},i+1)}$ at the next iteration. For brevity of notation, let $\rho^{(i)} = \rho^{(\mathrm{end},i)}$, with the starting state at iteration $i=1$ taken to be $\rho^{(0)} = \rho^{(\mathrm{start},1)} = \rho_{\mathrm{in}}$. In the i-th iteration, we set the failure probability in Kitaev's phase estimation to ϵ_i and the stepsize to $t_i = \frac{\zeta_i}{3T}$, which determines the success probability $p^{(\mathrm{succ},i)}$ of the ith iteration. The overlap with the principal eigenvector—the key figure of merit—after the i-th round is denoted $p_1^{(i)} \equiv p_1^{(\mathrm{end},i)} \equiv p_1^{(\mathrm{start},i+1)} := \langle\psi_1|\rho^{(i)}|\psi_1\rangle$. As long as we maintain $p_1^{(i)} \geq \lambda_1/2$, the bound in Eq. (82) holds, and for any value of i, we derive the following bound on the inverse overall success probability of all iterations $i+1, i+2, \ldots, \ell$:

$$1 \geq p_1^{(\ell)} \geq \prod_{k=i+1}^{\ell} \frac{1-\epsilon_k-\zeta_k}{p^{(\mathrm{succ},k)}} p_1^{(i)} \Longrightarrow \prod_{k=i+1}^{\ell} \frac{1}{p^{(\mathrm{succ},k)}} \leq \frac{1}{p_1^{(i)}} \prod_{k=i+1}^{\ell} \frac{1}{1-\epsilon_k-\zeta_k} \leq \frac{1}{p_1^{(i)}} \frac{1}{1-\sum_{k=i+1}^{\ell}(\epsilon_k+\zeta_k)}. \tag{83}$$

On the other hand, note that by Eq. (78) we have that for all $i = 1, \ldots, \ell$,

$$\begin{aligned} p_1^{(i)} = \frac{p_1^{(\mathrm{proj},i)}}{p_1^{(\mathrm{proj},i)} + \sum_{j>1} p_j^{(\mathrm{proj},i)}} &\geq \frac{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i)}{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i) + \sum_{j>1} p_j^{(\mathrm{proj},i)}} \quad \text{(by monotonicity of } \tfrac{x}{x+c} \text{ and Eq. (78))} \\ &\geq \frac{p_1^{(\mathrm{start},i)}}{p_1^{(\mathrm{start},i)} + \frac{\epsilon_i+\zeta_i}{1-\epsilon_i-\zeta_i}} = \frac{p_1^{(i-1)}}{p_1^{(i-1)} + \frac{\epsilon_i+\zeta_i}{1-\epsilon_i-\zeta_i}}. \end{aligned} \tag{84}$$

where in the last line we have used that $\sum_{j>1} p_j^{(\mathrm{proj},i)} \leq \epsilon_i + \zeta_i$ as a consequence of Eq. (79). Let $c := \frac{\epsilon_i+\zeta_i}{1-\epsilon_i-\zeta_i}$ and note that the function $\frac{p}{p+c}$ in Eq. (84) is monotone increasing on the interval $p \in [0,1]$ for all fixed $c \geq 0$, that is, when $\epsilon_i + \zeta_i < 1$; and also monotone decreasing in $c > 0$ for all fixed $p \in [0,1]$. For $p = \frac{1}{3}$ and $\epsilon_i + \zeta_i = \frac{1}{8}$ it evaluates to $\frac{p}{p+c} = \frac{7}{10} > \frac{2}{3}$, so we can conclude that if $\epsilon_i + \zeta_i \leq \frac{1}{8}$ and $p_1^{(i-1)} \geq \frac{1}{3}$, then $p^{(i)} > \frac{2}{3}$. If $p_1^{(i-1)} \leq \frac{1}{3}$ and $\epsilon_i + \zeta_i \leq \frac{1}{8}$, then we also have that $p^{(i)} \geq 2p_1^{(i-1)}$, so that $p_1^{(0)} \geq \min(\lambda_1, \frac{1}{2})$ implies that $p_1^{(i)} \geq \lambda_1/2$ holds for all $i \geq 0$.

The $\varepsilon_{\mathrm{dist}} > \frac{1}{3}$ case. Note that this is a subset of the case $\gamma < \frac{2}{3}$. Here we will conveniently set $\zeta_i := \epsilon_i \leq \frac{1}{16}$ so that starting with $\rho^{(0)} := \rho_{\mathrm{in}}$, for which $p_1^{(0)} \geq \gamma$ by assumption, we get that $\ell := \lceil \log_2(\frac{1}{3\gamma}) \rceil + 1$ successful iterations of the procedure suffice to achieve $p_1^{(\ell)} > \frac{2}{3}$. This follows by noting that $p_1^{(i)} \geq 2^i\gamma$ holds for every $i \in \{0, 1, \ldots, \ell-1\}$,

a fact that can be shown by induction and Eq. (84), which implies $p^{(\ell-1)} \geq 1/3$, and thus via the observation above also that $p^{(\ell)} \geq 2/3$.

We choose the failure probability and step size parameters to decrease with i as $\epsilon_i = \zeta_i := \frac{2^{(1-i)/2}}{16}$, so that $\sum_{i=1}^{\ell} \epsilon_i = \sum_{i=1}^{\ell} \zeta_i \leq \frac{1}{4}$, and Eq. (83) gives $\prod_{k=i+1}^{\ell} \frac{1}{p^{(\mathrm{succ},k)}} \leq \frac{2}{p_1^{(i)}}$.

Let $C^{(i)}$ denote the expected copy complexity to successfully complete the procedure up to and including the i-th iteration. We have $C^{(0)} = 1$ and $C^{(i+1)} = \frac{C^{(i)} + \Theta\left(\frac{\log^2(1/\epsilon_i)}{\delta^2 \epsilon_i}\right)}{p^{(\mathrm{succ},i)}}$, from which we can see by induction that

$$C^{(\ell)} = \Theta\left(\prod_{i=1}^{\ell} \frac{1}{p^{(\mathrm{succ},i)}} + \sum_{i=1}^{\ell} \frac{\log^2(1/\epsilon_i)}{\delta^2 \epsilon_i} \prod_{k=i}^{\ell} \frac{1}{p^{(\mathrm{succ},k)}}\right) \leq O\left(\frac{1}{p_1^{(0)}} + \sum_{i=1}^{\ell} \frac{\log^2(1/\epsilon_i)}{\delta^2 \epsilon_i p_1^{(i-1)}}\right) \leq O\left(\frac{1}{\gamma} + \sum_{i=1}^{\ell} \frac{i^2 2^{-i/2}}{\delta^2 \gamma}\right) = O\left(\frac{1}{\delta^2 \gamma}\right).$$

We can conclude that for $\gamma < \frac{2}{3}$ we can prepare a state $\rho_{\mathrm{out}} = \rho^{(\ell)}$ such that $\langle \psi_1 | \rho_{\mathrm{out}} | \psi_1 \rangle > \frac{2}{3}$ with an expected $O(\frac{1}{\delta^2 \gamma}) = O(\frac{1}{(1-\alpha)^2 \gamma^3})$ copy complexity.

The $\varepsilon_{\mathrm{dist}} \leq \frac{1}{3}$ case. If $\gamma < \frac{2}{3}$ we first run the above protocol which produces a state with $p_1 \geq \frac{2}{3}$, that we will call $\rho^{(0)}$ for the purposes of this subcase analysis. If $\gamma \geq \frac{2}{3}$ we can simply set $\rho^{(0)}$ to be ρ_{in}. In both cases, we can say that the preparation of $\rho^{(0)}$ has complexity $C = O(\frac{1}{(1-\alpha)^2 \gamma^3})$, since in the $\gamma \geq \frac{2}{3}$ case, it holds that $(1-\alpha) \geq \frac{1}{2}$ without loss of generality.

Let $\eta^{(i)} \equiv \eta^{(\mathrm{end},i)} := 1 - p_1^{(\mathrm{end},i)} \equiv 1 - p_1^{(i)}$ and $\eta^{(\mathrm{start},i)} := 1 - p_1^{(\mathrm{start},i)} \equiv 1 - p^{(i-1)}$. Similarly to Eq. (84), by Eqs. (78) and (79) we get for $i = 1, \ldots, \ell$ that

$$\begin{aligned}
\eta^{(i)} = 1 - \frac{p_1^{(\mathrm{proj},i)}}{p_1^{(\mathrm{proj},i)} + \sum_{j>1} p_j^{(\mathrm{proj},i)}} &\leq 1 - \frac{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i)}{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i) + \sum_{j>1} p_j^{(\mathrm{proj},i)}} \\
&\qquad \text{(by monotonicity of } \tfrac{x}{x+c} \text{ and Eq. (78))} \\
&\leq 1 - \frac{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i)}{p_1^{(\mathrm{start},i)}(1-\epsilon_i-\zeta_i) + (\epsilon_i+\zeta_i)(1-p_1^{(\mathrm{start},i)}) + \zeta_i(1-\lambda_1)} \quad \text{(by Eq. (79))} \\
&= \frac{(\epsilon_i+\zeta_i)\eta^{(\mathrm{start},i)} + \zeta_i(1-\lambda_1)}{(1-\eta^{(\mathrm{start},i)})(1-\epsilon_i-\zeta_i) + (\epsilon_i+\zeta_i)\eta^{(\mathrm{start},i)} + \zeta_i(1-\lambda_1)} \\
&\leq \frac{(\epsilon_i+\zeta_i)(\eta^{(\mathrm{start},i)}) + \zeta_i(1-\lambda_1)}{(1-\eta^{(\mathrm{start},i)})(1-\epsilon_i-\zeta_i)} = \frac{(\epsilon_i+\zeta_i)(\eta^{(i-1)}) + \zeta_i(1-\lambda_1)}{(1-\eta^{(i-1)})(1-\epsilon_i-\zeta_i)}.
\end{aligned}$$

Therefore,

$$\eta^{(i)} \leq \frac{\eta^{(i-1)} + \zeta_i(1-\lambda_1)}{4} \quad \text{as long as} \quad \eta^{(i-1)} \leq \frac{1}{3} \text{ and } \epsilon_i + \zeta_i \leq \frac{1}{8}. \tag{85}$$

We choose $\ell := \lceil \log_2(\frac{1-\gamma}{\varepsilon_{\mathrm{dist}}}) \rceil$, $\zeta_i := 2^{-3-i}$, and $\epsilon_i := 2^{-4-\ell+i}$. Since $\eta^{(0)} \leq 1-\gamma$, using Eq. (85) and induction we see that $\eta^{(i)} \leq 2^{-i}(1-\gamma)$ and therefore $\eta^{(\ell)} \leq \varepsilon_{\mathrm{dist}}$.

As $\sum_{i=1}^{\ell} \epsilon_i + \zeta_i \leq \frac{1}{4}$, we get by Eq. (83) that the sequence of iterations succeeds with probability $\prod_{i=1}^{\ell} p^{(\mathrm{succ},i)} \geq \frac{3}{4} p_1^{(0)} \geq \frac{1}{2}$. Therefore, the expected complexity can be bounded up to a constant factor by

$$C + \sum_{i=1}^{\ell} \frac{\log_2^2(1/\epsilon_i)}{\delta^2 \zeta_i} = C + \sum_{i=1}^{\ell} \frac{2^{i+3}(4+\ell-i)^2}{\delta^2} = C + \frac{2^{\ell+3}}{\delta^2} \sum_{i'=0}^{\ell-1} 2^{-i'}(i'+4)^2 = O\left(\frac{1}{(1-\alpha)^2\gamma^2}\left(\frac{1}{\gamma} + \frac{1-\gamma}{\varepsilon_{\mathrm{dist}}}\right)\right). \qquad \square$$

4.5 Resource state teleportation

Once we have prepared the resource state $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$ (up to low error), the next step is to consume this resource state to enact a transformation on the encoded address qubits. Suppose the n-qubit encoded register on which we

want to apply the operation $\overline{V(g)}$ is in an arbitrary pure state

$$|\overline{\alpha}\rangle = \sum_{w\in\{0,1\}^n} \alpha_w |\overline{w}\rangle\,, \qquad\qquad \sum_w |\alpha_w|^2 = 1\,. \tag{86}$$

The teleportation procedure calls for applying a set of n disjoint logical CNOT gates, each controlled by one of the logical qubits in the register holding $|\overline{\alpha}\rangle$ and targeting one of the logical qubits in the resource state, and then measuring the resource state. This was depicted in circuit (10), which is reproduced more compactly as

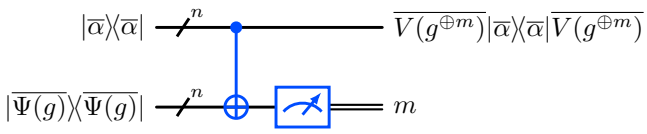

(87)

On the output label of the circuit, we have used the notation $g^{\oplus m}$ to represent the Boolean function defined for any x by the equation

$$g^{\oplus m}(x) = g(x \oplus m)\,. \tag{88}$$

To verify the circuit's output and the assertions in Section 2.3, note that the CNOT gates enact the transformation

$$|\overline{\alpha}\rangle|\overline{\Psi(g)}\rangle = \frac{1}{2^{n/2}} \sum_{w,z\in\{0,1\}^n} (-1)^{g(z)} \alpha_w |\overline{w}\rangle|\overline{z}\rangle \quad\longmapsto\quad \frac{1}{2^{n/2}} \sum_{w,z\in\{0,1\}^n} (-1)^{g(z)} \alpha_w |\overline{w}\rangle|\overline{w\oplus z}\rangle \tag{89}$$

$$= \frac{1}{2^{n/2}} \sum_{w,m\in\{0,1\}^n} (-1)^{g(w\oplus m)} \alpha_w |\overline{w}\rangle|\overline{m}\rangle\,. \tag{90}$$

A logical measurement is performed on the resource state register, obtaining outcome m. Regardless of the state $|\overline{\alpha}\rangle$ and the function g, the distribution over measurement outcomes m is the uniform distribution, and the post-measurement state on the first register is

$$\overline{V(g^{\oplus m})}|\overline{\alpha}\rangle = \frac{1}{2^{n/2}} \sum_{w\in\{0,1\}^n} (-1)^{g^{\oplus m}(w)} \alpha_w |\overline{w}\rangle = \frac{1}{2^{n/2}} \sum_{w\in\{0,1\}^n} (-1)^{g(w\oplus m)} \alpha_w |\overline{w}\rangle\,. \tag{91}$$

Examining the definition of $g^{\oplus m}$, we can see that it is defined by the same underlying table of data that defines g; however, the addresses to which the data items have been assigned are scrambled by adding (modulo 2) the random measurement outcome m to each address. Nevertheless, we may conclude that the teleportation procedure successfully inserts information about g into the phases of the state $|\overline{\alpha}\rangle$, albeit in an obfuscatedured way, due to the fact that the phase applied is $(-1)^{g(w\oplus m)}$ rather than the desired $(-1)^{g(w)}$, where m is uniformly random.

We now wish to understand what happens in the teleportation step if we use the state $\overline{\phi(g)}_{\text{dist}}$ (discussed in Section 4.4) as our resource state, which is an imperfect approximation to $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$. First, we define the ideal teleportation channel $\overline{\mathcal{T}(g)}$ as the procedure that perfectly prepares $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$, applies the logical CNOTs, and measures the resource register (equivalently, applies a completely dephasing channel). This has the action

$$\overline{\mathcal{T}(g)}[\rho] = \frac{1}{2^n} \sum_{m\in\{0,1\}^n} \overline{\mathcal{V}(g^{\oplus m})}[\rho] \otimes |m\rangle\langle m|\,, \tag{92}$$

where we have used the channel definition $\mathcal{V}(h)[\rho] = V(h)\,\rho\,V(h)^\dagger$ with the overlined version $\overline{\mathcal{V}(h)}$ denoting the logical version of the channel.

We may now argue that as long as the error of $\overline{\phi(g)}_{\text{dist}}$ is low, relative to $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$, the channel enacted is close to $\overline{\mathcal{T}(g)}$.

Proposition 11 (Teleportation with approximate resource state)**.** *Let $\overline{\mathcal{T}(g)}_{\text{appr}}$ be the channel realized by using the state $\overline{\phi(g)}_{\text{dist}}$ in place of $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$ in the teleportation protocol depicted in circuit (87). Then, we have*

$$\frac{1}{2}\|\overline{\mathcal{T}(g)} - \overline{\mathcal{T}(g)}_{\text{appr}}\|_\diamond \leq \frac{1}{2}\||\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\text{dist}}\|_1\,, \tag{93}$$

where $\|\cdot\|_\diamond$ indicates the diamond norm distance between channels.

Proof. Define the quantum operation $\overline{\mathcal{C}}$ as the operation on $2n$ logical qubits enacted by the CNOT gates and measurements (i.e., completely dephasing channel) in circuit (87) Thus, for arbitrary n-qubit state $\overline{\sigma}$ in the codespace of the first register, we have

$$\overline{\mathcal{T}(g)}[\overline{\sigma}] = \overline{\mathcal{C}}\left[\overline{\sigma} \otimes |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\right] \tag{94}$$

$$\overline{\mathcal{T}(g)}_{\rm appr}[\overline{\sigma}] = \overline{\mathcal{C}}\left[\overline{\sigma} \otimes \overline{\phi(g)}_{\rm dist}\right]. \tag{95}$$

Now, introduce an environment register of a qubits, and consider an arbitrary state $\overline{\rho}$ on $a+n$ logical qubits (whose reduced density matrix after tracing out the environment is in the codespace of the n-logical-qubit code). Let $\mathcal{I}_{\overline{E}}$ be the identity channel on the environment register. The diamond norm distance is defined as the supremum (over $\overline{\rho}$ for all possible sizes a of the environment) in the trace distance between the action of the two channels

$$\frac{1}{2}\|\overline{\mathcal{T}(g)} - \overline{\mathcal{T}(g)}_{\rm appr}\|_\diamond = \max_a \sup_{\overline{\rho}} \frac{1}{2}\|(\mathcal{I}_{\overline{E}} \otimes \overline{\mathcal{T}(g)})[\overline{\rho}] - (\mathcal{I}_{\overline{E}} \otimes \overline{\mathcal{T}(g)}_{\rm appr})[\overline{\rho}]\|_1 \tag{96}$$

$$= \max_a \sup_{\overline{\rho}} \frac{1}{2}\left\|(\mathcal{I}_{\overline{E}} \otimes \overline{\mathcal{C}})\left[\overline{\rho} \otimes \left(|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\rm dist}\right)\right]\right\|_1 \tag{97}$$

$$\leq \max_a \sup_{\overline{\rho}} \frac{1}{2}\|\overline{\rho} \otimes \left(|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\rm dist}\right)\|_1 \tag{98}$$

$$= \max_a \sup_{\overline{\rho}} \frac{1}{2}\||\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\rm dist}\|_1 = \frac{1}{2}\||\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\rm dist}\|_1, \tag{99}$$

where the inequality follows from monotonicity of the trace distance under quantum channels. This completes the proof. □

4.6 Adaptive correction and classical update rule

Recall that the goal is to implement the diagonal logical QRAM unitary $\overline{V(f)}$, given a data table (Boolean function) f. The previous section established that for any function g, the teleportation channel $\overline{\mathcal{T}(g)}$ receives a uniformly random measurement outcome m, and then, conditioned on m, enacts the unitary transformation $|\overline{\alpha}\rangle \mapsto \overline{V(g^{\oplus m})}|\overline{\alpha}\rangle$.

Our protocol repeats the teleportation process over a sequence of n rounds. In each round, the function g will be different, and the measurement outcome m will be sampled independently and uniformly at random. The value of g used in round j will be denoted $g^{(j)}$, and the value of m denoted by $m^{(j)}$.

We begin in round 1 setting $g = g^{(1)} = f$, receiving measurement outcome $m = m^{(1)} \in \{0,1\}^n$, and enacting logical unitary $\overline{V(g^{\oplus m})}$. The key observation is that $\overline{V(g^{\oplus m})}^2 = \overline{\mathbb{I}}$, the logical identity operation, and hence

$$\overline{V(g)} = \overline{V(g)}\ \overline{V(g^{\oplus m})}\ \overline{V(g^{\oplus m})} = \overline{V(g \oplus g^{\oplus m})}\ \overline{V(g^{\oplus m})}, \tag{100}$$

where we have invoked the composition rule $V(h)V(h') = V(h \oplus h')$. Thus, given that we have already implemented $\overline{V(g^{\oplus m})}$, we must now implement the *correction unitary* $\overline{V(g \oplus g^{\oplus m})}$, which is also a member of the family of diagonal (logical) QRAM operators. We thus compute a new Boolean function by the *classical update rule* UR, which takes as input a data table g and a bit string $m \in \{0,1\}^n$ and outputs a new data table

$$\mathrm{UR}(g, m) = g \oplus g^{\oplus m}, \tag{101}$$

which is depicted as a circuit as

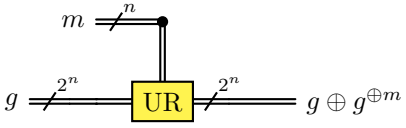

(102)

and we update $g \leftarrow g^{(2)} = \mathrm{UR}(g^{(1)}, m^{(1)})$ to use for round 2. Consequently, the correction unitary is now equal to $\overline{V(g)}$ and the goal of round 2 is again simply to implement the unitary $\overline{V(g)}$, just as it was in round 1. As before, we receive a new measurement outcome $m = m^{(2)}$, we compute $g^{(3)} = \mathrm{UR}(g^{(2)}, m^{(2)})$, and we update $g \leftarrow g^{(3)}$ for round 3.

We then iterate this process a number of times. We note that, if after applying the update rule we ever obtain $g = \mathbf{0}$, the zero function, then we may terminate the procedure, because the correction unitary will be $\overline{V(\mathbf{0})} = \bar{\mathbb{I}}$ at the next round. Moreover, $g = \mathbf{1}$ (the constant function that outputs 1 on all inputs), then the correction unitary is $-\bar{\mathbb{I}}$, which is equivalent to $\bar{\mathbb{I}}$ up to an unphysical global sign. We claim that after at most n rounds, we will be certain to obtain $g \in \{\mathbf{0}, \mathbf{1}\}$, based on the following proposition.

Proposition 12. *Let g be an n-bit Boolean function. Define $\deg(g)$ to be the degree of g when it is expanded as a polynomial of its input bits over the field $\mathbb{F}_2$. Suppose that $\deg(g) = d$. Let $m \in \{0,1\}^n$, and let $h = \mathrm{UR}(g, m)$, as defined in Eq. (101). Then, we have*

$$\deg(h) \leq d - 1\,. \tag{103}$$

Proof. This is a consequence of the reasoning in Appendix D, specifically the cancellation in Eq. (233) and the reasoning underneath. □

The proposition establishes that each application of the update rule $g \leftarrow \mathrm{UR}(g, m)$ decreases the degree of g by at least 1. Recall that in round 1, we have $\deg(g) = \deg(f) \leq n$, simply by virtue of the fact that f is an n-bit function. Thus, we may assert that in round j we have $g = g^{(j)}$ and

$$\deg(g^{(j)}) \leq n + 1 - j\,. \tag{104}$$

In particular, after applying the update rule in round n with $g = g^{(n)}$ and $m = m^{(n)}$, we are guaranteed to obtain $h = \mathrm{UR}(g, m)$, which leads to a degree $\deg(h) = 0$. If the degree of h is 0, this implies that either $h = \mathbf{0}$ or $h = \mathbf{1}$.

4.7 Total complexity of protocol

We have now explained the action of each component of the protocol, and may we state its overall complexity.

Theorem 2 (Main result)**.** *Let f be an arbitrary dataset (n-bit Boolean function) for which we wish to implement the fault-tolerant diagonal QRAM unitary $\overline{V(f)}$ of Eq. (1), and let ε be an error parameter. Suppose that we have access to a noisy physical QRAM device subject to dataset-independent noise (Definition 1) which on input g produces state $\widetilde{\psi}(g)$ on n physical qubits achieving fidelity $\langle \Psi(g)|\widetilde{\psi}(g)|\Psi(g)\rangle \geq F$ for all g. Suppose further that we have the capability to reload the QRAM device with a new dataset, and that we have the capability to move the n-qubit output state to a fault-tolerant quantum processor subject to circuit-level stochastic noise (Definition 2) with error rate p. If p is below a constant threshold p_0 determined by the QEC code family, and separately if pn^2 is below a different constant threshold related to the fault-tolerant encoding procedure, then there exists an adaptive distillation–teleportation procedure that implements a quantum channel $\mathcal{P}_{\mathrm{DT}}$ for which*

$$\frac{1}{2}\|\mathcal{P}_{\mathrm{DT}} - \overline{V(f)}[\cdot]\overline{V(f)}^\dagger\|_\diamond \leq \varepsilon\,. \tag{105}$$

The protocol uses (in expectation over internal randomness) Q queries to the noisy physical QRAM device, Q applications of the encoding process $\mathcal{E}_{\mathrm{FT}}$ (Corollary 1), and Q' additional fault-tolerant operations (controlled-SWAP, Hadamard, CNOT, single-qubit logical state preparations, and single-qubit logical measurements), where

$$Q = O\left(\frac{n(1-F)}{F^2}\left(\frac{n}{\varepsilon} + \frac{1}{F}\right)\right) \tag{106}$$

$$Q' = O\big(n^2 Q\big)\,. \tag{107}$$

Additionally, the protocol applies the classical update rule (Eq. (101)) at most n times, and the classical partial Clifford twirling operation $g \mapsto g_C$ (Eq. (35)) Q times, each time on a data table of size 2^n classical bits.

Proof. The protocol is depicted as a quantum circuit in Fig. 2. We begin with correctness. In each of the n rounds indexed by $j = 1, \ldots, n$, it implements a channel $\overline{\mathcal{T}(g^{(j)})}$. As discussed in Section 4.6, if all resource states are prepared perfectly, the procedure is guaranteed to implement the unitary $\overline{V(f)}$, up to a global sign which does not impact the channel $\overline{V(f)}[\cdot]\overline{V(f)}^\dagger$.

However, the teleportation channel $\overline{\mathcal{T}(g^{(j)})}$ is not implemented perfectly by the protocol. To ensure the overall diamond norm error is ε, it suffices to choose parameters such that $\overline{\mathcal{T}(g)}$ is implemented up to ε/n diamond distance

for all g, since the errors from each of the n channel applications add linearly in the worst case, when performed in succession. By Proposition 11, it suffices to distill resource states $\overline{\phi(g)}_{\text{dist}}$ that satisfy

$$\frac{1}{2}\| |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| - \overline{\phi(g)}_{\text{dist}}\|_1 \leq \varepsilon_{\text{dist}} \tag{108}$$

with distillation error $\varepsilon_{\text{dist}} = \varepsilon/n$. Meanwhile, by Proposition 6, this is accomplished by the distillation protocol using $O(\frac{1-F_{\min}}{F_{\min}^2}(\frac{1}{\varepsilon_{\text{dist}}} + \frac{1}{F_{\min}}))$ copies of the input states $\overline{\phi(g)}_{\text{twirl}}$, defined in Eq. (43), provided that for all g $\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq F_{\min}$ for some $F_{\min}$. We are guaranteed from Corollary 1 that $F_{\min} \geq (1 - O(np) - O(n\sqrt{p}))F$, which can be replaced by $\Omega(F)$ as long as pn^2 is below a certain constant. The number of gates required by distillation is a factor of $O(n)$ larger than the number of copies.

Since there are n rounds, we require n calls to the distillation procedure. Thus, the total number of queries to the noisy QRAM device and the poly(n)-cost encoding procedure $\mathcal{E}_{\text{FT}}$ is

$$Q = O\left(\frac{n(1-F)}{F^2}\left(\frac{n}{\varepsilon} + \frac{1}{F}\right)\right) \tag{109}$$

The total fault-tolerant gate complexity from distillation is $O(nQ)$. The teleportation procedure also requires n fault-tolerant CNOT gates in each of the n rounds. Finally, the Clifford twirling step requires the application of $O(n^2)$ fault-tolerant Clifford gates for each of the Q copies. In total, these Clifford gates dominate the fault-tolerant gate count, which is $Q' = O(n^2 Q)$. Each round requires only one call to the update rule, and each of the Q copies requires classically applying a partial Clifford update $g \mapsto g_C$. This completes the proof. □

5 Complexity of the classical update rule

The classical update rule calls for updating a data table g to the data table $h = \text{UR}(g, m) = g \oplus g^{\oplus m}$, for a certain fixed measurement outcome $m \in \{0,1\}^n$, as in Eq. (101). That is, the entry at address x in the data table should be updated from $g(x)$ to $g(x) \oplus g(x \oplus m)$. In this section, we analyze the complexity of this transformation under several different frameworks. The guiding question is to understand the ways in which the classical update rule is a more complex operation than a RAM query.

We note that the partial Clifford twirling step also requires substantial classical computation to randomly transform the dataset. We do not specifically analyze this step here, because we view it as less fundamental to our protocol. For example, if the physical QRAM device and encoding step were error free (or if they have errors but the principal eigenvalue of the resulting state is correct), then partial Clifford twirling is not necessary, but the need for the update rule remains.

5.1 Classical circuit complexity

The update rule takes $2^n + n$ bits as input and produces 2^n bits as output, as in circuit (102). It is straightforward to see that a classical circuit built from elementary gates (NOT, AND, NAND, etc.) would require $\Omega(2^n)$ gates to implement the update rule. Observe that for any fixed nonzero value $m \neq 0^n$, we have for every x

$$h(x) = h(x \oplus m) = g(x) \oplus g(x \oplus m) \tag{110}$$

That is, the 2^n addresses are partitioned into pairs $\{x, x \oplus m\}$, where h takes the same value on both elements of the pair, and its value is equal to the parity of the input bits at locations x and $x \oplus m$. Each of these 2^{n-1} parity calculations is independent and requires at least one elementary gate.

5.2 Complexity in a classical RAM model

The need for $\Omega(2^n)$ circuit complexity does not alone entail that the update rule is an expensive classical calculation. After all, the standard RAM operation also has $\Omega(2^n)$ circuit complexity, but it is commonplace to consider a model of classical computation where RAM has unit cost.

However, the single-bit RAM operation takes as input n bits (an address) and returns just 1 bit, so it appears to be a much simpler operation than the classical update rule. Since (i) the output of the update rule has 2^n bits, (ii) the output depends on all 2^n input bits $g(x)$, and (iii) each RAM query can access only 1 of the bits, we

conclude that, in a model where the input data g can only be accessed via RAM queries, implementing the update rule requires $\Omega(2^n)$ RAM queries.

5.3 Classical circuit depth in an all-to-all model

The structured nature of the update rule suggests it may still be amenable to some degree of parallelization. Ideally, one could design a specialized shallow circuit that directly implements the update, rather than relying on RAM. Indeed, if one imposes no restrictions on spatial layout of the $2^n + n$ input bits and 2^n output bits (i.e., one allows all-to-all gates), then one can perform the update rule with a classical circuit of depth $O(n)$ comprised of elementary gates each acting on only $O(1)$ bits.

We provide one possible way to accomplish this. The construction requires $O(n2^n)$ ancilla bits and has the following steps; we provide an $n = 3$ example to assist with understanding each step in Fig. 5 (and in the description of each step, we reference the colors in that figure for clarity of explanation). Let e_i refer to the length-n bit string with a 1 in the i-th position and 0 in the other $n-1$ positions.

1. For each $x \in \{0,1\}^n$ in parallel, the (blue) bit holding $g(x)$ is copied into a (yellow) ancilla bit, accomplished in depth 1.
2. For each $i = 1, \ldots, n$ in parallel, the (green) input bit holding m_i is copied into $2^{n-1} - 1$ (gray) ancilla bits, accomplished in depth $n-1$ (using a tree-like approach, the number of copies of m_i can be doubled with each additional circuit layer).
3. For each $i = 1, \ldots, n$ in series, and for each of the 2^{n-1} address pairs $\{x, x \oplus e_i\}$ in parallel, if $m_i = 1$ then the (yellow) ancilla bit which held the copy of $g(x)$ at the beginning of this step is swapped with the (yellow) ancilla bit which held the copy of $g(x \oplus e_i)$ at the beginning of this step. Each value of i requires depth 1: the availability of 2^{n-1} (green and gray) copies of the m_i bit created in step 2 allows parallelization of the m_i-controlled $\{x, x \oplus e_i\}$ swaps. Thus, the overall depth of this step is n. Over the course of the n steps, the (yellow) ancilla bit that originally held $g(x)$ before this step undergoes the transformations

$$g(x) \overset{i=1}{\mapsto} g(x \oplus m_1 e_1) \overset{i=2}{\mapsto} g(x \oplus m_1 e_1 \oplus m_2 e_2) \overset{i=3}{\mapsto} \cdots \overset{i=n}{\mapsto} g(x \oplus \bigoplus_{i=1}^{n} m_i e_i) = g(x \oplus m)\,, \tag{111}$$

   that is, for each x, the (yellow) ancilla bit originally holding $g(x)$ now holds $g(x \oplus m)$.
4. For each $x \in \{0,1\}^n$ in parallel, the (yellow) ancilla bit that originally held the copy of $g(x)$ after step 1 (which now holds $g(x \oplus m)$) is added modulo 2 into the (blue) bit holding $g(x)$, incurring depth 1.
5. The original 2^n (blue) input bits are taken to be the 2^n output bits; the ancilla bits are discarded.

The registers (blue) holding the original bits $g(x)$ have now been updated to $h(x) = g(x) \oplus g(x \oplus m)$, which is the desired output of the update rule.

5.4 Classical circuit depth in a spatially local model

If the input and output bits of the shallow circuit are embedded into d spatial dimensions, then the $O(n)$-depth circuit above requires spatially nonlocal gates. For example, we may recognize the action of step 3 as a re-arrangement of the 2^n (yellow) ancilla bits by traversing edges of the n-dimensional Boolean hypercube. In $d = n$ spatial dimensions, this could be done using spatially local gates, and each (yellow) classical bit need only interact with $O(n)$ of the 2^n other (yellow) bits (its neighbors on the hypercube). Unfortunately, real classical circuits must be embedded into $d \leq 3$ spatial dimensions. In this case, step 3 cannot be performed exclusively with spatially local gates for more than d of the values of $i \in \{1, \ldots, n\}$.

In fact, we can show that if gates are local in d spatial dimensions, then the depth required is at least $\Omega(2^{n/d})$. This follows from the fact that without knowing the value of m, the circuit must be prepared to connect the bit at address x to all $2^n - 1$ of the other bits; for any pair $x, y \in \{0,1\}^n$, if $m = x \oplus y$, then the circuit must be able to compute the parity of $g(x)$ and $g(y)$. If the bits storing $g(x)$ and $g(y)$ live on opposite sides of the d-dimensional array, computing this parity will require a classical circuit with depth at least $\Omega(2^{n/d})$. Ultimately, this essentially amounts to a speed of light–type restriction, where depth is restricted due to the fact that information can only move so quickly through space. This kind of argument could also be used to show that classical RAM requires a circuit of depth $\Omega(2^{n/d})$ in a local model, in d spatial dimensions, so it does not represent a fundamental limitation that is unique to QRAM.

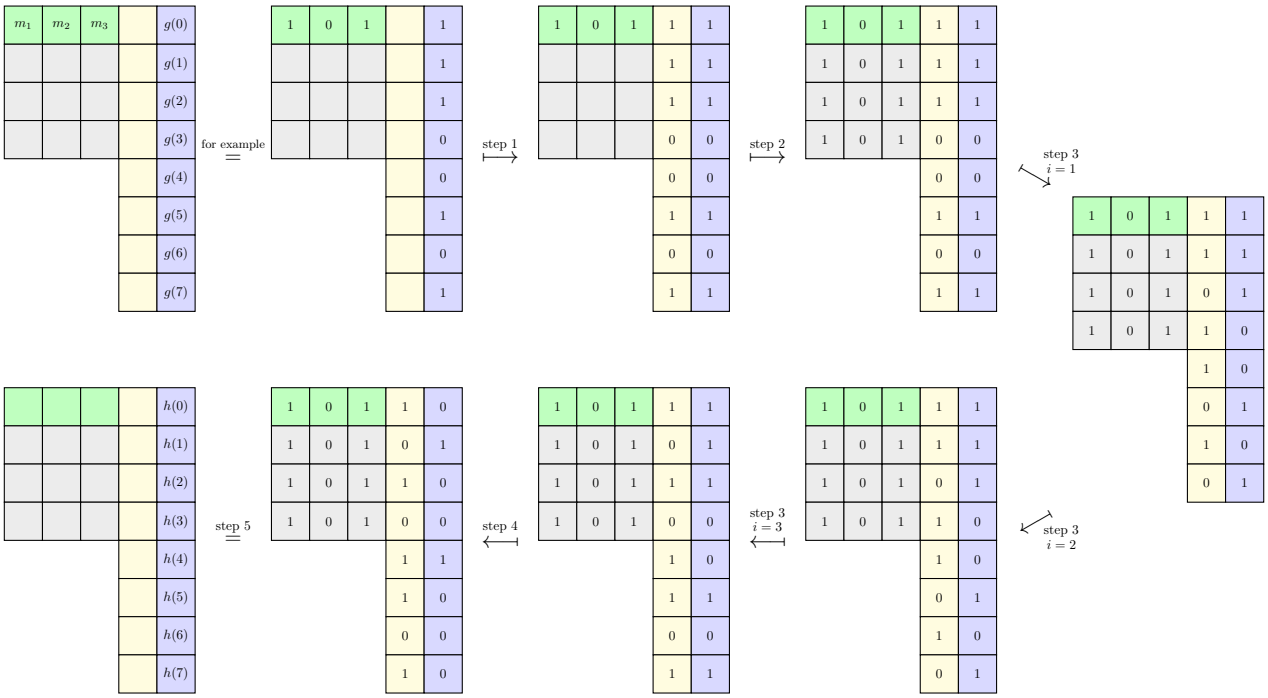

Figure 5: Step-by-step action of the $O(n)$-depth classical circuit that implements the update rule $g \mapsto h = g \oplus g^{\oplus m}$, for a particular $n = 3$ example input with $m = (1, 0, 1)$. Each box stores one bit of information, and each step modifies some subset of these bits with parallelized layers of local gates (note that step 2 requires $n - 1$ layers to create all $2^{n-1} - 1$ copies of each m_i). The inputs to the update rule are the n bits of m (green) and the 2^n bits in the classical data table storing g (blue). The circuit utilizes $2^n + n(2^{n-1} - 1)$ ancilla bits (gray and yellow).

5.5 Wire density

While speed of light–type restrictions can be relevant for RAM at large scales, they are not a factor in practice at small or intermediate scales. If one ignores the speed of light, one can simply build long wires into the circuit and enable all-to-all connectivity. These long wires should not be thought of as completely free, however. For example, electronic ciruits typically dissipate energy and lead to heating in proportion to their total wire length. The need to cool electronic chips is a key limitation in practical computer systems. Thus, a cost metric worth considering [8] is the wire density of the circuit, which we define as the total wire length divided by the total spacetime of the circuit, where the spacetime is defined as the number of bits the circuit uses (henceforth referred to as the circuit width) multiplied by the circuit depth.

Classical circuits for RAM can have depth $O(n)$, width $O(2^n)$, and total wire length $O(n2^n)$ [8], even when embedded in one spatial dimension. Thus, the wire density is a constant with respect to n, suggesting the circuit can be scaled without causing heating issues.

We now examine the circuit for the update rule described in Section 5.3. It has depth $O(n)$ and width $O(n2^n)$. Steps 1 and 2 perform copying of the bits and contribute wire length $O(n2^n)$. If the circuit is embedded in n spatial dimensions, then step 3 can also be accomplished with total wire length $O(n2^n)$, as each gate is local and has $O(1)$ wire length. However, in $d \leq 3$ spatial dimensions, the wire length is asymptotically larger. To implement step 3, the (yellow) ancilla bits storing the copy of $g(x)$ must be connected to the (yellow) ancilla bits initially storing the copies of $g(x \oplus e_1)$, $g(x \oplus e_2)$, $\ldots$, $g(x \oplus e_n)$—essentially, an embedding of the n-dimensional Boolean hypercube into d dimensions. Consequently, the total wire length of gates acting on each (yellow) ancilla bit will be at least $\Omega(2^{n/d})$. Since there are 2^n (yellow) ancilla bits, the overall wire length of the circuit is at least $\Omega(2^{n(1+1/d)})$ and thus the wire density grows with n as $\Omega(2^{n/d}/\mathrm{poly}(n))$, a fundamentally different outcome than the case of classical RAM.

5.6 Relation to matrix-vector multiplication and the Walsh–Hadamard transform

matrix-vector multiplication for $2^n \times 2^n$ matrices is an operation with 2^n inputs (the entries of the input vector) and 2^n outputs (the entries of the output vector). This feature is similar to the update rule, although the inputs and

outputs for matrix multipliation would typically each be multiple bits (e.g., an integer or floating point number), rather than just a single bit. Moreover, the analysis of Ref. [8] demonstrated how classical sparse matrix-vector multiplication requires a growing wire density, consistent with the observation above for the update rule.

Here, we will argue that the update rule is in a certain sense equivalent to a sparse matrix-vector multiplication, and specifically it is equivalent to the Walsh–Hadamard (WH) transform up to factors of poly(n). This equivalence holds under a parallel model of computation. Namely, we assume that we have 2^n classical co-processors, each with poly(n)-size local memory, which may perform local arithmetic on their memory in parallel, but cannot communicate with one another, except through joint application of the update rule or through joint application of a sparse matrix-vector multiplication. When jointly applying the update rule, each of the 2^n processors supplies one entry $g(x)$ and receives the output $h(x) = g(x) \oplus g(x \oplus m)$ for a fixed global m. When jointly applying sparse matrix-vector multiplication, each processor provides one of the 2^n entries of the input vector and receives one of the 2^n entries of the output vector. The remainder of this subsection aims to justify this claim of equivalence.

5.6.1 The (fast) Walsh–Hadamard transform

The WH transform is the multiplication of a length-2^n vector by a $2^n \times 2^n$ matrix denoted by H, where the matrix element associated with the transition from n-bit input address y to n-bit output address x is given by

$$H_{xy} = \frac{1}{2^{n/2}}(-1)^{x \cdot y} \tag{112}$$

We can see that H_{xy} is a dense matrix—all of its entries are nonzero—however, it can be shown that it is the product of n sparse matrices. Specifically, let $H^{(i)}$ be defined as[9]

$$H^{(i)}_{xy} = \begin{cases} 1 & \text{if } x = y \oplus e_i \\ (-1)^{x_i y_i} & \text{if } x = y \\ 0 & \text{otherwise} \end{cases} \tag{113}$$

Then, it holds that

$$H = \frac{1}{2^{n/2}} H^{(n)} H^{(n-1)} \cdots H^{(2)} H^{(1)}\,, \tag{114}$$

This fact can be verified by noting that the nonzero offdiagonal entries of $H^{(i)}$ are matrix elements $H^{(i)}_{uv}$ where u and v differ only on the i-th bit. Thus, any offdiagonal transition element H_{xy} can only be obtained in the product $H^{(n)} \cdots H^{(1)}$ by choosing the corresponding offdiagonal entry of $H^{(i)}$, (equal to 1) whenever $x_i \neq y_i$ differ, and the diagonal entry $(-1)^{x_i y_i}$ of $H^{(i)}$ whenever $x_i = y_i$. This gives precisely the quantity $(-1)^{\sum_i x_i y_i} = (-1)^{x \cdot y}$.

Matrix multiplication by $H^{(i)}$ requires only $O(2^n)$ arithmetic operations, since each row of $H^{(i)}$ has only 2 nonzero entries. The fast WH transform utilizes this decomposition to implement multiplication by H in classical time poly$(n)2^n$, much smaller than the $\Omega(2^{2n})$ time required to multiply general dense matrices that do not have this kind of decomposition.

5.6.2 Reduction from update rule to Walsh–Hadamard transform

Now, we show how the update rule can be accomplished with two applications of the WH transform and parallel local arithmetic. Let $\mathbf{g}$ denote the length-2^n vector with entry $g(x)$ at index x, and let $\mathbf{h}$ denote the length-2^n vector with entry $g(x) + g(x \oplus m)$ at entry x. Note here that we are using normal addition $+$, rather than modular addition $\oplus$, and we allow $\mathbf{h}$ to take integer values in $\{0, 1, 2\}$. We have that the x-th entry of the matrix-vector

[9] We note that the matrix $H^{(i)}$ is proportional (by a factor $\sqrt{2}$) to the transformation applied to the amplitudes of a quantum state when a single-qubit Hadamard gate is applied to the i-th qubit, and identity is applied to the other $n-1$ qubits. The full WH transform is the product of the $H^{(i)}$ matrices, that is, the Hadamard gate on all n qubits.

product $H\mathbf{h}$ is

$$(H\mathbf{h})_x = \frac{1}{2^{n/2}} \sum_{y\in\{0,1\}^n} (-1)^{x\cdot y}(g(y) + g(y\oplus m)) \tag{115}$$

$$= \left[\frac{1}{2^{n/2}} \sum_{y\in\{0,1\}^n} (-1)^{x\cdot y} g(y)\right] + (-1)^{m\cdot x}\left[\frac{1}{2^{n/2}} \sum_{y\in\{0,1\}^n} (-1)^{x\cdot(y\oplus m)} g(y\oplus m)\right] \tag{116}$$

$$= (H\mathbf{g})_x + (-1)^{m\cdot x}(H\mathbf{g})_x \tag{117}$$

$$= \begin{cases} 2(H\mathbf{g})_x & \text{if } m\cdot x = 0 \\ 0 & \text{if } m\cdot x = 1 \end{cases} \tag{118}$$

That is, the WH transform $H\mathbf{h}$ of the update rule output $\mathbf{h}$ can be easily computed from the WH transform $H\mathbf{g}$ of the update rule input $\mathbf{g}$.

Let $M^{(m)}$ be the diagonal matrix for which $M^{(m)}_{xx} = 2$ if $m\cdot x = 0$ and $M^{(m)}_{xx} = 0$ if $m\cdot x = 1$. Then, we may use the equation above to say that $H\mathbf{h} = MH\mathbf{g}$. Since $H^2 = \mathbb{I}$, we then have

$$\mathbf{h} = HH\mathbf{h} = HM^{(m)}H\mathbf{g} \tag{119}$$

This gives a straightforward way to transform $\mathbf{g} \mapsto \mathbf{h}$ in the model where 2^n classical co-processors can perform local arithmetic in parallel or jointly perform the WH transform. First, we assume that 2^n processors each locally store the bit $g(x)$ for one address x, and an n-bit copy of m (note that a copy of m could be pre-distributed to each of the 2^n processors via a tree-like circuit of constant wire density, as in step 2 of Section 5.3). Then, $\mathbf{g} \mapsto \mathbf{h}$ is accomplished by applying the WH transform, applying the diagonal transformation $M^{(m)}$, and finally applying the WH transform again. The x-th entry of the diagonal matrix $M^{(m)}$ can be computed in $O(1)$ depth locally by the x-th processor acting on its internal memory. We recall that $\mathbf{h}$ may have entries in $\{0,1,2\}$; to recover binary entries, we simply take every entry modulo 2 via parallel local arithmetic, which does not discard any important information, owing to the fact that $(-1)^{g(x)+g(x\oplus m)} = (-1)^{g(x)\oplus g(x\oplus m)}$.

The conclusion is that $O(1)$ applications of the WH transform, along with parallel local arithmetic, is sufficient to implement the classical update rule, or in other words, the classical update rule is no harder than the WH transform.

5.6.3 Reduction from WH transform to update rule

Now, we show the converse: that the WH transform can be implemented using poly(n) applications of the update rule along with parallel local arithmetic. Specifically, we aim to implement the matrix transformation $H^{(i)}$. Suppose we are given a vector $\mathbf{g}$ of integers each represented by at most n bits.[10] Let $g(x)$ denote the integer corresponding to the entry of $\mathbf{g}$ at address x, which is stored in the local memory of one of the 2^n processors. Let $g_j(x)$ denote the j-th bit of the integer $g(x)$. We perform the following steps

1. For each address x, the co-processor storing $g(x)$ makes a copy of $g(x)$ in its local memory.
2. We fix global n-bit string $m = e_i$ (known by all processors) and for each j, the 2^n co-procesoors jointly apply the update rule $\mathrm{UR}(g_j, e_i)$ onto the set of 2^n bits $g_j(x)$ (i.e., bit-wise application of update rule). For each $x \in \{0,1\}^n$, and each j, the co-processor at address x now has in its memory the value $g_j(x)$ and the value $g_j(x) \oplus g_j(x\oplus e_i)$.
3. For each x and each j, the co-processor at address x performs local arithmetic to add (modulo 2) the local register holding $g_j(x)$ into the local register holding $g_j(x)\oplus g_j(x\oplus e_i)$ so that the two registers now hold $g_j(x)$ and $g_j(x\oplus e_i)$. Effectively, this step and the previous step have together performed a swap $g(x) \leftrightarrow g(x\oplus e_i)$ between the various co-processors.
4. For each address x, the co-processor at address x takes the local register holding integer $g(x)$ and flips it to $-g(x)$ only if $x_i = 1$.
5. Finally, for each address x, the co-processor at address x adds the local register holding $(-1)^{x_i}g(x)$ into the local register holding $g(x\oplus e_i)$, such that the latter register now holds $g(x\oplus e_i) + (-1)^{x_i}g(x)$.

[10] The assumption of integer entries is without loss of generality. If the entries are not integers but rather non-integer numbers expressed in binary with a finite number of bits of precision, then we can always multiply by a power of 2 so that all the entries are integers, and divide by that power of 2 after performing the calculation.

The latter register in the local memory of co-processor x now holds precisely the value $(H^{(i)}\mathbf{g})_x$, indicating that the co-processors have successfully managed to apply one step of the fast WH transform. Accomplishing this required one application of the update rule for each bit of the entries of the vector. Assuming the integer entries have at most poly(n) bits, this means $H^{(i)}$ can be accomplished with poly(n) applications of the update rule, and local arithmetic, specifically, copying (step 1), bit-wise addition mod 2 (step 3), negation (step 4), and integer addition (step 5).

Since H is the product of n matrices $H^{(i)}$ up to a proportionality constant, we conclude that poly(n) applications of the update rule are sufficient to implement the WH transform in this model of computation; in other words, the WH transform is no harder than the update rule, up to a factor of poly(n).

5.6.4 Performing general sparse matrix-vector multiplication using the update rule

Multiplication by $H^{(i)}$ is in some sense easier than multiplication by an arbitrary sparse matrix, since (i) the location of the nonzero entries is highly structured and (ii) all entries are ± 1, avoiding the need for any integer or floating-point multiplications. The only arithmetic required is addition and subtraction. Here we discuss how the update rule can also be used for arbitrary sparse matrix-vector multiplications.

Lemma A.3 of Ref. [8] examines performing general sparse matrix-vector multiplication using parallel classical co-processors, each capable of local addition and multiplication. If the parallel processors can communicate with each other via links that form a sorting network, then the sparse matrix-vector multiplication can be accomplished in roughly the time required to perform a sort using the sorting network. As shown in steps 2 and 3 of the procedure in Section 5.6.3, the update rule and parallel bitwise xor together allow a swap of data at location x and location $x \oplus e_i$ for all x in parallel. By copying the data before performing the swap, and then choosing whether to discard the original copy or the swapped copy, one can use the update rule to swap in parallel any *subset* of the location pairs $(x, x \oplus e_i)$. Thus, the ability to perform the update rule enables access to parallel swaps along a Boolean hypercube connectivity in n dimensions, a model for which it is known that sorting can be completed in poly(n) time [107, 108]. Together with Ref. [8, Lemma A.3] this shows how poly(n) applications of the update rule and poly(n) rounds of parallel local arithmetic on the database entries enables arbitrary sparse matrix-vector multiplication.

Unlike for the WH transform, this procedure for arbitrary sparse matrix-vector multiplication also requires the local arithmetic to include multiplication, rather than just addition. However, poly(n)-bit integer multiplication can be accomplished with only poly(n) integer additions. In any case, this suggests that the update rule is essentially equivalent to a general sparse matrix-vector multiplication, at least in this model where parallel local arithmetic is possible. We note that it could still be possible that in practice that the constant and poly(n) prefactors for the update rule are substantially better than those for general sparse matrix-vector multiplication.

5.7 Concluding comments on parallelization of the update rule

The above arguments establish that the update rule can be parallelized to $O(n)$ depth, but only in a model where all-to-all gates are possible. Embedding this all-to-all circuit into a finite number of spatial dimensions causes the circuit to have wire density that grows exponentially with n, a feature that suggests the parallelized update rule is less scalable than classical RAM. It is unclear whether this growing wire density is problematic for practical sizes of n.

Relatedly, we have shown that in a parallel model of computation, the ability to apply the update rule is roughly equivalent to the ability to apply a sparse matrix-vector multiplication, and there is a particularly close connection to the WH transform. In many instances, sparse matrix-vector multiplication can be parallelized very successfully in practice, for example, by leveraging graphics processing units (GPUs), where the parallelization is hardwired into the chip. This comes despite the fact that, asymptotically speaking, the wire density of a sparse matrix-vector multiplication would increase exponentially with n [8]. Since the classical update rule seems to be an operation that is no harder than a sparse matrix-vector multiplication, we are hopeful that it would be possible to effectively parallelize the classical update rule in practice.

However, this argument also clarifies the opportunity cost of the classical resources dedicated to performing the update rule. For example, if the quantum algorithm requiring fault-tolerant QRAM aims to solve a certain linear algebra problem, one must consider whether the classical device that performs the classical update rule is capable of solving that problem on its own, without the need for a quantum computer at all. After all, many linear algebra problems, such as solving sparse linear systems, can be solved using a small number of sparse matrix-vector multiplications; see Section 6.2 and Ref. [8].

6 Applications

In this section, we make a coarse attempt at estimating the resources required by our protocol in several applications. The goal is to shed light on which aspects of our protocol are most in need of improvement for it to be useful.

Generally, these applications come in two flavors—first, there are big-data appliations that heavily rely on QRAM and require the cheap QRAM assumption from Section 1 to have a significant quantum speedup. For these applications, conceptually speaking, the limiting aspect of using our protocol is the exponential *classical* computation required to do the update rule (and the Clifford twirling), as this prevents a true exponential speedup and full justification of the cheap QRAM assumption. One must always consider that the classical resources required to perform the update rule could be re-purposed directly toward solving the computational problem; in Section 5, we discussed how the classical update rule is in a sense equivalent in complexity to a sparse matrix-vector multipliation.

The second flavor of application are those where the QRAM operation of Eq. (1) (or its b-bit generalization, discussed in Appendix A) is required, but the cheap QRAM assumption is not essential—cryptanalysis and chemistry, below. In fact, in these applications, the fault-tolerant QRAM operation is typically compiled as a space-efficient quantum circuit, requiring only $O(n)$ logical qubits and $\Omega(2^n)$ depth—in this context, the operation is often referred to as QROM (quantum read-only memory) [74], rather than QRAM. Assigning QRO(A)M cost $\Omega(2^n)$ does not necessarily jeopardize the possibility of quantum advantage. This is an appealing place to apply our protocol, because it may be viewed as offloading exponential resources from the quantum processor to the classical processor, which is typically a favorable trade. The issue we face here is that the $O(1/\varepsilon)$ overhead for distillation in our method compares unfavorably to the polylog$(1/\varepsilon)$ overhead achieved for distillation of T and CCZ magic states, and ε may need to be taken quite small if QRAM is applied many times. Furthermore, the values of n encountered in practice may not be sufficiently large for the asymptotic advantage of our method to kick in, although future improvements could cut down on the overhead of our protocol.

6.1 Arbitrary quantum state preparation

Our protocol for fault-tolerant QRAM could be used to prepare an arbitrary n-qubit state, given a list of its 2^n amplitudes stored in classical memory. This subroutine could be useful for preparing initial states for the simulation of the dynamics of many-body systems or ansatz states for quantum phase estimation or variational algorithms. It would also be useful in certain algorithms for quantum machine learning or solving differential equations.

The process for creating an arbitrary state with QRAM goes roughly as follows [14]. First, one classically pre-processes the 2^n complex amplitudes that define the state (comprising $2^{n+1}-2$ independent real degrees of freedom, accounting for normalization and an unphysical global phase) to compute a list of $2^{n+1}-2$ rotation angles. The arbitrary state can be prepared through a sequence of steps labeled by $i = 1, \ldots, n$. On step i, a single-qubit rotation is performed on qubit i by one of 2^{i-1} angles; which angle to use depends on the setting of qubits $1, \ldots, i-1$. If the angles are accessible via QRAM, the quantum algorithm can compute the angle with a single query by using qubits $1, \ldots, i-1$ as the address. Once the angle has been read in, the single-qubit rotation on qubit i can be efficiently performed, and then a second query to the QRAM can be made to uncompute the angle. For more details on this strategy, see for example Refs. [11, 14, 38, 109]. At most $2n$ fault-tolerant QRAM queries would be required, two queries each to datasets of sizes $1, 2, 4, \ldots, 2^n$ angles. These $O(n)$ queries could be implemented fault-tolerantly by our protocol—by Theorem 2, each of the n fault-tolerant queries can be implemented up to error ε/n (so that the total error is ε) using $O(n^3/\varepsilon)$ queries to the physical QRAM device that can coherently access datasets of size up to 2^n, giving a total of $O(n^4/\varepsilon)$ physical queries. However, we note that it is possible the n dependence could be improved by leveraging the fact that only a small fraction of these queries truly require the full 2^n-size QRAM.

Thus, our protocol could represent a great improvement over the $\Omega(2^n)$ fault-tolerant quantum gates (of which at least $\Omega(\sqrt{2^n})$ must be non-Clifford gates [38]) required by a standard circuit approach to state preparation. Since state preparation is typically only one component of a larger algorithm, whether this would be a worthwhile approach within end-to-end applications may depend on the details of the application.

6.2 Quantum machine learning

The quantum linear system algorithm [21] prepares a quantum state $|\mathbf{x}\rangle$ encoding the solution to a well-conditioned $2^n \times 2^n$ linear system $A\mathbf{x} = \mathbf{b}$ using only poly(n) queries to data access oracles for the entries of the matrix A and the vector $\mathbf{b}$. Thus, remarkably, the number of queries needed can be exponentially smaller than the size of the matrices and vectors themselves, due to the ability of the quantum algorithm to explore the exponentially large Hilbert space in superposition. This insight, and the broader framework of quantum linear algebra [31, 104, 110],

has led to a number of quantum algorithms in the realm of machine learning [9, 10], where linear algebra problems are ubiquitous.

When the entries of A and $\mathbf{b}$ have succinct formulas, the data access oracles can be implemented with $\text{poly}(n)$ gate complexity directly on a fault-tolerant quantum processor. However, in big-data applications, it is more relevant to consider A and $\mathbf{b}$ as having arbitrary entries determined by the data, which is stored in classical memory. In these cases, the data access oracles, for example, a unitary block-encoding of A [104, 109, 110] or a state-preparation unitary for $|\mathbf{b}\rangle$ (see Section 6.1) can be implemented with $\text{poly}(n)$ QRAM queries. Since these algorithms will require many quantum gates and many calls to the data access oracles to solve an end-to-end machine learning problem, they will only be possible once fault-tolerant quantum computers are available, and furthermore, to potentially provide exponential speedups, they will require the ability to perform QRAM fault-tolerantly at cost $\text{poly}(n)$ (cheap QRAM assumption from Section 1).

It is worth briefly mentioning that many of these quantum machine learning algorithms have been dequantized [111–116], in the sense that quantum-inspired classical algorithms can also achieve $\text{poly}(n)$ complexity for problems involving datasets of size 2^n, using the ability to query individual entries of the dataset (e.g., via RAM) and also to sample an entry with probability proportional to its magnitude (a classical analogue of arbitrary state preparation). In these cases, the quantum algorithm cannot provide an exponential speedup. Polynomial speedups may still be possible if the cheap QRAM assumption holds. Moreover, these classical methods do not apply in every case; notably, when the matrices involved are sparse and high rank, the possibility of exponential quantum speedup persists [117], provided that the cheap QRAM assumption is true.

To see how our protocol impacts the outlook of these applications, we suppose generically that a quantum machine learning algorithm requires $\text{poly}(n)$ fault-tolerant quantum gates and $\text{poly}(n)$ fault-tolerant queries to QRAM to complete its task. By implementing QRAM using our fault-tolerant QRAM protocol with error parameter $\varepsilon = 1/\text{poly}(n)$, the problem can be solved with $\text{poly}(n)$ fault-tolerant gates and $\text{poly}(n)$ calls to the faulty physical QRAM device, which keeps open the possibility of superpolynomial speedup in quantum resources, here assuming that the computational cost of the physical query is also $\text{poly}(n)$.

However, our protocol also requires classical computations of complexity $O(2^n)$ to perform the update rule. Although this classical complexity may be parallelized, as we discussed in Section 5, it is essentially equivalent to the ability to perform a sparse matrix-vector multiplication for a vector of size 2^n. This is a crucial caveat for machine learning, since sparse matrix-vector multiplication often suffices to efficiently solve the problem in the first place (see discussion in Ref. [8]). For example, the classical conjugate gradient method [118, 119] can invert well-conditioned, sparse linear systems $A\mathbf{x} = \mathbf{b}$ with $\text{poly}(n)$ sparse matrix-vector multiplications. In fact, the number of sparse matrix-vector multiplications required by conjugate gradient has better asymptotic complexity (scaling as the square root of the condition number for positive semidefinite A) than the number of QRAM calls required by the quantum linear system algorithm (scaling at least linearly in the condition number [120]). This suggests that in a generic instance of the sparse linear system problem, it would be more efficient to apply $\Omega(2^n)$ classical computational resources directly toward solving the linear algebra problem, rather than to perform the update rule required by our protocol.

That said, a few opportunities remain. For instance, one can consider the case that the sparse $2^n \times 2^n$ matrix A has repeated entries or some other kind of high-level structure, such that the number of classical degrees of freedom is asymptotically less than 2^n. For concreteness, suppose that the size of the dataset determining A is only $\sqrt{2^n}$. Then, the QRAM operation needs only to be applied for size-$\sqrt{2^n}$ datasets, and the classical complexity of the update rule is only $O(\sqrt{2^n})$, quadratically cheaper than the number of arithmetic operations required for a full matrix-vector multiplication by the matrix A. In this situation, our protocol may enable an end-to-end solution with $\text{poly}(n)$ quantum cost and $O(\sqrt{2^n})$ classical cost, which could represent a quadratic speedup in terms of classical complexity and an exponential speedup in terms of quantum complexity. While quadratic *quantum* speedups are typically thought of as insufficient to overcome the large disadvantage in constant prefactor for quantum computation [73], a quadratic *classical* speedup faces no such obstacles, and could be considered quite large. Generalizing this line of thinking, we may expect to find end-to-end speedups in situations where the best achievable classical complexity is asymptotically greater than the number of classical degrees of freedom of the problem. Toward this end, another example worth exploring may be the inversion of dense matrices with $O(2^{2n})$ degrees of freedom, which generally requires $2^{\omega n}$ classical arithmetic operations with $\omega > 2$. It remains to find concrete end-to-end examples where quantum advantages of this kind may come to fruition.

6.3 Cryptanalysis

Shor's algorithm for factoring or computing the discrete logarithm relies on performing coherent modular arithmetic. Gidney's windowed quantum arithmetic [121] has been used to reduce the cost of Shor's algorithm by replacing some of this costly arithmetic with coherent reads from a quantum lookup table of classically precomputed values [122, 123]. The algorithm can be expressed as repeated blocks of a coherent lookup table read, each followed by a coherent addition.[11] We investigate the performance of our distillation-based QRAM scheme for implementing the coherent lookup table reads in Shor's algorithm. Elliptic curve cryptography (ECC) is a more attractive target than factoring because of the smaller number of calls to the lookup table, which allows a less stringent target ε, as well as the high cost of elliptic curve addition compared to modular addition. We follow the presentation of Ref. [124], where Shor's algorithm for ECC is presented as two applications of quantum phase estimation on unitaries of the form:

$$U_X|R\rangle = |R + X\rangle, \tag{120}$$

where $R = (x, y)$ is an elliptic curve point, and X denotes either the base point P or the public key Q. The goal is to compute integer j such that $Q = [j]P$, where the notation $[j]$ indicates that we are performing elliptic curve scalar multiplication of the point P j times (see Ref. [124] for full definitions of the addion and multiplication operations). Quantum phase estimation uses controlled unitaries of the form $U_X^{2^j}$ for $0 \leq j \leq k-1$, where k denotes the number of bits in the ECC scheme (e.g., ECC-256). In Ref. [124], windowed quantum arithmetic is used to replace blocks of 16 controlled unitaries with a single QROM load of 2^{16} classically pre-computed values. The load is followed by an elliptic curve addition operation. The cost of loading N pieces of classical data from a QROM is N Toffoli gates. For $N = 2^{16}$, this is approximately 6.5×10^4 Toffolis. We note that this is much smaller than the cost of the elliptic curve addition operation, which is approximately 8.34×10^6. The minimum Toffoli cost of the algorithm is obtained by loading groups of 2^{19} values, while the minimum active volume cost is obtained at 2^{16} values.

For the sake of calculation, we suppose that the physical QRAM device produces states with minimum fidelity of $F = 50\%$, and we seek to achieve error $\varepsilon_{\text{tot}} = 0.1$ (where ε_{tot} is the sum of the errors of all QRAM loads in the algorithm). We use the generalized b-bit version of our protocol from Appendix A. If we choose the iterated swap test distillation protocol for its simplicity (each swap test requiring exactly $n + b$ non-Clifford Toffoli gates), for non-vanishing F, the number of non-Clifford gates for distillation scales roughly as $O((1 - F)n^2(n + b)/\varepsilon))$ with ε the error per fault-tolerant QRAM query (note that the twirling and teleportation steps are entirely Clifford). For $F \sim 50\%$, and assuming for simplicity a constant prefactor of 1, we estimate the non-Clifford cost as $n^2(n+b)/(2\varepsilon)$. We use this expression to calculate that the QRAM-based approach achieves its minimum Toffoli and active volume cost at a group size of 2^{64} values, which is impractically large for the classical update step of our scheme. A more realistic size of 2^{32} values only increases the costs by approximately 10%. Nevertheless, both the minimum Toffoli and active volume costs of the existing QROM-based approach are approximately a factor of 1.8× and 1.6× lower than the minimum costs of our QRAM-scheme, respectively, and we have not even considered the computational cost of applying the QRAM device itself. A major contribution to the cost of our method stems from the large amount of data to be loaded ($b = 512$), which features multiplicatively in our Toffoli costs. In contrast, the Toffoli cost of the QROM scheme is independent of the value of b. Improved methods for distillation, especially when b is large, could make our scheme more competitive in this application.

6.4 Chemistry

Coherent data access using a quantum lookup table has become a key subroutine in modern algorithms for quantum chemistry [74]. These algorithms assume access to a block-encoding of the Hamiltonian. This block-encoding is typically implemented by writing the Hamiltonian as a linear combination of unitaries $H = \sum_{j=0}^{L-1} c_j U_j$ and using oracles of the form:

$$\text{PREPARE}|0^{\lceil \log_2(L) \rceil}\rangle = \frac{1}{\sqrt{\lambda}} \sum_{j=0}^{L-1} \sqrt{|c_j|}|j\rangle \tag{121}$$

$$\text{SELECT} = \sum_{j=0}^{L-1} |j\rangle\langle j| \otimes \text{sign}(c_j) U_j + \sum_{j=L}^{2^{\lceil \log_2(L) \rceil} - 1} |j\rangle\langle j| \otimes \mathbb{I}, \tag{122}$$

[11] This could be regular addition, modular addition, or elliptic curve point addition.

where $\lambda = \sum_j |c_j|$ is the normalization factor of the block-encoding. The LCU approach to block-encodings still works if PREPARE results in each computational basis state $|j\rangle$ being entangled with a garbage register. The technique of coherent alias sampling, introduced in Ref. [74], provides an efficient approach for implementing PREPARE in this way, with a cost of $O(L + \log(1/\delta))$ Toffoli gates using QROM, where δ is the largest error in $\sqrt{|c_j|}$. The dependence on L can be improved quadratically using the techniques of Ref. [38]. In the most straightforward application of these techniques (referred to as "sparse qubitization" [125]) the complexity of PREPARE dominates the algorithm, and its cost depends on the cost to load the Hamiltonian coefficients from a quantum lookup table.

Sparse qubitization requires $O(\lambda/\Delta)$ calls to the block-encoding of the Hamiltonian, where Δ is the precision on the energy estimate of the Hamiltonian. For typical λ values of small molecules, in the range 10^2–10^4 Hartree, and $\Delta = 10^{-3}$ Hartree, this implies at least 10^5 calls to the quantum lookup table. If we were to replace the calls to the quantum lookup table with our distillation-based QRAM scheme for the example of FeMo-co ($\lambda = 7614$, $N = L = 179{,}498$, $b = 84$ [125]), using the same methodology as in the previous section, we conclude that we would require at least 2.5×10^{11} Toffoli gates per call to the block-encoding, substantially larger than the value of 10^4 Toffolis in Ref. [125]. We note that even if the dependence of our approach on the error ε per query could be reduced to $O(\log(1/\varepsilon))$, it is unclear whether our approach would improve over existing methods, as the amount of data (i.e., Lb) loaded for the chemistry Hamiltonian is sufficiently modest that the QROM scaling of $O(\sqrt{Lb})$ is comparable to the scaling of our protocol $O(\log^2(L)(\log(L) + b))$ (Theorem 3).

The sparse qubitization approach has been superseded in some applications by more efficient methods [126, 127], which also use a lookup table to coherently load angles for basis rotations that are used in the SELECT oracle. We refer to Refs. [127, 128] for a detailed accounting of the contribution of the lookup table reads to the total gate and space complexity. It is unlikely that our aproach to implementing QRAM would reduce the costs of these algorithms, at least in its current form.

7 Outlook on the cheap QRAM assumption

A central goal of this research program is to determine whether QRAM can be considered equally cheap as RAM—at least in an abstract, asymptotic sense—or whether its quantum nature makes QRAM fundamentally more difficult, preventing a justification of the cheap QRAM assumption from Section 1. The central aspect of QRAM that differentiates it from RAM is the need to protect the quantum information about which address (or superposition of addresses) is being queried, even when the hardware has errors. To emphasize this distinction, it is worth noting that a fault-tolerant RAM could be easily constructed from a faulty RAM device using a simple repetition code: by repeatedly querying the faulty RAM device on the same input address, one can efficiently boost the probability of a successful RAM query, provided that as the device has nonzero bias in favor of the correct answer. In contrast, this same strategy fails for QRAM because even a single query to a faulty QRAM device may leak the address information and decohere the address register. Formally speaking, we might say that logical (Q)RAM is a transversal gate for the repetition code, but that this is not sufficient for fault-tolerant QRAM, since the repetition code does not protect against phase-flip errors. For fault-tolerant QRAM, a more sophisticated strategy is required.

Our protocol provides such a method for fault-tolerant QRAM, successfully protecting which address state $\sum_x \alpha_x |\overline{x}\rangle$ is being queried without performing QEC on the $\Omega(2^n)$ components of the QRAM device. It does this by relying on resource states (see Eq. (9)) that are equal superpositions of all 2^n addresses, independent of which superposition of addresses one wants to query—the noisy device is prevented from direct interaction with the address information. However, these resource states have exponentially small amplitude on any individual address. Consequently, if the data at one address is modified, the ideal resource state barely changes. This begs the question: if these resource states are insensitive to the underlying bits in the dataset, how, then, can they be used to insert information about the dataset into the quantum state? Our protocol achieves this by iteratively and adaptively inserting *global* information about the dataset f—that is, properties of f that depend on all 2^n data entries—into the quantum state. One way to see this is by decomposing $f(x)$ into a sum over global, oscillatory contributions of the form $(-1)^{k\cdot x}$, that is, its Fourier expansion

$$f(x) = \frac{1}{2^{n/2}} \sum_{k\in\{0,1\}^n} (-1)^{k\cdot x} \tilde{f}(k), \tag{123}$$

where $\tilde{f}$ is the Walsh–Hadamard transform of f. When we teleport the resource state $|\overline{\Psi(f)}\rangle$, we enact the QRAM unitary $\overline{V(f^{\oplus m})}$ instead of $\overline{V(f)}$, where m is uniformly random and $f^{\oplus m}$ is defined by $f^{\oplus m}(x) = f(x \oplus m)$. While the datasets $f^{\oplus m}$ and f are not guaranteed to agree on any of the addresses, they have a close relationship in

frequency space. One can see from Eq. (123) that

$$\tilde{f}^{\oplus m}(k) = (-1)^{k\cdot m}\tilde{f}(k)\,. \tag{124}$$

Thus, $\tilde{f}^{\oplus m}(k) = \tilde{f}(k)$ for half the values of k and $\tilde{f}^{\oplus m}(k) = -\tilde{f}(k)$ for the other half (except in the unlikely case that $m = 0^n$, in which case they would agree on all values). In a sense, we may say that by implementing $\overline{V(f^{\oplus m})}$, we have successfully inserted half of the information about the function f into the quantum state, regardless of which random m is obtained. To insert the other half of the information, the protocol updates the dataset to $f' = f \oplus f^{\oplus m}$. We may observe that f' is periodic in translation by m (i.e., $f'(x) = f'(x \oplus m)$ for all x), and thus $\tilde{f}'(k) = 0$ for any k for which $k \cdot m = 1$ (half the values of k). Each successive correction function will have half as many nonzero Fourier components as the previous one (assuming the n-bit measurement outcomes in previous rounds form a linearly independent set), until finally after n rounds there is no remaining frequency information left to insert. This global approach to information insertion appears crucial for correctly applying QRAM on any input state, even while not knowing or learning what that input state is.

The cost of this global approach is adaptivity and classical computation. At each iteration, we do not have control over which half of the frequency information we insert into the quantum state; this is determined by a uniformly random measurement outcome m. After receiving m, the protocol must adaptively update the *entire* dataset to essentially remove the half of the global information that has already been applied to the quantum state. This removal requires touching all 2^n entries of the dataset, and then giving the physical QRAM device access to the new dataset. As discussed in Section 5, this updating of frequency information is a non-negligible classical computation, essentially equivalent in complexity to performing the Walsh–Hadamard transformation on the dataset. While the Walsh–Hadamard transform may be parallelizable, it appears to be a harder calculation than a normal RAM query; for example, it requires fundamentally greater wire density than RAM.

The main theoretical open question, then, is to determine if the classical complexity of the protocol can be reduced to a quantity more similar to a RAM query, or otherwise find a way to justify that the overall (i.e., both classical and quantum) cost of QRAM is poly(n). Our protocol makes some progress in this direction, and it offers clear advantages over actively error corrected circuit QRAM at the practical level—enabling the usage of a specialized low-fidelity QRAM device, and eliminating the need for massively parallel QEC. However, in a theoretical sense, both our protocol and circuit QRAM ultimately require the same $\Omega(2^n)$ scaling of classical resources to protect the address information from decoherence. One is left to wonder whether this scaling of classical complexity could be a fundamental requirement for achieving fault-tolerant QRAM.

Acknowledgments

The authors thank Mario Berta, Joe Iverson, Sam Jaques, Michael Kastoryano, Fernando Pastawski, Samson Wang, and John Wright for helpful conversations. We also thank Simone Severini, James Hamilton, Nafea Bshara, Peter DeSantis, and Andy Jassy for their involvement and support of the research activities at the AWS Center for Quantum Computing.

A Generalization of the protocol to multiple output bits

The main text provided a protocol to implement the single-bit diagonal QRAM operation of Eq. (1), where the classical data table f is an n-bit Boolean function consisting of a single bit stored at each of 2^n addresses. The operation applies a sign to each corresponding basis state $|x\rangle \mapsto (-1)^{f(x)}|x\rangle$.

In many applications, there are $b > 1$ bits stored at each of the 2^n addresses in memory, and often we wish to coherently read all b bits into a separate bus register, that is, perform the operation (cf. Eq. (1))

$$\sum_{x\in\{0,1\}^n}\sum_{u\in\{0,1\}^b}\alpha_{x,u}|x\rangle|u\rangle \quad \overset{U(f)}{\longmapsto} \quad \sum_{x\in\{0,1\}^n}\sum_{u\in\{0,1\}^b}\alpha_{x,u}|x\rangle|u\oplus f(x)\rangle \tag{125}$$

We refer to the first register storing $|x\rangle$ as the address register, and the second register storing the data as the bus register. Our protocol can also straightoforwardly handle this case, although it requires introducing a bit more notation and slightly more overhead.

A.1 Definitions

First, let $\mathcal{F}_n$ denote the set of n-bit Boolean functions $f\colon \{0,1\}^n \to \{0,1\}$ (considered as data tables in the main text). Next, let $\mathcal{F}_n^{(b)}$ denote the set of n-bit functions with b bits of output $f\colon \{0,1\}^n \to \{0,1\}^b$. We wish to implement the QRAM operation $U(f)$ of Eq. (125) for arbitrary data tables $f \in \mathcal{F}_n^{(b)}$.

To do so, we will need to generalize the set $\mathcal{F}_n^{(b)}$ to include an extra "sign" bit. These are essentially equivalent to n-bit functions with $b+1$ bits of output, that is, the set $\mathcal{F}_n^{(b+1)}$, but where the sign bit plays a different role than the other b bits. Formally, we let $\mathcal{F}_n^{(b\pm)} = \mathcal{F}_n \times \mathcal{F}_n^{(b)}$, and we denote elements of $f \in \mathcal{F}_n^{(b\pm)}$ by a pair $f = (f_\pm, f_\square)$, where $f_\pm$ is an n-bit Boolean function that dictates the "sign" and $f_\square$ is an n-bit function with b bits of output.

For any $f \in \mathcal{F}_n^{(b\pm)}$, we define a unique Boolean function $\hat{f} \in \mathcal{F}_{n+b}$ to act on inputs of $n+b$ bits denoted by pairs (x,u) with $x \in \{0,1\}^n$ and $u \in \{0,1\}^b$:

$$\hat{f}(x,u) = f_\pm(x) \oplus \left[\bigoplus_{i=1}^{b} u_i f_i(x)\right] = f_\pm(x) \oplus (u \cdot f_\square(x))\,, \tag{126}$$

where f_i is the i-th output bit of $f_\square$, and in the final expression, $\cdot$ denotes the standard dot product (modulo 2) between a pair of b-bit strings, viewed as vectors. We can see that any $\hat{f}$ obtained from f via Eq. (126) contains all the same information as f, just represented in a different form. We can also see that

$$\deg(\hat{f}) = \max\left(\deg(f_\pm), 1 + \max_{i\in[b]} \deg(f_i)\right) \leq n+1\,. \tag{127}$$

This is an important observation—since $\hat{f}$ is a Boolean function on $n+b$ bits, one would naively expect its degree could be as large as $n+b$, but due to the structure of Eq. (126), the actual degree is independent of b.

Having defined signed Boolean functions, we can generalize the standard QRAM operation $U(f)$ of Eq. (125) slightly by allowing the QRAM to also conditionally apply signs to the computational basis states, controlled by another Boolean function $f_\pm$. Specifically, given a signed n-bit function $f = (f_\pm, f_\square) \in \mathcal{F}_n^{(b\pm)}$, we define

$$U(f)|x\rangle|u\rangle = (-1)^{f_\pm(x)}|x\rangle|u \oplus f_\square(x)\rangle\,. \tag{128}$$

We may think of $f \in \mathcal{F}_n^{(b\pm)}$ as a data table storing $b+1$ bits of data at each of 2^n addresses. In fact, we may think of the main text as the special case of Eq. (128) where $b=0$ and there is only a sign bit.

In this appendix, we will be working with the more general signed QRAM function from Eq. (128) and assuming $f \in \mathcal{F}_n^{(b\pm)}$; we may always take $f_\pm = \mathbf{0}$ (the zero function) to recover the standard QRAM from Eq. (125).

A.2 Generalized diagonal QRAM

The key to implenting $U(f)$ will be to convert it to a diagonal operation generalizing $V(f)$ from Eq. (1), so that the same approach as was used in the main text can be applied. We define $V(f)$ for $f \in \mathcal{F}_n^{(b\pm)}$ to be equal to the QRAM operation $U(f)$ from Eq. (128) conjugated by the Hadamard transform on the bus register

$$V(f) = (\mathbb{I} \otimes H^{\otimes b})\, U(f)\, (\mathbb{I} \otimes H^{\otimes b})\,. \tag{129}$$

A straightforward calculation shows that $V(f)$ is diagonal, with the diagonal entries given by $(-1)^{\hat{f}(z)}$ for z a 2^{n+b}-bit string:

$$V(f)|x\rangle|u\rangle = (-1)^{\hat{f}(x,u)}|x\rangle|u\rangle\,. \tag{130}$$

This fact is what motivated the definition of $\hat{f}$ in Eq. (126). Thus, this definition of $V(f)$ is seen to be an immediate generalization of Eq. (1), with the function f replaced by $\hat{f}$.

As one example, if $b=1$ and $f_\pm = \mathbf{0}$, then the function $\hat{f}(x,u) = u_1 f_\square(x)$ is an $(n+1)$-bit Boolean function. The operation $V(f)$ using the definition of $V(f)$ from Eq. (130) is equivalent to $V(\hat{f})$ for the definition of $V(f)$ from Eq. (1). Furthermore, the operation can be thought of as a controlled $V(f_\square)$ operation, with the single bus qubit acting as the control, as alluded to in Footnote 2.

We will also need to generalize the update rule of Eq. (101). Suppose we modify the function $\hat{f}$ by adding (modulo 2) a fixed bit string $m \in \{0,1\}^{n+b}$ to the address and bus registers before evaluating the function. We may decompose $m = (m_A, m_B)$ where $m_A \in \{0,1\}^n$ and $m_B \in \{0,1\}^b$. We observe that

$$\hat{f}(x \oplus m_A, u \oplus m_B) = f_\pm(x \oplus m_A) \oplus \left[\bigoplus_{i=1}^{b} (u_i \oplus m_{B,i}) f_i(x \oplus m_A) \right]. \tag{131}$$

Next, for any signed n-bit function $f = (f_\pm, f_1, \ldots, f_b) \in \mathcal{F}_n^{(b\pm)}$ and any $m \in \{0,1\}^{n+b}$, we may define a unique function $f^{\oplus m} = (f_\pm^{\oplus m}, f_1^{\oplus m}, \ldots, f_b^{\oplus m}) \in \mathcal{F}_n^{(b\pm)}$ by the relations

$$\begin{aligned} f_\pm^{\oplus m}(x) &= f_\pm(x \oplus m_A) \oplus \left[\bigoplus_{i=1}^{b} m_{B,i} f_i(x \oplus m_A) \right] \\ f_i^{\oplus m}(x) &= f_i(x \oplus m_A) \qquad \text{for } i = 1, \ldots, b\,. \end{aligned} \tag{132}$$

This definition generalizes the one from Eq. (88), and it was chosen to ensure that

$$\hat{f}(z \oplus m) = \hat{f}^{\oplus m}(z)\,. \tag{133}$$

A.3 Protocol for generalized QRAM

It is now simple to apply the ideas from the main text to fault-tolerantly implement the generalized $\overline{V(f)}$ from Eq. (130), which amounts to applying a sign $(-1)^{\hat{f}(x,u)}$ to each basis state $|\overline{x}\rangle|\overline{u}\rangle$. The unitary $\overline{U(f)}$ can thus be implemented by conjugating $\overline{V(f)}$ with fault-tolerant Hadamard operations on the bus register (Eq. (129)).

To implement the sign $(-1)^{\hat{f}(x,u)}$, we could simply view $\hat{f}$ as a Boolean function on $n+b$ bits of degree at most $n+1$ (Eq. (127)). We could store all 2^{n+b} outputs of this function $\hat{f}$ in classical memory, and apply the protocol from the main text. The drawback of this approach is that it does not take advantage of the structure of the function $\hat{f}$, which has only $(b+1)2^n$ binary degrees of freedom rather than 2^{n+b}. If b is large it would be classically intractable to store all 2^{n+b} outputs of $\hat{f}$ in classical memory, and to perform the update rule directly as $g \leftarrow \mathrm{UR}(g, m)$ when the data table g has 2^{n+b} classical bits.

Fortunately, we may apply our protocol without storing and manipulating 2^{n+b} classical bits. Rather than keeping track of $\hat{g} \in \mathcal{F}_{n+b}$ explicitly and applying an update rule $\hat{g} \leftarrow \mathrm{UR}(\hat{g}, m)$, we simply maintain a more compact and natural representation of $g = (g_\pm, g_\square)$ without any redundancy. We preserve the form of the update rule from Eq. (101)

$$\mathrm{UR}(g, m) = g \oplus g^{\oplus m} \tag{134}$$

except that we interpret $g^{\oplus m}$ using Eq. (132). Due to Eq. (133), this update rule ensures that if $h = \mathrm{UR}(g, m)$ then $\hat{h}(z) = \hat{g}(z) \oplus \hat{g}(z \oplus m)$, even though we are not storing all 2^{n+b} outputs of $\hat{h}$ in classical memory. Indeed, by examining Eq. (132), we see that we may implement the update rule by applying the original update rule of Eq. (101) to each of the $b+1$ bits of g using the bit string $m_A \in \{0,1\}^n$, and then additinally updating $g_\pm$ by adding (modulo 2) the i-th bit of $g_\square$ to $g_\pm$ for each i where the i-th bit of m_B is 1.

Due to the fact that the degree of $\hat{g}$ is at most $n+1$, the number of rounds required is at most $n+1$, regardless of the size of b.

Besides the update rule, the protocol proceeds as if the function being applied were an $(n+b)$-bit function, rather than an n-bit function. The resource states created by the oracle are $(n+b)$-qubit physical states (cf. Eq. (9))

$$|\Psi(g)\rangle = \frac{1}{2^{(n+b)/2}} \sum_{x \in \{0,1\}^n} \sum_{u \in \{0,1\}^b} (-1)^{\hat{g}(x,u)} |x\rangle|u\rangle \tag{135}$$

$$= \frac{1}{2^{(n+b)/2}} \sum_{z \in \{0,1\}^{n+b}} (-1)^{\hat{g}(z)} |z\rangle \tag{136}$$

which can be created by applying the physical unitary $V(g)$ onto the state $|+\rangle^{\otimes(n+b)}$. The encoding protocol must now encode $n+b$ physical qubits into $n+b$ logical qubits, and the encoding error grows linearly in $n+b$ as

$\varepsilon_{\text{enc}} = O(\sqrt{p}(n+b))$ or better (straightforward generalization of the results of Section 4.2). The distillation protocol based on the iterated swap test or quantum PCA with fractional swap gates proceeds identically, except that each controlled swap operation now requires $n+b$ fault-tolerant controlled swap gates on qubit systems, instead of only n. The teleportation procedure also now requires $n+b$ CNOT gates, rather than only n. As a result, we can state the following generalized theorem.

Theorem 3. *Let $f \in \mathcal{F}_n^{(b\pm)}$ be an arbitrary signed n-bit function with b output bits, for which we wish to implement the fault-tolerant diagonal QRAM unitary $\overline{U(f)}$ of Eq. (128), and let ε be an error parameter. The same results hold as in Theorem 2, provided that the physical QRAM device can perform the physical version of the generalized QRAM operation of Eq. (130) on $n+b$ qubits with the same guarantees, and that the physical error rate of the main processor satisfies $p(n+b)^2 = O(1)$. The complexity of the protocol is the same with an identical asymptotic expression for Q, and*

$$Q' = O\big((n+b)^2 Q\big)\,. \tag{137}$$

Additionally, the protocol applies the (generalized) classical update rule $n+1$ times (rather than only n), and the (generalized) classical partial Clifford twirling operation $g \mapsto g_C$ (Eq. (35)) Q times, each time on a data table of size $(b+1)2^n$ classical bits.

Proof. This follows straightforwardly from the observation that the same number of resource states are needed, but the Clifford twirling, distillation, and teleportation steps now act on states of $n+b$ qubits rather than only n. □

B Delayed proofs for encoding error

Here we restate and provide the full proofs of the two propositions in Section 4.2.

Proposition 1 (Restated)**.** *Denote the fidelity of the physical state by $F(g)_{\text{phys}} = \langle\Psi(g)|\widetilde{\psi}(g)|\Psi(g)\rangle$. Suppose the processor is subject to circuit-level stochastic noise with strength p (Definition 2), and et $\mathcal{E}'_{\text{FT}}$ be a fault-tolerant encoding channel with encoding error ε_{enc} (as in Definition 3). Then, there exists another fault-tolerant encoding channel $\mathcal{E}_{\text{FT}}$ (formed by Pauli-twirling $\mathcal{E}'_{\text{FT}}$), for which*

$$\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq (1-p)^n\Big((1-3\varepsilon_{\text{enc}})F(g)_{\text{phys}} - 2\Gamma(\mathcal{E})\Big)\,, \tag{29}$$

where $\overline{\phi(g)}$ is defined from $\mathcal{E}_{\text{FT}}$ as in Eq. (25), and $\Gamma(\mathcal{E})$ is a quantity that vanishes with increasing code size, provided the physical error rate p is below a constant threshold, as discussed in Section 3.3.

Proof. There are broadly two components to this proof. First, we will define $\mathcal{E}_{\text{FT}}$ based on a Pauli twirl of $\mathcal{E}'_{\text{FT}}$ and show that it is a stochastic channel with a large identity component, up to a small correction. Second, we will show how the channel being stochastic implies the claimed fidelity statement.

We begin with the Pauli twirl. Formally, we let $\hat{\mathbb{P}}$ denote the set of n-qubit Pauli operators, that is, the subset of the signed Pauli set $\mathbb{P}$ defined in Definition 5, where the sign bit s is fixed to $s=0$. We define $\mathcal{E}_{\text{FT}}$ to be the procedure that (i) chooses a Pauli $G \in \hat{\mathbb{P}}$ uniformly at random, (ii) applies physical G, (iii) applies $\mathcal{E}'_{\text{FT}}$ to encode the state, (iv) applies the fault-tolerant QEC gadget $\mathcal{Q}_{\text{FT}}$, (v) applies the fault-tolerant gadget for logical $\overline{G}$ on the encoded state. In the absence of noise, the physical G and logical $\overline{G}$ will cancel out, and encoding will be perfect. The purpose of including them is to be able to guarantee that in the presence of noise the overall channel is a stochastic channel, as we will explain shortly. Let the physical channel implemented in step (ii) be denoted by $\widetilde{\mathcal{G}}$, and let the channel enacted in step (v) on the encoded system be denoted by $\widetilde{\mathcal{G}}_{\text{FT}}$. The channel implemented by the noisy encoding procedure is thus given by an average over all choices of G (denoted $\mathbb{E}_{G\sim\hat{\mathbb{P}}}$), as follows

$$\widetilde{\mathcal{E}}_{\text{FT}} = \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \widetilde{\mathcal{G}}_{\text{FT}} \circ \widetilde{\mathcal{Q}}_{\text{FT}} \circ \widetilde{\mathcal{E}}'_{\text{FT}} \circ \widetilde{\mathcal{G}}\,. \tag{138}$$

Since we have assumed circuit-level stochastic noise (Definition 2) with strength p, and the physical G has n circuit locations, we may write

$$\widetilde{\mathcal{G}} = (1-p)^n\mathcal{G} + (1-(1-p)^n)\mathcal{N}_G \tag{139}$$

for some CPTP map $\mathcal{N}_G$, where $\mathcal{G}[\cdot] = G[\cdot]G^\dagger$ is the ideal channel (note that for $G \in \hat{\mathbb{P}}$, we have $G = G^\dagger$). Furthermore, since $\mathcal{G}_{\mathrm{FT}}$ is fault-tolerant, we have that

$$\frac{1}{2}\left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{G}}_{\mathrm{FT}} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} - \overline{\mathcal{G}} \circ \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \right\|_\diamond \leq \Gamma(\mathcal{E}) \,, \tag{140}$$

where $\Gamma(\mathcal{E})$ vanishes with increasing code size, as in Eq. (19). Let the superoperator inside the diamond norm in the equation above be denoted by $\mathcal{R}$, which satisfies $\|\mathcal{R}\|_\diamond \leq 2\Gamma(\mathcal{E})$. Combining the above equations, we may write

$$\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}} = \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \left[(\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{G}}_{\mathrm{FT}} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}}) \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}} \circ \widetilde{\mathcal{G}} \right] = \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \left[(\mathcal{R} + \overline{\mathcal{G}} \circ \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}}) \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}} \circ \widetilde{\mathcal{G}} \right] \tag{141}$$

$$= (1-p)^n \left(\mathcal{R}_1 + \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \overline{\mathcal{G}} \circ (\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}}) \circ \mathcal{G} \right) + (1 - (1-p)^n) \mathcal{R}_2 \tag{142}$$

where $\mathcal{R}_1 = \mathbb{E}_{G\in\hat{\mathbb{P}}} \mathcal{R} \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}} \circ \mathcal{G}$, which also satisfies $\|\mathcal{R}_1\|_\diamond \leq 2\Gamma(\mathcal{E})$, and the term $\mathcal{R}_2$ is a CPTP map that collects the contribution $\mathcal{N}_G$ from Eq. (139). This equation is a stochastic mixture of two CPTP channels (viewing the term in parentheses as one channel and $\mathcal{R}_2$ as the other), and each contributes non-negatively to the overall fidelity. Thus, recalling that $\overline{\phi(g)} = \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)]$, we can ignore the contribution of the second term and say

$$\langle \overline{\Psi(g)} | \overline{\phi(g)} | \overline{\Psi(g)} \rangle = \langle \overline{\Psi(g)} | \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\widetilde{\psi}(g)] | \overline{\Psi(g)} \rangle \tag{143}$$

$$\geq (1-p)^n \langle \overline{\Psi(g)} | \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \overline{G} \Big(\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}}[G\widetilde{\psi}(g)G] \Big) \overline{G} | \overline{\Psi(g)} \rangle - (1-p)^n 2\Gamma(\mathcal{E}) \,, \tag{144}$$

since $\langle \overline{\Psi(g)} | \mathcal{R}_1[\widetilde{\psi}(g)] | \overline{\Psi(g)} \rangle \leq \left\| \mathcal{R}_1[\widetilde{\psi}(g)] \right\|_1 \leq \|\mathcal{R}_1\|_\diamond$.

We may generically decompose the channel $\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}}[\cdot] = \sum_{P,P'\in\hat{\mathbb{P}}} \chi_{P,P'} \overline{P} \mathcal{E}[\cdot] \overline{P'}$ as a sum over logical Pauli operators acting on the left and right of the encoded state, weighted by numbers $\chi_{P,P'}$ arranged into a matrix. Now, defining the channel $\mathcal{N}[\rho] = \mathbb{E}_{G\sim\hat{\mathbb{P}}} \overline{G}(\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\mathrm{FT}} \circ \widetilde{\mathcal{E}}'_{\mathrm{FT}}[G\rho G])\overline{G}$, we have

$$\mathcal{N}[\rho] = \sum_{P,P'\in\hat{\mathbb{P}}} \chi_{P,P'} \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \overline{G}\,\overline{P}\mathcal{E}[G\rho G]\overline{P'}\,\overline{G} = \sum_{P,P'\in\hat{\mathbb{P}}} \chi_{P,P'} \mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \overline{G}\,\overline{P}\,\overline{G}\mathcal{E}[\rho]\overline{G}\,\overline{P'}\,\overline{G} \,, \tag{145}$$

where we have noted that the physical G acting prior to perfect encoding is equivalent to logical $\overline{G}$ acting after perfect encoding. Examining one term from this expansion, we have

$$\mathop{\mathbb{E}}_{G\sim\hat{\mathbb{P}}} \overline{G}\,\overline{P}\,\overline{G}\,\mathcal{E}[\rho]\,\overline{G}\,\overline{P'}\,\overline{G} = \delta_{P,P'} \overline{P}\,\mathcal{E}[\rho]\,\overline{P'} \tag{146}$$

where $\delta_{P,P'}$ is the Kronecker delta that equals 1 if and only if $P = P'$. This is true because $\overline{G}\,\overline{P}\,\overline{G} = \pm\overline{P}$ for every G, and if $P \neq P'$, then exactly $1/2$ of the choices for G will lead to a net positive sign and $1/2$ will lead to a net negative sign.

Since the "offdiagonal" terms with $P \neq P'$ vanish, this establishes that the channel $\mathcal{N}$ is a stochastic Pauli channel and we may write it as $\mathcal{N}[\rho] = (1-\delta)\mathcal{E}[\rho] + \sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}} \chi_{P,P} \overline{P}\mathcal{E}[\rho]\overline{P}$, where $1 - \delta = \chi_{\mathbb{I},\mathbb{I}}$. Since $\mathcal{N}$ is CPTP by construction, we have $\chi_{P,P} \geq 0$ for all P [94, Lemma 5.2.4] and $\sum_{P\in\hat{\mathbb{P}}} \chi_{P,P} = 1$, implying $0 \leq \delta \leq 1$. We would now like to bound its contribution to the fidelity. We have

$$\langle \overline{\Psi(g)} | \mathcal{N}[\widetilde{\psi}(g)] | \overline{\Psi(g)} \rangle = (1-\delta) \langle \overline{\Psi(g)} | \mathcal{E}[\widetilde{\psi}(g)] | \overline{\Psi(g)} \rangle + \sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}} \chi_{P,P} \langle \overline{\Psi(g)} | \overline{P}\,\mathcal{E}[\widetilde{\psi}(g)]\,\overline{P} | \overline{\Psi(g)} \rangle \tag{147}$$

$$= (1-\delta) \langle \Psi(g) | \widetilde{\psi}(g) | \Psi(g) \rangle + \sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}} \chi_P \langle \Psi(g) | P\,\widetilde{\psi}(g)\,P | \Psi(g) \rangle \tag{148}$$

$$\geq (1-\delta) \langle \Psi(g) | \widetilde{\psi}(g) | \Psi(g) \rangle = (1-\delta) F(g)_{\mathrm{phys}} \,. \tag{149}$$

In the second part of the proof, we aim to bound δ in terms of $\varepsilon_{\mathrm{enc}}$ by using the fact that $\mathcal{N}$ is close to the identity channel. By assumption, $\mathcal{E}'_{\mathrm{FT}}$ has encoding error $\varepsilon_{\mathrm{enc}}$ (recall Definition 3). We begin by observing that the

twirled channel $\mathcal{N}$ can only have smaller encoding error than $\mathcal{E}'_{\mathrm{FT}}$.

$$\varepsilon_{\mathrm{enc}} = \sup_\rho \frac{1}{2}\Big\|\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}'_{\mathrm{FT}}[\rho]-\overline{\rho}\Big\|_1 = \sup_\rho \frac{1}{2}\Big\|\overline{\mathcal{G}}\circ\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}'_{\mathrm{FT}}[\rho]-\overline{G}\,\overline{\rho}\,\overline{G}\Big\|_1 \tag{150}$$

$$= \sup_\rho \frac{1}{2}\Big\|\overline{\mathcal{G}}\circ\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}'_{\mathrm{FT}}[G\rho G]-\overline{\rho}\Big\|_1 = \mathop{\mathbb{E}}_{G\in\hat{\mathbb{P}}}\sup_\rho \frac{1}{2}\Big\|\overline{\mathcal{G}}\circ\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}'_{\mathrm{FT}}[G\rho G]-\overline{\rho}\Big\|_1 \tag{151}$$

$$\geq \sup_\rho \frac{1}{2}\Big\|\mathop{\mathbb{E}}_{G\in\hat{\mathbb{P}}}\overline{\mathcal{G}}\circ\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}'_{\mathrm{FT}}[G\rho G]-\overline{\rho}\Big\|_1 = \sup_\rho \frac{1}{2}\|\mathcal{N}[\rho]-\overline{\rho}\|_1 \tag{152}$$

Then, we may substitute the Pauli form of $\mathcal{N}$ to evaluate

$$\varepsilon_{\mathrm{enc}} \geq \sup_\rho \frac{1}{2}\Big\|-\delta\overline{\rho}+\sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}}\chi_{P,P}\overline{P}\overline{\rho}\overline{P}\Big\|_1 \geq \mathop{\mathbb{E}}_{|\psi\rangle\sim\mathrm{Haar}}\frac{1}{2}\Big\|-\delta|\overline{\psi}\rangle\langle\overline{\psi}|+\sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}}\chi_{P,P}\overline{P}|\overline{\psi}\rangle\langle\overline{\psi}|\overline{P}\Big\|_1 \tag{153}$$

where expectation value denotes drawing $|\psi\rangle$ uniformly from the Haar measure over n-qubit states. Generally, for any Hermitian operator M and any pure state $|\xi\rangle$, we have $\|M\|_1 \geq |\langle\xi|M|\xi\rangle|$. Thus, we may write

$$\varepsilon_{\mathrm{enc}} \geq \mathop{\mathbb{E}}_{|\psi\rangle\sim\mathrm{Haar}}\frac{1}{2}\Big|\langle\overline{\psi}|\Big(-\delta|\overline{\psi}\rangle\langle\overline{\psi}|+\sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}}\chi_{P,P}\overline{P}|\overline{\psi}\rangle\langle\overline{\psi}|\overline{P}\Big)|\overline{\psi}\rangle\Big| \tag{154}$$

$$\geq \frac{1}{2}\delta-\sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}}\frac{\chi_{P,P}}{2}\mathop{\mathbb{E}}_{|\psi\rangle\sim\mathrm{Haar}}\mathrm{Tr}\big((|\overline{\psi}\rangle\langle\overline{\psi}|\otimes|\overline{\psi}\rangle\langle\overline{\psi}|)(\overline{P}\otimes\overline{P})\big) \tag{155}$$

$$= \frac{1}{2}\delta-\sum_{P\in\hat{\mathbb{P}}\setminus\{\mathbb{I}\}}\frac{\chi_{P,P}}{2}\frac{1}{2^n+1} = \frac{1}{2}\delta-\frac{1}{2(2^n+1)}\delta = \frac{2^n}{2^n+1}\frac{\delta}{2}\geq\frac{\delta}{3} \tag{156}$$

where we have used that the expectation over the Haar measure of $|\overline{\psi}\rangle\langle\overline{\psi}|^{\otimes 2}$ is $(\overline{\mathbb{I}}\otimes\overline{\mathbb{I}}+\overline{\mathbb{S}})/(d(d+1))$, with $\mathbb{S}$ the swap operator and $d=2^n$ [129]. Then, we have evaluated $\mathrm{Tr}\big(\overline{P}\otimes\overline{P}\big)=0$ and $\mathrm{Tr}\big(\overline{\mathbb{S}}(\overline{P}\otimes\overline{P})\big)=\mathrm{Tr}\big(\overline{\mathbb{I}}\big)=d$. This confirms that $\delta\leq 3\varepsilon_{\mathrm{enc}}$.

To conclude, we combine the fact that $\delta\leq 3\varepsilon_{\mathrm{enc}}$ with Eq. (144) and Eq. (149) to say

$$\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq (1-p)^n\Big((1-3\varepsilon_{\mathrm{enc}})F(g)_{\mathrm{phys}}-2\Gamma(\mathcal{E})\Big) \tag{157}$$

which implies the proposition statement. □

Proposition 2 (Restated)**.** *Consider a family of quantum error-correcting codes encoding n logical qubits into a codespace, and suppose that this family has a threshold p_0 with respect to local stochastic noise (implying Eq. (18)). Correspondingly, for a particular instance of the family (labeled by its encoding map $\mathcal{E}$), let $\mathcal{Q}$ be the ideal QEC projector, $\mathcal{Q}_{\mathrm{FT}}$ be the fault-tolerant QEC gadget, and $\Gamma(\mathcal{E})$ be the logical error suppression function (see Section 3.3). Then, there exists a fault-tolerant encoding procedure $\mathcal{E}_{\mathrm{FT}}$ of size $|\mathcal{E}|\cdot\mathrm{poly}(k)$, such that, when implemented under the circuit-level stochastic noise model (Definition 2), the encoding error as defined in Definition 3 satisfies*

$$\varepsilon_{\mathrm{enc}} = \sup_\rho \frac{1}{2}\|\mathcal{Q}\circ\widetilde{\mathcal{Q}}_{\mathrm{FT}}\circ\widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho]-\mathcal{E}[\rho]\|_1 \ \leq \Gamma(\mathcal{E})+2\sqrt{cp}n+2|\mathcal{E}|(cp)^k,$$

where C is an absolute constant and k can take values from a sequence of geometrically increasing integers, provided that the physical error rate p is below some constant threshold.

Proof. We will use the fault-tolerant quantum input-output framework introduced by Christandl, Fawzi, and Goswami [59], which tells us that, intuitively, a quantum circuit can be implemented fault-tolerantly on a noisy device except at the input and output layers of the circuit, during which correlated errors can occur with a strength that depends on the underlying circuit-level noise model. Specifically, Ref. [59, Theorem 59] states that, given an ideal quantum circuit of size $|\mathcal{E}|$ that implements map $\mathcal{E}$, we can efficiently find a fault-tolerant circuit $\mathcal{E}_{\mathrm{FT}}$ of size $|\mathcal{E}|\cdot\mathrm{poly}(k)$ with the same input and output systems, such that its noisy implementation $\widetilde{\mathcal{E}}_{\mathrm{FT}}$ subject to circuit-level

stochastic noise with parameter p (Definition 2) satisfies

$$\left\| \widetilde{\mathcal{E}}_{\mathrm{FT}} - \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right) \right\|_\diamond \leq 4|\mathcal{E}|(cp)^k, \tag{158}$$

for some absolute constants $c > 0$, integer $k \geq 1$, and a suitable environment channel $\mathcal{F}$, and where n and n' are the number of qubits at the input and output of $\mathcal{E}$, respectively. Here, each $\mathcal{N}_i$ and $\mathcal{W}_j$ informally represent local noise channels acting on the input and output qubits, respectively, each of which introduces an error rate bounded by $\varepsilon_0 \leq 2cp$ for some absolute constant c, in the sense specified by Eq. (163) and Eq. (170) below. Therefore, we have

$$\begin{aligned}
\varepsilon_{\mathrm{enc}} &= \frac{1}{2} \sup_\rho \left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho] - \mathcal{E}[\rho] \right\|_1 && (159)\\
&\leq \frac{1}{2} \sup_\rho \left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho] - \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right)[\rho] \right\|_1 \\
&\quad + \frac{1}{2} \sup_\rho \left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right)[\rho] - \mathcal{E}[\rho] \right\|_1 && (160)\\
&\leq \frac{1}{2} \sup_\rho \left\| \widetilde{\mathcal{E}}_{\mathrm{FT}}[\rho] - \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right)[\rho] \right\|_1 \\
&\quad + \frac{1}{2} \sup_\rho \left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right)[\rho] - \mathcal{E}[\rho] \right\|_1 && (161)\\
&\leq 2|\mathcal{E}|(cp)^k + \frac{1}{2} \sup_\rho \left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) \circ (\mathcal{E} \otimes \mathcal{F}) \circ \left(\bigotimes_{i=1}^{n} \mathcal{N}_i \right)[\rho] - \mathcal{E}[\rho] \right\|_1, && (162)
\end{aligned}$$

where first we have used the triangle inequality, then the data processing inequality, and lastly the fact that the 1-norm is not greater than the diamond norm combined with the inequality in Eq. (158). In what follows, we bound the second term by getting rid of the channels $\mathcal{N}_i$ and invoking the QEC guarantee on the code $\mathcal{E}$ to handle the channels $\mathcal{W}_j$.

From Ref. [59, Theorem 59], we have that each $\mathcal{N}_i$ acts on the i-th input qubit and $\mathcal{N}_i : \mathcal{B}(A_i) \to \mathcal{B}(A_i \otimes F_i)$, where A_i and F_i are Hilbert spaces associated with the i-th input qubit and the environment, respectively, and $\mathcal{B}(\cdot)$ denotes linear operators on the Hilbert space. Note that in this notation, $\mathcal{E} : \mathcal{B}(\bigotimes_{i=1}^n A_i) \to \mathcal{B}(\bigotimes_{i=1}^{n'} A_i)$ and $\mathcal{F} : \mathcal{B}(\bigotimes_{i=1}^n F_i) \to \mathcal{B}(\bigotimes_{i=1}^{n'} F_i)$. Furthermore, it holds that

$$\mathrm{Tr}_{F_i} \circ \mathcal{N}_i = (1 - \varepsilon_0)\mathcal{I}_{A_i} + \varepsilon_0 \mathcal{Z}_{A_i}, \tag{163}$$

where $\varepsilon_0 \leq 2cp$ for some absolute constant c, $\mathcal{Z}_{A_i}$ is a CPTP map, and $\mathcal{I}_X$ denotes the identity map on Hilbert space X. Intuitively, this equation tells us that the channel $\mathcal{N}_i$ is close to the identity channel when restricted to the i-th qubit, meaning that its action on the environment qubits must be close to a channel that appends a fixed reference state τ. This can be formalized with the continuity theorem of Stinespring's dilation [130, Theorem 1] as follows. First, consider the Stinespring dilation $\mathcal{N}_i[\cdot] = \mathrm{Tr}_{G_i}\ (V_{A_i \to A_i F_i G_i}[\cdot] V^\dagger_{A_i \to A_i F_i G_i})$, where $V_{A_i \to A_i F_i G_i}$ is an isometry. The continuity theorem states that for two quantum channels $\mathcal{T}_1, \mathcal{T}_2 : \mathcal{B}(A) \to \mathcal{B}(B)$, if V_1 and V_2 are the Stinespring dilating isometries with the same dilating Hilbert space E, then

$$\inf_U \|(\mathbb{I}_B \otimes U_E)V_1 - V_2\|^2 \leq \|\mathcal{T}_1 - \mathcal{T}_2\|_\diamond \leq 2 \inf_U \|(\mathbb{I}_B \otimes U_E)V_1 - V_2\|, \tag{164}$$

where $\mathbb{I}_X$ is the identity operator on Hilbert space X, and $\|\cdot\|$ denotes the operator norm, and the infimum is taken over unitaries U in $\mathcal{B}(E)$. We apply this theorem by taking $A = B = A_i$, $E = F_i G_i$, $\mathcal{T}_1 = \mathrm{Tr}_{F_i} \circ \mathcal{N}_i$, $\mathcal{T}_2 = \mathcal{I}_{A_i}$, $V_1 = V_{A_i \to A_i F_i G_i}$ and V_2 to be an isometry that for all ρ_{A_i} satisfies $V_2 \rho_{A_i} V_2^\dagger = \rho_{A_i} \otimes \tau_{F_i G_i}$ for some fixed state

$\tau_{F_iG_i}$. Thus, the continuity theorem implies the existence of a unitary $U_{F_iG_i}$, such that

$$\|(\mathbb{I}_{A_i}\otimes U_{F_iG_i})V_{A_i\to A_iF_iG_i}-V_2\|\leq\|\operatorname{Tr}_{F_i}\circ\mathcal{N}_i-\mathcal{I}_{A_i}\|_\diamond^{1/2}=\sqrt{2\varepsilon_0}. \tag{165}$$

It follows that for each i we can approximate $\mathcal{N}_i$ with a channel $\mathcal{R}_i$ defined by $\mathcal{R}_i[\rho_{A_i}]=\rho_{A_i}\otimes\sigma_{F_i}$, where $\sigma_{F_i}=\operatorname{Tr}_{G_i}(U_{F_iG_i}^\dagger\tau_{F_iG_i}U_{F_iG_i})$ and satisfies

$$\|\mathcal{N}_i-\mathcal{R}_i\|_\diamond\leq 2\|V_{A_i\to A_iF_iG_i}-(\mathbb{I}_{A_i}\otimes U_{F_iG_i}^\dagger)V_2\|\leq 2\sqrt{2\varepsilon_0}\leq 4\sqrt{cp}. \tag{166}$$

Thus, we obtain

$$\left\|\bigotimes_{i=1}^n\mathcal{N}_i-\bigotimes_{i=1}^n\mathcal{R}_i\right\|_\diamond=\left\|\sum_{i=1}^n\mathcal{R}_1\otimes\cdots\otimes\mathcal{R}_{i-1}\otimes(\mathcal{N}_i-\mathcal{R}_i)\otimes\mathcal{N}_{i+1}\otimes\cdots\otimes\mathcal{N}_n\right\|_\diamond\leq\sum_{i=1}^n\|\mathcal{N}_i-\mathcal{R}_i\|_\diamond\leq 4n\sqrt{cp}, \tag{167}$$

allowing us to get rid of the channels $\mathcal{N}_i$ in Eq. (162). Namely, by invoking the triangle inequality and the fact that the 1-norm is not greater than the diamond norm we obtain

$$\varepsilon_{\text{enc}}\leq\frac{1}{2}\sup_\rho\left\|\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ\left(\bigotimes_{j=1}^{n'}\mathcal{W}_j\right)\circ(\mathcal{E}\otimes\mathcal{F})\circ\bigotimes_{i=1}^n\mathcal{R}_i[\rho]-\mathcal{E}[\rho]\right\|_1+2n\sqrt{cp}+2|\mathcal{E}|(cp)^k, \tag{168}$$

$$=\frac{1}{2}\sup_\rho\left\|\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ\left(\bigotimes_{j=1}^{n'}\mathcal{W}_j\right)[\mathcal{E}[\rho]\otimes\sigma]-\mathcal{E}[\rho]\right\|_1+2n\sqrt{cp}+2|\mathcal{E}|(cp)^k. \tag{169}$$

where we defined a state $\sigma=\mathcal{F}(\otimes_{i=1}^n\sigma_{F_i})$ on the Hilbert space $\otimes_{j=1}^{n'}F_j$.

We now turn our attention to the channels $\mathcal{W}_j$. From Ref. [59, Theorem (59)], we know that each $\mathcal{W}_j$ acts on the j-th output qubit and $\mathcal{W}_j:\mathcal{B}(A_j\otimes F_j)\to\mathcal{B}(A_j)$, where A_j and F_j are Hilbert spaces associated with the j-th output qubit and the environment, respectively. Furthermore, we have

$$\mathcal{W}_j=(1-\varepsilon_0)\mathcal{I}_{A_j}\otimes\operatorname{Tr}_{F_j}+\varepsilon_0\mathcal{Y}_{A_j}, \tag{170}$$

where $\varepsilon_0\leq 2cp$ and c is the same absolute constant from Eq. (163) and $\mathcal{Y}_{A_j}$ is a CPTP map. We would like to use the above form of $\mathcal{W}_j$ and the stochastic-error decoding guarantee of the code $\mathcal{E}$ to bound the first term in Eq. (169). We have

$$\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ\left(\bigotimes_{j=1}^{n'}\mathcal{W}_j\right)[\mathcal{E}[\rho]\otimes\sigma]=\sum_{S\subseteq[n']}(1-\varepsilon_0)^{n'-|S|}\varepsilon_0^{|S|}\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ\left(\bigotimes_{j\notin S}(\mathcal{I}_j\otimes\operatorname{Tr}_{F_j})\bigotimes_{j'\in S}\mathcal{Y}_{A_{j'}}\right)[\mathcal{E}[\rho]\otimes\sigma] \tag{171}$$

$$=\sum_{S\subseteq[n']}(1-\varepsilon_0)^{n'-|S|}\varepsilon_0^{|S|}\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ(\mathcal{I}_{S^c}\otimes\mathcal{W}'_S)[\mathcal{E}[\rho]], \tag{172}$$

where $\mathcal{W}'_S[\cdot]:=\left(\bigotimes_{j'\in S}\mathcal{Y}_{A_{j'}}\right)\left[[\cdot]\otimes(\bigotimes_{j\notin S}\operatorname{Tr}_{F_j})[\sigma]\right]$ is a channel acting on the qubits $S\subseteq[n']$. We now split the above sum into two parts depending whether or not the subset S is correctable. For a correctable set S, we have

$$\mathcal{Q}\circ\widetilde{\mathcal{Q}}\circ(\mathcal{I}_{S^c}\otimes\mathcal{W}'_S)[\mathcal{E}[\rho]]=\mathcal{E}[\rho]. \tag{173}$$

On the other hand, according to the decoding guarantee of the fault-tolerant QEC gadget $\mathcal{Q}_{\text{FT}}$ on the code $\mathcal{E}$, as long as p (and accordingly $\varepsilon_0\leq 2cp$) is sufficiently below a threshold value, the contribution of the uncorrectable sets is bounded follows

$$p_{\text{uncorrectable}}:=\sum_{S\text{ uncorrectable}}(1-\varepsilon_0)^{n'-|S|}\varepsilon_0^{|S|}\leq\Gamma(\mathcal{E}), \tag{174}$$

where $\Gamma(\mathcal{E})$ is a quantity that can be exponentially driven to zero by choosing larger codes from the QEC family, as discussed in Section 3.3 (see Eq. (18)). Hence, we have that

$$\left\| \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ \left(\bigotimes_{j=1}^{n'} \mathcal{W}_j \right) [\mathcal{E}[\rho] \otimes \sigma] - \mathcal{E}[\rho] \right\|_1 \tag{175}$$

$$\leq \left\| \sum_{S \text{ uncorrectable}} (1-\varepsilon_0)^{n'-|S|} \varepsilon_0^{|S|} \mathcal{Q} \circ \widetilde{\mathcal{Q}} \circ (\mathcal{I}_{S^c} \otimes \mathcal{W}'_S)[\mathcal{E}[\rho]] - p_{\text{uncorrectable}} \cdot \mathcal{E}[\rho] \right\|_1 \tag{176}$$

$$\leq 2\Gamma(\mathcal{E}). \tag{177}$$

Collecting terms, we therefore conclude that

$$\varepsilon_{\text{enc}} \leq \Gamma(\mathcal{E}) + 2\sqrt{cp}n + 2|\mathcal{E}|(cp)^k, \tag{178}$$

which completes the proof. □

C Delayed proofs for partial Clifford twirling

First, we state a couple of propositions used in the proofs of the claims in the main text.

Proposition 13. *Suppose $P, P' \in \mathbb{P}$ with $P \neq P'$ and $P \neq -P'$. Let $\mathbb{E}_{C\sim\mathbb{T}}$ denote expectation value over drawing C randomly from $\mathbb{T}$ as in Definition 4. Then*

$$\mathop{\mathbb{E}}_{C\sim\mathbb{T}} CPC^\dagger \otimes CP'C^\dagger = 0 \tag{179}$$

Proof. Denote P and P' in canonical form as $P = \mathrm{i}^{a\cdot b}(-1)^s X^b Z^a$, and $P' = \mathrm{i}^{a'\cdot b'}(-1)^{s'} X^{b'} Z^{a'}$. If $P \neq \pm P'$, then it either must be the case that $a \neq a'$ or $b \neq b'$. Consider a fixed $C \in \mathbb{T}$ defined by (A, B, u, v). We define another $C' \in \mathbb{T}$ for each case separately. If $a \neq a'$ then $a \oplus a' \neq 0^n$ and we can find a w such that $w \cdot (a \oplus a') = 1$ (note that w is actually independent of C). Then we define C' by the choice $(A, B, u \oplus w, v)$. If $b \neq b'$ then we can find a w such that $w \cdot (A^{-1}b \oplus A^{-1}b') = 1$ (w depends on C only through A), and we define C' by $(A, B, u, v \oplus w)$. In both cases, referring to Eq. (191) from the proof of Proposition 4, we have $CPC^\dagger = \pm C'PC'^\dagger$ and $CP'C^\dagger = \mp C'P'C'^\dagger$, or in other words,

$$CPC^\dagger \otimes CP'C^\dagger = -C'PC'^\dagger \otimes C'P'C'^\dagger \,. \tag{180}$$

However, adding w to either u or to v does not change the random distribution over $C \sim \mathbb{T}$. This implies the expectation value is equal to its negation, and therefore it is zero. □

Proposition 14. *Let $P \in \mathbb{P}_{\text{odd}}$ (from Definition 5). Then, for any g,*

$$P|\Psi(g)\rangle\langle\Psi(g)|P|\Psi(g)\rangle = 0 \,. \tag{181}$$

Proof. Write P in canonical form as $\mathrm{i}^{a\cdot b}(-1)^s X^b Z^a$. Since $P \in \mathbb{P}_{\text{odd}}$, we have that $a \cdot b = 1$. Hence, while P is Hermitian, it also satisfies $P = -P^*$, where the * denotes complex conjugation. In other words, all of the matrix elements of P in the computational basis are imaginary. The Hermiticity of P implies that for any state $|\xi\rangle$, we have that $\langle\xi|P|\xi\rangle$ is real, i.e., $\langle\xi|P|\xi\rangle = (\langle\xi|P|\xi\rangle)^*$. However, if the entries of $|\xi\rangle$ in the computational basis are real, then

$$(\langle\xi|P|\xi\rangle)^* = \langle\xi|P^*|\xi\rangle = -\langle\xi|P|\xi\rangle \,, \tag{182}$$

which implies that $\langle\xi|P|\xi\rangle = 0$. It suffices to observe that the entries of $|\Psi(g)\rangle$ are real. □

C.1 Proof of uniform Pauli spreading

Here we restate and prove Proposition 4 from Section 4.3.

Proposition 4 (Restated)**.** *Let $C \sim \mathbb{T}$ denote choosing C randomly from $\mathbb{T}$ as described in Definition 4. Given a fixed $P \in \mathbb{P}$, for any $C \in \mathbb{T}$, $CPC^\dagger \in \mathbb{P}$ since C is Clifford. Furthermore, let $Q \in \mathbb{P}$ be a random variable formed by*

choosing $C \sim \mathbb{T}$ and defining $Q = CPC^\dagger$. Then, the distribution over Q is the uniform distribution over the subset of $\mathbb{P}$ (i.e., $\mathbb{P}_0$, $\mathbb{P}_1$, $\mathbb{P}_Z$, $\mathbb{P}_{\text{even}}$, or $\mathbb{P}_{\text{odd}}$) to which P belongs.

Proof. Recall that choosing $C \sim \mathbb{T}$ means generating uniformly random A, B, u, v and defining $C = Z^v Q_B M_A^\dagger X^u$. Let $P \in \mathbb{P}$ be a fixed Pauli string $P = \mathrm{i}^{a \cdot b}(-1)^s X^b Z^a$. We examine the effect of conjugation by each component of C sequentially. We leave off the i until the end since it is simply a constant.

1. Conjugation by X gates: We have $X^u\big[(-1)^s X^b Z^a\big]X^u = (-1)^{s \oplus (u \cdot a)} X^b Z^a$.
2. Conjugation by CNOT gates: First, we note that for any computational basis state $|x\rangle$, we have

$$M_A^\dagger X^b Z^a M_A |x\rangle = M_A^\dagger X^b Z^a |Ax\rangle = (-1)^{Ax \cdot a} M_A^\dagger |Ax \oplus b\rangle = (-1)^{A^\top a \cdot x} |x \oplus A^{-1}b\rangle \tag{183}$$

$$= (-1)^{A^\top a \cdot x} X^{A^{-1}b} |x\rangle = X^{A^{-1}b} Z^{A^\top a} |x\rangle \tag{184}$$

   From this, we have

$$M_A^\dagger\big[(-1)^{s \oplus (u \cdot a)} Z^a X^b\big] M_A = (-1)^{s \oplus (u \cdot a)} X^{A^{-1}b} Z^{A^\top a} \tag{185}$$

3. Conjugation by CZ: The Q_B gate is diagonal and Hermitian and commutes with Pauli-Z strings. We notice for any $p \in \mathbb{F}_2^n$ and any $|x\rangle$ that

$$Q_B X^p Q_B |x\rangle = (-1)^{x^\top B x} Q_B X^p |x\rangle = (-1)^{x^\top B x} Q_B |x \oplus p\rangle = (-1)^{\left[x^\top B x\right] \oplus \left[(x \oplus p)^\top B (x \oplus p)\right]} |x \oplus p\rangle \tag{186}$$

$$= (-1)^{\left[x^\top B x\right] \oplus \left[(x \oplus p)^\top B (x \oplus p)\right]} X^p |x\rangle = (-1)^{\left[(Bp \oplus B^\top p) \cdot x\right] \oplus \left[p^\top B p\right]} X^p |x\rangle \tag{187}$$

$$= (-1)^{p^\top B p} X^p Z^{(Bp \oplus B^\top p)} |x\rangle \tag{188}$$

   and thus

$$Q_B \Big[(-1)^{s \oplus (u \cdot a)} X^{A^{-1}b} Z^{A^\top a}\Big] Q_B = (-1)^{s \oplus \left[u \cdot a\right] \oplus \left[b^\top A^{-1\top} B A^{-1} b\right]} X^{A^{-1}b} Z^{A^\top a \oplus \left[(B \oplus B^\top) A^{-1} b\right]} \tag{189}$$

4. Conjugation by Z: Finally, the Z^v gate is commuted, which only modifies the sign

$$Z^v \Big[\mathrm{i}^{a \cdot b}(-1)^{s \oplus \left[u \cdot a\right] \oplus \left[b^\top A^{-1\top} B A^{-1} b\right]} X^{A^{-1}b} Z^{A^\top a \oplus \left[(B \oplus B^\top) A^{-1} b\right]}\Big] Z^v \tag{190}$$

$$= \mathrm{i}^{a \cdot b}(-1)^{s \oplus \left[u \cdot a\right] \oplus \left[b^\top A^{-1\top} B A^{-1} b\right] \oplus \left[v \cdot A^{-1} b\right]} X^{A^{-1}b} Z^{A^\top a \oplus \left[(B \oplus B^\top) A^{-1} b\right]} \tag{191}$$

$$= CPC^\dagger \tag{192}$$

Now, we verify that the claims of the uniformity of the distribution over $Q = CPC^\dagger$ when A, B, u, v are chosen uniformly at random. We begin with the easy cases.

- If $P \in \mathbb{P}_0$ or $P \in \mathbb{P}_1$, then the statement is trivial, as the subset has only one element, and substituting $a = b = 0^n$ confirms that $CPC^\dagger = P$.
- If $P \in \mathbb{P}_Z$, then $b = 0^n$ and $a \neq 0^n$. We have

$$CPC^\dagger = \mathrm{i}^{a \cdot b}(-1)^{s \oplus \left[u \cdot a\right]} Z^{A^\top a} \tag{193}$$

  For A a uniformly random invertible binary matrix, $A^\top a$ is a uniformly random nonzero element of $\mathbb{F}_2^n$. Moreover, the uniformly random choice of $u \in \mathbb{F}_2^n$ ensures that the sign is uniformly random. Together, these facts confirm that $CPC^\dagger$ is uniformly random over $\mathbb{P}_Z$.

To handle the final two cases, we now examine the term $(B \oplus B^\top)A^{-1}b$. We have defined B as a random upper triangular matrix. Note, however, that for the matrix $B \oplus B^\top$, one obtains the same distribution regardless of whether B is upper triangular or arbitrary, so long as the entries of B are chosen uniformly at random from $\mathbb{F}_2$. Furthermore, if V is an invertible matrix, then $B \mapsto V^\top B V$ is a bijective map (its inverse is $B \mapsto V^{\top -1} B V^{-1}$). Thus, if B is chosen uniformly at random from all $n \times n$ matrices in $\mathbb{F}_2$ then the distribution over $V^\top B V$ is also the uniform distribution over all matrices in $\mathbb{F}_2$, so we may replace B with $V^\top B V$ and not change the distribution.

For fixed A and nonzero b, we can find an invertible V such that $VA^{-1}b = e$, where $e = (1, 0, 0, \ldots, 0) \in \mathbb{F}_2^n$ is the string with a single 1 in its first entry. For this choice of V, we have

$$(V^\top BV \oplus V^\top B^\top V)A^{-1}b = V^\top (B \oplus B^\top)e \tag{194}$$

We observe that $(B \oplus B^\top)e$ has a zero in its first entry, regardless of the choice of B, but the other $n-1$ entires are uniformly random when B is uniformly random, that is, $(B \oplus B^\top)e$ is distributed uniformly over the subset of $\mathbb{F}_2^n$ orthogonal to e (in the sense of having dot product zero). Letting $c = (B \oplus B^\top)e$, we have $0 = c^\top e = (V^\top c)^\top V^{-1}e = (V^\top c)^\top A^{-1}b$, so $A^{-1}b$ is orthogonal to $V^\top c$. In fact, since c is distributed uniformly over all vectors orthogonal to e and V is bijective, $V^\top c$ is distributed uniformly over all vectors orthogonal to $A^{-1}b$. As noted above, since V is a bijective map, the distribution of $V^\top BV$ is the same as that of B itself, so we can replace $V^\top BV$ with B in Eq. (194), and hence we see that $(B \oplus B^\top)A^{-1}b = V^\top c$. Thus, the quantity $(B \oplus B^\top)A^{-1}b$ is also distributed uniformly over all vectors orthogonal to $A^{-1}b$. We now continue with the final two cases.

- If $P \in \mathbb{P}_{\text{odd}}$, then $a, b \neq 0^n$ (otherwise $a \cdot b = 1$ would be impossible). The quantity $CPC^\dagger$ can be written in canonical form as $\mathrm{i}^{a' \cdot b'}(-1)^{s'} X^{b'} Z^{a'}$, where s', b', a' are given by Eq. (191). First, we immediately observe that randomizing over $u \in \mathbb{F}^n$ leads $u \cdot a$ to be uniformly random, and thus s' is uniformly random over $\mathbb{F}_2$, independent of a' and b'. Since A is distributed randomly over invertible matrices, $b' = A^{-1}b$ is distributed uniformly at random over $\mathbb{F}_2^n \setminus \{0^n\}$. Next, for fixed A (and thus fixed b'), the above reasoning establishes that,

$$a' = A^\top a \oplus \left[(B \oplus B^\top)A^{-1}b\right] = A^\top a \oplus r \tag{195}$$

  where r is distributed uniformly at random over vectors orthogonal to $b' = A^{-1}b$. Recall that $(A^\top a) \cdot (A^{-1}b) = a \cdot b = 1$. This means that

$$a' \cdot b' = \left[(A^\top a) \cdot (A^{-1}b)\right] \oplus \left[r \cdot A^{-1}b\right] = a \cdot b = 1 \tag{196}$$

  Since r is distributed at random among all vectors for which $b' \cdot r = 0$, the distribution over a' is precisely the uniform distribution over all vectors for which $a' \cdot b' = 1$. This confirms that $CPC^\dagger$ is distributed uniformly over $\mathbb{P}_{\text{odd}}$.

- If $P \in \mathbb{P}_{\text{even}}$, then $b \neq 0^n$ and $a \cdot b = 0$. As before, we have $CPC^\dagger = \mathrm{i}^{a' \cdot b'}(-1)^{s'} X^{b'} Z^{a'}$, and $b' = A^{-1}b$ is distributed uniformly at random over $\mathbb{F}_2^n \setminus \{0^n\}$. Furthermore, since $v \in \mathbb{F}_2^n$ is uniformly and indepenently chosen, the quantity $v \cdot (A^{-1}b)$ is uniformly random, meaning s' is uniformly random and independent of a' and b'. Next, we can again write

$$a' = A^\top a \oplus r \tag{197}$$

  where r is orthogonal to b', and $A^\top a$ is a fixed vector for which $b' \cdot (A^\top a) = b \cdot a = 0$. Thus, the distribution over a' is uniform over vectors orthogonal to b', confirming that $CPC^\dagger$ is uniform over the set $\mathbb{P}_{\text{even}}$.

□

C.2 Proof of correct top eigenvector

Here we restate and prove Proposition 5 from Section 4.3.

Proposition 5 (Restated)**.** *Suppose that for every g, the state $\overline{\phi(g)}$ defined in Eq. (25) satisfies $\langle\overline{\Psi(g)}|\overline{\phi(g)}|\overline{\Psi(g)}\rangle \geq F_{\min}$, and suppose that the faulty QRAM device is subject to dataset-independent noise (Definition 1). Let $C \sim \mathbb{T}$ denote drawing C randomly from the twirling set as described in Definition 4, and let $\mathbb{E}$ denote expectation value. For each g, let $\overline{\phi(g)}_{\text{twirl}}$ be defined as in Eq. (43). Then $\overline{\phi(g)}_{\text{twirl}}$ satisfies the eigenvalue equation*

$$\overline{\phi(g)}_{\text{twirl}}|\overline{\Psi(g)}\rangle = \lambda_{\text{twirl}}|\overline{\Psi(g)}\rangle\,, \tag{44}$$

with $\lambda_{\text{twirl}} \geq F_{\min}$. Furthermore, all other eigenvalues of $\overline{\phi(g)}_{\text{twirl}}$ are no larger than 2^{-n+1}.

Proof. From the definition of dataset-independent noise (Definition 1) and the definiton of $\overline{\phi(g)}$ in Eq. (25), we have that

$$\overline{\phi(g)} = \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\text{FT}} \circ \widetilde{\mathcal{E}}_{\text{FT}} \circ \mathcal{N}_2 \circ \mathcal{V}(g) \circ \mathcal{N}_1[|+\rangle\langle+|^{\otimes n}]\,. \tag{198}$$

Since the ideal encdoing map $\mathcal{E}$ is injective and $\mathcal{Q}$ projects onto its image (i.e., the codespace), we may define the channel $\mathcal{N}_{\text{enc}} = \mathcal{E}^{-1} \circ \mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\text{FT}} \circ \widetilde{\mathcal{E}}_{\text{FT}}$ (where $\mathcal{E}^{-1}$ is well defined on inputs in the codespace) and note that

$$\mathcal{Q} \circ \widetilde{\mathcal{Q}}_{\text{FT}} \circ \widetilde{\mathcal{E}}_{\text{FT}} = \mathcal{E} \circ \mathcal{N}_{\text{enc}} \,. \tag{199}$$

Thus, we have

$$\overline{\phi(g)} = \mathcal{E}\Big[\mathcal{N}_{\text{enc}}\Big[\mathcal{N}_2\Big[V(g)\,\mathcal{N}_1\Big[|+\rangle\langle+|^{\otimes n}\Big]V(g)\Big]\Big]\Big] \tag{200}$$

The 2^n states $\{Z^u|+\rangle, u \in \{0,1\}^n\}$ form an orthonormal basis, and we may write the state $\mathcal{N}_1\Big[|+\rangle\langle+|^{\otimes n}\Big]$ as a fixed g-independent density matrix in this basis

$$\mathcal{N}_1\Big[|+\rangle\langle+|^{\otimes n}\Big] = \sum_{u,v\in\{0,1\}^n} M_{uv} Z^u |+\rangle\langle+|^{\otimes n} Z^v \tag{201}$$

The operation $V(g)$ is diagonal and commutes with Z^u and Z^v, so we have

$$V(g)\Big[\mathcal{N}_1\Big[|+\rangle\langle+|^{\otimes n}\Big]\Big]V(g) = \sum_{u,v\in\{0,1\}^n} M_{uv} Z^u |\Psi(g)\rangle\langle\Psi(g)| Z^v \tag{202}$$

The channel $\mathcal{N}_{\text{enc}} \circ \mathcal{N}_2$ in general has a Pauli decomposition

$$\mathcal{N}_{\text{enc}} \circ \mathcal{N}_2[\sigma] = \sum_{P,P'\in\mathbb{P}} \chi_{PP'} P\sigma P' \tag{203}$$

where χ is some noise matrix, and $\mathbb{P}$ denotes the set of signed Pauli strings, as in Definition 5. Moreover, $\mathcal{E}$ is the physical-to-logical mapping which is equivalent to putting an overline on all the objects. Thus, we have

$$\overline{\phi(g)} = \sum_{P,P'\in\mathbb{P}} \sum_{u,v\in\{0,1\}^n} \chi_{PP'} M_{uv} \overline{P}\,\overline{Z}^u |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| \overline{Z}^v\,\overline{P}' \,. \tag{204}$$

Now we compute the expectation over the twirled state, recalling the equality from Proposition 3.

$$\begin{aligned}
\overline{\phi(g)}_{\text{twirl}} &= \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \sum_{P,P'\in\mathbb{P}} \sum_{u,v\in\{0,1\}^n} \chi_{PP'} M_{uv} \overline{C}\,\overline{P}\,\overline{Z}^u |\overline{\Psi(g_C)}\rangle\langle\overline{\Psi(g_C)}| \overline{Z}^v\,\overline{P}'\,\overline{C}^\dagger && (205)\\
&= \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \sum_{P,P'\in\mathbb{P}} \sum_{u,v\in\{0,1\}^n} \chi_{PP'} M_{uv} \overline{C}\,\overline{P}\,\overline{Z}^u\,\overline{C}^\dagger |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| \overline{C}\,\overline{Z}^v\,\overline{P}'\,\overline{C}^\dagger && (206)\\
&= \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \sum_{P,P'\in\mathbb{P}} \chi'_{PP'} \overline{C}\,\overline{P}\,\overline{C}^\dagger |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| \overline{C}\,\overline{P}'\,\overline{C}^\dagger && (207)
\end{aligned}$$

where we have absorbed the $\overline{Z}^u$ and $\overline{Z}^v$ into $\overline{P}$ and $\overline{P'}$, leading to a redefinition of $\chi_{PP'}$ to $\chi'_{PP'}$. We may move the expectation value inside the sum. The result of Proposition 13 ensures that if $P \neq P'$, the expectation value vanishes, leaving only diagonal terms where $P = P'$ (if $P = -P'$ then the term does not immediately vanish, but without loss of generality we may take $\chi'_{-P,P} = 0$ for all P by appropriately redefining the value of $\chi'_{P,P}$). This gives

$$\overline{\phi(g)}_{\text{twirl}} = \sum_{P\in\mathbb{P}} \chi'_{PP} \mathop{\mathbb{E}}_{C\sim\mathbb{T}} \overline{C}\,\overline{P}\,\overline{C}^\dagger |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| \overline{C}\,\overline{P}\,\overline{C}^\dagger \,. \tag{208}$$

Recall from Definition 5 that we can partition the set $\mathbb{P}$ into sets $\mathbb{P}_0$, $\mathbb{P}_1$, $\mathbb{P}_Z$, $\mathbb{P}_{\text{even}}$, and $\mathbb{P}_{\text{odd}}$. Note that $\mathbb{P}_0$ and $\mathbb{P}_1$ each contain a single element, which differs by a sign that cancels out in the expression above. Without loss of generality, we may assume that $\chi'_{PP} = 0$ for $P = -\mathbb{I}$ and ignore the set $\mathbb{P}_1$ henceforth. The fact that $\overline{\phi(g)}_{\text{twirl}}$ has trace 1 means that $\sum_{P\in\mathbb{P}} \chi'_{PP} = 1$; we define χ'_0, χ'_Z, χ'_{even}, and χ'_{odd} as the sum of χ'_{PP} over P from the relevant set.

The result of Proposition 4 states that for each P, $CPC^\dagger$ is distributed uniformly over the subset of $\mathbb{P}$ that contains P. Thus, for all P in the same subset the quantity $\mathbb{E}_{C\sim\mathbb{T}}\, \overline{C}\,\overline{P}\,\overline{C}^\dagger |\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| \overline{C}\,\overline{P}\,\overline{C}^\dagger$ is the same. We may

thus define

$$\overline{\rho(g)}_Z = \mathop{\mathbb{E}}_{P\sim\mathbb{P}_Z} \overline{P}|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{P} \tag{209}$$

$$\overline{\rho(g)}_{\text{even}} = \mathop{\mathbb{E}}_{P\sim\mathbb{P}_{\text{even}}} \overline{P}|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{P} \tag{210}$$

$$\overline{\rho(g)}_{\text{odd}} = \mathop{\mathbb{E}}_{P\sim\mathbb{P}_{\text{odd}}} \overline{P}|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{P} \tag{211}$$

and write

$$\overline{\phi(g)}_{\text{twirl}} = \chi'_0|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| + \big[\chi'_Z\overline{\rho(g)}_Z\big] + \big[\chi'_{\text{even}}\overline{\rho(g)}_{\text{even}}\big] + \big[\chi'_{\text{odd}}\overline{\rho(g)}_{\text{odd}}\big] \tag{212}$$

Next, we explicitly compute the term of Z-like Paulis.

$$\overline{\rho(g)}_Z = \frac{1}{2^n-1}\sum_{0^n\neq u\in\{0,1\}^n} \overline{Z}^u|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|\overline{Z}^u = \frac{\overline{\mathbb{I}}}{2^n-1} - \frac{|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|}{2^n-1}\,, \tag{213}$$

which follows since the set of $\overline{Z}^u|\overline{\Psi(g)}\rangle$ form an orthonormal basis for $u\in\{0,1\}^n$, and sampling a random vector from an orthonormal basis yields the maximally mixed state.

Additionally, we can use the fact that the Pauli matrices form a 1-design [129] to say that if we draw a P uniformly at random from all of $\mathbb{P}$, conjugating an arbitrary state $|\overline{\xi}\rangle\langle\overline{\xi}|$ by P yields the maximally mixed state after averaging over P. Weighting the different subsets by their sizes (e.g., a $(2^n-1)/2^{2n}$ fraction of all signed Pauli strings are in $\mathbb{P}_Z$), this fact is equivalent to

$$\frac{1}{2^{2n}}|\overline{\xi}\rangle\langle\overline{\xi}| + \frac{2^n-1}{2^{2n}}\mathop{\mathbb{E}}_{P\in\mathbb{P}_Z}\overline{P}|\overline{\xi}\rangle\langle\overline{\xi}|\overline{P} + \frac{2^n-1}{2^{n+1}}\mathop{\mathbb{E}}_{P\in\mathbb{P}_{\text{even}}}\overline{P}|\overline{\xi}\rangle\langle\overline{\xi}|\overline{P} + \frac{2^n-1}{2^{n+1}}\mathop{\mathbb{E}}_{P\in\mathbb{P}_{\text{odd}}}\overline{P}|\overline{\xi}\rangle\langle\overline{\xi}|\overline{P} = \frac{\overline{\mathbb{I}}}{2^n} \tag{214}$$

We apply this facts to the state $|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}|$, substitute Eq. (213), and rearrange to write

$$\overline{\rho(g)}_{\text{even}} = 2\frac{\overline{\mathbb{I}}}{2^n} - \overline{\rho(g)}_{\text{odd}} \tag{215}$$

Next, we plug this into Eq. (212) to get rid of $\rho(g)_{\text{even}}$ and say

$$\overline{\phi(g)}_{\text{twirl}} = (\chi'_0 - \frac{\chi'_Z}{2^n-1})|\overline{\Psi(g)}\rangle\langle\overline{\Psi(g)}| + \big[\big(\frac{2^n}{2^n-1}\chi'_Z + 2\chi'_{\text{even}}\big)\frac{\mathbb{I}}{2^n}\big] + \big[(\chi'_{\text{odd}} - \chi'_{\text{even}})\overline{\rho(g)}_{\text{odd}}\big] \tag{216}$$

We are now ready to conclude. A consequence of Proposition 14 is that $\overline{\rho(g)}_{\text{odd}}|\overline{\Psi(g)}\rangle = 0$. Thus, we have

$$\overline{\phi(g)}_{\text{twirl}}|\overline{\Psi(g)}\rangle = \lambda_{\text{twirl}}|\overline{\Psi(g)}\rangle \tag{217}$$

with $\lambda_{\text{twirl}} = \chi'_0 + 2^{-n+1}\chi'_{\text{even}}$. From the definition of $\overline{\phi(g)}_{\text{twirl}}$ and the relation in Proposition 3, we have

$$\lambda_{\text{twirl}} = \langle\overline{\Psi(g)}|\overline{\phi(g)}_{\text{twirl}}|\overline{\Psi(g)}\rangle = \mathop{\mathbb{E}}_{C\sim\mathbb{T}}\langle\overline{\Psi(g)}|\overline{C}\,\overline{\phi(g_C)}\,\overline{C}^\dagger|\overline{\Psi(g)}\rangle = \mathop{\mathbb{E}}_{C\sim\mathbb{T}}\langle\overline{\Psi(g_C)}|\overline{\phi(g_C)}^\dagger|\overline{\Psi(g_C)}\rangle \geq F_{\min}\,. \tag{218}$$

Finally, we reason about the other eigenvalues of $\overline{\phi(g)}_{\text{twirl}}$. We may note from their definitions in Eq. (210) and Eq. (211) that $\overline{\rho(g)}_{\text{even}}$ and $\overline{\rho(g)}_{\text{odd}}$ are positive semidefinite operators. From Eq. (215), this means the eigenvalues of $\overline{\rho(g)}_{\text{odd}}$ cannot exceed 2^{-n+1}. We now turn back to Eq. (216) to bound any other eigenvalue λ_{other} of an eigenvector $|\overline{\lambda_{\text{other}}}\rangle$ of $\overline{\phi(g)}_{\text{twirl}}$. We have

$$\overline{\phi(g)}_{\text{twirl}}|\overline{\lambda_{\text{other}}}\rangle = \big[\big(\frac{2^n}{2^n-1}\chi'_Z + 2\chi'_{\text{even}}\big)\frac{1}{2^n}\big]|\overline{\lambda_{\text{other}}}\rangle + \big[(\chi'_{\text{odd}} - \chi'_{\text{even}})\overline{\rho(g)}_{\text{odd}}\big]|\overline{\lambda_{\text{other}}}\rangle\,. \tag{219}$$

Thus, $|\overline{\lambda_{\text{other}}}\rangle$ must also be an eigenvector of $\overline{\rho(g)}_{\text{odd}}$; assume the associated eigenvalue is $C \geq 0$. We have,

$$\lambda_{\text{other}} = \left(\frac{2^n}{2^n-1}\chi'_Z + 2\chi'_{\text{even}}\right)\frac{1}{2^n} + (\chi'_{\text{odd}} - \chi'_{\text{even}})C \tag{220}$$

$$= \frac{1}{2^n-1}\chi'_Z + \chi'_{\text{even}}\left(\frac{1}{2^{n-1}} - C\right) + \chi'_{\text{odd}}C \tag{221}$$

Suppose $\chi'_{\text{even}} \geq \chi'_{\text{odd}}$. In that case, the last term in Eq. (220) is negative (or zero), so we may bound

$$\lambda_{\text{other}} \leq \frac{1}{2^n-1}\chi'_Z + \frac{2}{2^n}\chi'_{\text{even}}. \tag{222}$$

On the other hand, suppose $\chi'_{\text{even}} < \chi'_{\text{odd}}$. In that case, we have that

$$\lambda_{\text{other}} \leq \frac{1}{2^n-1}\chi'_Z + \chi'_{\text{odd}}\left(\frac{1}{2^{n-1}} - C\right) + C\chi'_{\text{odd}} \tag{223}$$

$$= \frac{1}{2^n-1}\chi'_Z + \chi'_{\text{odd}}\frac{1}{2^{n-1}}. \tag{224}$$

Thus, we can combine the expressions and conclude that

$$\lambda_{\text{other}} \leq \frac{1}{2^n-1}\chi'_Z + \frac{1}{2^{n-1}}\max(\chi'_{\text{even}}, \chi'_{\text{odd}}). \tag{225}$$

Finally, since $\chi'_0 + \chi'_Z + \chi'_{\text{even}} + \chi'_{\text{odd}} = 1$ with each term in the sum non-negative, it immediately follows that no eigenvalue λ_{other} can be larger than 2^{-n+1}. □

D QRAM is in the Clifford hierarchy

Here we show that for any f, the n-qubit QRAM gate $\overline{V(f)}$ from Eq. (1) is in the n-th level of the logical Clifford hierarchy, $\mathcal{C}_n$ from Eq. (8), and can therefore be teleported via the strategy sketched in Section 2.2. This can be seen directly as a consequence of Ref. [46, Theorem 3], but here we give a self-contained proof.

We view the classical dataset f as an n-bit Boolean function $f\colon \{0,1\}^n \to \{0,1\}$, which we expand as a polynomal of its input bits over the field $\mathbb{F}_2$ (i.e., addition and multiplication are done modulo 2). We consider all possible 2^n monomials, denoted by x^e, where $e \in \{0,1\}^n$ dictates the exponents of each bit

$$x^e = x_1^{e_1} x_2^{e_2} \cdots x_n^{e_n}. \tag{226}$$

Then, we may uniquely write

$$f(x_1, \ldots, x_n) = \bigoplus_{e\in\{0,1\}^n} c_e x^e \tag{227}$$

where $c_e \in \{0,1\}$ is the binary coefficient for the monomial x^e, and $\oplus$ denotes addition modulo 2. The degree of f is then given by the maximum degree of any monomial that appears in this expansion, that is,

$$\deg(f) = \max_{e\in\{0,1\}^n} c_e|e|, \tag{228}$$

where $|e| = \sum_{i=1}^n e_i$ is the Hamming weight of e.

We will prove by induction that $\overline{V(f)} \in \mathcal{C}_{\deg(f)}$. As the base case, we observe that if $\deg(f) = 1$, then there is a bit string $u \in \{0,1\}^n$ for which $f(x) = \bigoplus_{i=1}^n u_i x_i$. Thus, we have that $\overline{V(f)} = \overline{Z}^u$, where $\overline{Z}^u$ denotes the (logical) Pauli operator with $\overline{Z}$ in positions where $u_i = 1$ and $\overline{\mathbb{I}}$ in other positions. That is, for degree-1 functions f, the QRAM operation is a Pauli operator, $\overline{V(f)} \in \mathcal{C}_1$.

Next, we assume for induction that for any degree-$(d-1)$ function g, the gate $\overline{V(g)} \in \mathcal{C}_{d-1}$. We consider a degree-$d$ function f, and we would like to show that this implies $\overline{V(f)} \in \mathcal{C}_d$. Note that $\overline{V(f)} = \overline{V(f)}^\dagger$, and that the

family of QRAM operations obeys the general composition rule

$$\overline{V(f)}\,\overline{V(h)} = \overline{V(f \oplus h)} \tag{229}$$

for any pair of n-bit Boolean functions f, h, where $f \oplus h$ is the function for which $(f \oplus h)(x) = f(x) \oplus h(x)$. Thus, based on the decomposition in Eq. (227), we have

$$\overline{V(f)} = \prod_{e \in \{0,1\}^n} \overline{V(x^e)}^{c_e} \tag{230}$$

Consider a factor $\overline{V(x^e)}$ associated with a degree-$|e|$ monomial x^e. For any X-type Pauli string X^m, we compute

$$\overline{X}^m \overline{V(x^e)} = \overline{V((x \oplus m)^e)}\,\overline{X}^m\,, \tag{231}$$

and thus

$$\overline{V(x^e)}\,\overline{X}^m \overline{V(x^e)} = \overline{V(x^e)}\,\overline{V((x \oplus m)^e)}\,\overline{X}^m = \overline{V(x^e \oplus (x \oplus m)^e)}\,\overline{X}^m\,. \tag{232}$$

The crucial observation is that for any fixed m, the Boolean function $x^e \oplus (x \oplus m)^e$, viewed as a function of $x_1, \ldots, x_n$, has degree at most $|e| - 1$. We see this by inspection of the expression

$$x^e \oplus (x \oplus m)^e = \left[x_1^{e_1} \cdots x_n^{e_n}\right] \oplus \left[(x_1 \oplus m_1)^{e_1} (x_2 \oplus m_2)^{e_2} \cdots (x_n \oplus m_n)^{e_n}\right], \tag{233}$$

noting that distributing the multiplication of the second term gives $2^{|e|}$ terms, one of which has degree $|e|$ and will cancel the first term. The rest of the terms have degree $|e| - 1$ or lower. Since f is a sum of monomials as in Eq. (227) and $V(f)$ decomposes into commuting diagonal factors as $V(f) = \prod_{e \in \{0,1\}^n} \overline{V(x^e)}^{c_e}$ from Eq. (230), we have

$$\overline{V(f)}\,\overline{X}^m\,\overline{V(f)} = \left[\prod_{e \in \{0,1\}^n} \overline{V(x^e \oplus (x \oplus m)^e)}^{c_e}\right] \overline{X}^m = \overline{V(g)}\,\overline{X}^m\,, \tag{234}$$

where g is a Boolean function satisfying $g(x) = \bigoplus_{e \in \{0,1\}^n} c_e (x^e \oplus (x \oplus m)^e)$. Since f has degree d, g is the sum of degree $d-1$ functions, and thus has degree at most $d-1$. By the inductive assumption $\overline{V(g)} \in \mathcal{C}_{d-1}$, which implies that $\overline{V(g)}\,\overline{X}^m \in \mathcal{C}_{d-1}$ as well. Furthermore, since $\overline{V(f)}$ is diagonal, it is straightforward to see that $\overline{V(f)}\,\overline{Z}^m\,\overline{V(f)} = \overline{Z}^m \in \mathcal{C}_{d-1}$ for any m. Finally, since any Pauli string $\overline{P}$ is proportional to $\overline{X}^a\,\overline{Z}^b$ for some $a, b \in \{0,1\}^n$, these facts together imply $\overline{V(f)}\,\overline{X}^a\,\overline{Z}^b\,\overline{V(f)} = \overline{V(g)}\,\overline{X}^m\,\overline{Z}^b \in \mathcal{C}_{d-1}$; thus, $\overline{V(f)}\,\overline{P}\,\overline{V(f)} \in \mathcal{C}_{d-1}$ for any Pauli string $\overline{P}$. By induction, we conclude that $\overline{V(f)} \in \mathcal{C}_{\deg(f)}$ for all f. Since the maximum degree of any f is n, we have $\overline{V(f)} \in \mathcal{C}_n$.

References

[1] Williams, F. C. and Kilburn, T. "Electronic digital computers." *Nature* **162** (1948), 487–487.

[2] C., W. F. and T., K. "A storage system for use with binary-digital computing machines." *Proceedings of the IEE—part II: power engineering* **96** (1949), 183–200.

[3] Engineering and Technology History Wiki. *Milestones: Manchester University "Baby"' computer and its derivatives, 1948–1951.* https://ethw.org/Milestones:Manchester_University_%22Baby%22_Computer_and_its_Derivatives,_1948-1951, accessed 2025-03-17. (2022).

[4] Kim, S., Hooper, C., Wattanawong, T., Kang, M., Yan, R., Genc, H., Dinh, G., Huang, Q., Keutzer, K., Mahoney, M. W., Shao, S., and Gholami, A. "Full stack optimization of transformer inference." In: *Architecture and System Support for Transformer Models (ASSYST)* (2023). https://openreview.net/forum?id=GtyQbLUUagE, accessed: 2025-03-17. arXiv:2302.14017.

[5] Silvano, C., Ielmini, D., Ferrandi, F., et al. "A survey on deep learning hardware accelerators for heterogeneous HPC platforms." *ACM Comput. Surv.* (2025). arXiv:2306.15552.

[6] Infineon. *CY7C1069G-10BVXIT.* https://www.infineon.com/cms/en/product/memories/sram-static-ram/asynchronous-sram/cy7c1069g-10bvxit/, accessed: 2025-03-17. (2025).

[7] Giovannetti, V., Lloyd, S., and Maccone, L. "Quantum random access memory." *Phys. Rev. Lett.* **100** (2008), 160501. arXiv:0708.1879.

[8] Jaques, S. and Rattew, A. G. "QRAM: A survey and critique." arXiv:2305.10310 (2023).
[9] Ciliberto, C., Herbster, M., Ialongo, A. D., Pontil, M., Rocchetto, A., Severini, S., and Wossnig, L. "Quantum machine learning: A classical perspective." *Proc. R. Soc. A* **474** (2018), 20170551. arXiv:1707.08561.
[10] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. "Quantum machine learning." *Nature* **549** (2017), 195–202. arXiv:1611.09347.
[11] Dalzell, A. M., McArdle, S., Berta, M., Bienias, P., Chen, C.-F., Gilyén, A., Hann, C. T., Kastoryano, M. J., Khabiboulline, E. T., Kubica, A., Salton, G., Wang, S., and Brandão, F. G. S. L. *Quantum algorithms: A survey of applications and end-to-end complexities*. Cambridge University Press (2025). arXiv:2310.03011.
[12] Rebentrost, P., Mohseni, M., and Lloyd, S. "Quantum support vector machine for big data classification." *Phys. Rev. Lett.* **113** (2014), 130503. arXiv:1307.0471.
[13] Zhao, Z., Fitzsimons, J. K., and Fitzsimons, J. F. "Quantum-assisted Gaussian process regression." *Phys. Rev. A* **99** (2019), 052331. arXiv:1512.03929.
[14] Kerenidis, I. and Prakash, A. "Quantum recommendation systems." In: *ITCS* (2017), 49:1–49:21. arXiv:1603.08675.
[15] Berry, D. W. "High-order quantum algorithm for solving linear differential equations." *J. Phys. A* **47** (2014), 105301. arXiv:1010.2745.
[16] Berry, D. W., Childs, A. M., Ostrander, A., and Wang, G. "Quantum algorithm for linear differential equations with exponentially improved dependence on precision." *Commun. Math. Phys.* **356** (2017), 1057–1081. arXiv:1701.03684.
[17] Childs, A. M. and Liu, J.-P. "Quantum spectral methods for differential equations." *Commun. Math. Phys.* **375** (2020), 1427–1457. arXiv:1901.00961.
[18] Krovi, H. "Improved quantum algorithms for linear and nonlinear differential equations." *Quantum* **7** (2023), 913. arXiv:2202.01054.
[19] Jennings, D., Lostaglio, M., Lowrie, R. B., Pallister, S., and Sornborger, A. T. "The cost of solving linear differential equations on a quantum computer: Fast-forwarding to explicit resource counts." *Quantum* **8** (2024), 1553. arXiv:2309.07881.
[20] Berry, D. W. and C. S. Costa, P. "Quantum algorithm for time-dependent differential equations using Dyson series." *Quantum* **8** (2024), 1369. arXiv:2212.03544.
[21] Harrow, A. W., Hassidim, A., and Lloyd, S. "Quantum algorithm for linear systems of equations." *Phys. Rev. Lett.* **103** (2009), 150502. arXiv:0811.3171.
[22] Brandão, F. G. S. L. and Svore, K. M. "Quantum speed-ups for solving semidefinite programs." In: *FOCS* (2017), 415–426. arXiv:1609.05537.
[23] Brandão, F. G. S. L., Kalev, A., Li, T., Lin, C. Y.-Y., Svore, K. M., and Wu, X. "Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning." In: *ICALP* (2019), 27:1–27:14. arXiv:1710.02581.
[24] van Apeldoorn, J., Gilyén, A., Gribling, S., and de Wolf, R. "Quantum SDP-solvers: Better upper and lower bounds." *Quantum* **4** (2020), 230. Earlier version in *FOCS'17*. arXiv:1705.01843.
[25] van Apeldoorn, J. and Gilyén, A. "Quantum algorithms for zero-sum games." arXiv:1904.03180 (2019).
[26] van Apeldoorn, J. and Gilyén, A. "Improvements in quantum SDP-solving with applications." In: *ICALP* (2019), 99:1–99:15. arXiv:1804.05058.
[27] Kerenidis, I. and Prakash, A. "A quantum interior point method for LPs and SDPs." *ACM Trans. Quantum Comput.* **1** (2020). arXiv:1808.09266.
[28] Kerenidis, I., Prakash, A., and Szilágyi, D. "Quantum algorithms for second-order cone programming and support vector machines." *Quantum* **5** (2021), 427. arXiv:1908.06720.
[29] Augustino, B., Nannicini, G., Terlaky, T., and Zuluaga, L. F. "Quantum interior point methods for semidefinite optimization." *Quantum* **7** (2023), 1110. arXiv:2112.06025.
[30] Kerenidis, I., Prakash, A., and Szilágyi, D. "Quantum algorithms for portfolio optimization." In: *AFT* (2019), 147–155. arXiv:1908.08040.
[31] Dalzell, A. M., Clader, B. D., Salton, G., Berta, M., Lin, C. Y.-Y., Bader, D. A., Stamatopoulos, N., Schuetz, M. J. A., Brandão, F. G. S. L., Katzgraber, H. G., and Zeng, W. J. "End-to-end resource analysis for quantum interior-point methods and portfolio optimization." *PRX Quantum* **4** (2023), 040325. arXiv:2211.12489.
[32] Arunachalam, S., Gheorghiu, V., Jochym-O'Connor, T., Mosca, M., and Srinivasan, P. V. "On the robustness of bucket brigade quantum RAM." *New J. Phys.* **17** (2015), 123010. arXiv:1502.03450.
[33] Steiger, D. S. and Troyer, M. "Racing in parallel: Quantum versus classical." In: *APS March Meeting Abstracts* (2016), H44–010.

[34] Di Matteo, O., Gheorghiu, V., and Mosca, M. "Fault-tolerant resource estimation of quantum random-access memories." *IEEE Trans. Quantum Eng.* **1** (2020), 1–13. arXiv:1902.01329.
[35] Paler, A., Oumarou, O., and Basmadjian, R. "Parallelizing the queries in a bucket-brigade quantum random access memory." *Phys. Rev. A* **102** (2020), 032608. arXiv:2002.09340.
[36] Hann, C. T., Lee, G., Girvin, S., and Jiang, L. "Resilience of quantum random access memory to generic noise." *PRX Quantum* **2** (2021), 020311. arXiv:2012.05340.
[37] Mukhopadhyay, P. "A quantum random access memory (qram) using a polynomial encoding of binary strings." *Sci. Rep.* **15** (2025), 11002. arXiv:2408.16794.
[38] Low, G. H., Kliuchnikov, V., and Schaeffer, L. "Trading T gates for dirty qubits in state preparation and unitary synthesis." *Quantum* **8** (2024), 1375. arXiv:1812.00954.
[39] Jaeger, R. C. and Blalock, T. N. *Microelectronic circuit design*. McGraw-Hill New York (2016).
[40] Wang, Y., Alexeev, Y., Jiang, L., Chong, F. T., and Liu, J. "Fundamental causal bounds of quantum random access memories." *npj Quant. Inf.* **10** (2024), 71. arXiv:2307.13460.
[41] Zeng, B., Chen, X., and Chuang, I. L. "Semi-Clifford operations, structure of $\mathcal{C}_k$ hierarchy, and gate complexity for fault-tolerant quantum computation." *Phys. Rev. A* **77** (2008), 042313. arXiv:0712.2084.
[42] Hastings, M. B. and Haah, J. "Distillation with sublogarithmic overhead." *Phys. Rev. Lett.* **120** (2018), 050504. arXiv:1709.03543.
[43] Kubica, A. and Beverland, M. E. "Universal transversal gates with color codes: A simplified approach." *Phys. Rev. A* **91** (2015), 032330. arXiv:1410.0069.
[44] Koutsioumpas, S., Banfield, D., and Kay, A. "The smallest code with transversal T." arXiv:2210.14066 (2022).
[45] Kuperberg, G. "Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem." In: *TQC* (2013), 20–34. arXiv:1112.3333.
[46] Cui, S. X., Gottesman, D., and Krishna, A. "Diagonal gates in the Clifford hierarchy." *Phys. Rev. A* **95** (2017), 012329. arXiv:1608.06596.
[47] Horsman, D., Fowler, A. G., Devitt, S., and Meter, R. V. "Surface code quantum computing by lattice surgery." *New J. Phys.* **14** (2012), 123011. arXiv:1111.4022.
[48] Li, Y. "A magic state's fidelity can be superior to the operations that created it." *New J. Phys.* **17** (2015), 023037. arXiv:1410.7808.
[49] Łodyga, J., Mazurek, P., Grudka, A., and Horodecki, M. "Simple scheme for encoding and decoding a qubit in unknown state for various topological codes." *Sci. Rep.* **5** (2015), 8975. arXiv:1404.2495.
[50] Litinski, D. "Magic state distillation: Not as costly as you think." *Quantum* **3** (2019), 205. arXiv:1905.06903.
[51] Bravyi, S. and Kitaev, A. "Universal quantum computation with ideal Clifford gates and noisy ancillas." *Phys. Rev. A* **71** (2005), 022316. arXiv:quant-ph/0403025.
[52] Knill, E. "Fault-tolerant postselected quantum computation: Schemes." arXiv:quant-ph/0402171 (2004).
[53] Bombín, H. and Martin-Delgado, M. A. "Topological quantum distillation." *Phys. Rev. Lett.* **97** (2006), 180501. arXiv:quant-ph/0605138.
[54] Kubica, A., Yoshida, B., and Pastawski, F. "Unfolding the color code." *New J. Phys.* **17** (2015), 083026. arXiv:1503.02065.
[55] Moussa, J. E. "Transversal Clifford gates on folded surface codes." *Phys. Rev. A* **94** (2016), 042316. arXiv:1603.02286.
[56] Litinski, D. "A game of surface codes: large-scale quantum computing with lattice surgery." *Quantum* **3** (2019), 128. arXiv:1808.02892.
[57] Gidney, C. and Fowler, A. G. "Flexible layout of surface code computations using AutoCCZ states." arXiv:1905.08916 (2019).
[58] Gottesman, D. and Chuang, I. L. "Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations." *Nature* **402** (1999), 390–393. arXiv:quant-ph/9908010.
[59] Christandl, M., Fawzi, O., and Goswami, A. "Fault-tolerant quantum input/output." arXiv:2408.05260 (2024).
[60] Wills, A., Hsieh, M.-H., and Yamasaki, H. "Constant-overhead magic state distillation." arXiv:2408.07764 (2024).
[61] Nguyen, Q. T. "Good binary quantum codes with transversal ccz gate." arXiv:2408.10140 (2024).
[62] Golowich, L. and Guruswami, V. "Asymptotically good quantum codes with transversal non-Clifford gates." arXiv:2408.09254 (2024).
[63] Cirac, J. I., Ekert, A. K., and Macchiavello, C. "Optimal purification of single qubits." *Phys. Rev. Lett.* **82** (1999), 4344–4347. arXiv:quant-ph/9812075.

[64] Keyl, M. and Werner, R. F. "The rate of optimal purification procedures." *Annales Henri Poincaré* **2** (2001), 1–26. arXiv:quant-ph/9910124.
[65] Fiurášek, J. "Optimal probabilistic cloning and purification of quantum states." *Phys. Rev. A* **70** (2004), 032308. arXiv:quant-ph/0403165.
[66] Fu, H. "Quantum state purification." MA thesis: University of Waterloo (2016).
[67] Childs, A. M., Fu, H., Leung, D., Li, Z., Ozols, M., and Vyas, V. "Streaming quantum state purification." *Quantum* **9** (2025), 1603. arXiv:2309.16387.
[68] Li, Z., Fu, H., Isogawa, T., and Chuang, I. "Optimal quantum purity amplification." arXiv:2409.18167 (2024).
[69] Grier, D., Leung, D., Li, Z., Pashayan, H., and Schaeffer, L. "Streaming quantum state purification for general mixed states." arXiv:2503.22644 (2025).
[70] Irani, S., Natarajan, A., Nirkhe, C., Rao, S., and Yuen, H. "Quantum search-to-decision reductions and the state synthesis problem." In: *CCC* (2022). arXiv:2111.02999.
[71] Lloyd, S., Mohseni, M., and Rebentrost, P. "Quantum principal component analysis." *Nat. Phys.* **10** (2014), 631–633. arXiv:1307.0401.
[72] Kimmel, S., Lin, C. Y.-Y., Low, G. H., Ozols, M., and Yoder, T. J. "Hamiltonian simulation with optimal sample complexity." *npj Quant. Inf.* **3** (2017), 13. arXiv:1608.00281.
[73] Babbush, R., McClean, J. R., Newman, M., Gidney, C., Boixo, S., and Neven, H. "Focus beyond quadratic speedups for error-corrected quantum advantage." *PRX Quantum* **2** (2021), 010103. arXiv:2011.04149.
[74] Babbush, R., Gidney, C., Berry, D. W., Wiebe, N., McClean, J., Paler, A., Fowler, A., and Neven, H. "Encoding electronic spectra in quantum circuits with linear T complexity." *Phys. Rev. X* **8** (2018), 041015. arXiv:1805.03662.
[75] Giovannetti, V., Lloyd, S., and Maccone, L. "Architectures for a quantum random access memory." *Phys. Rev. A* **78** (2008), 052310. arXiv:0807.4994.
[76] Nielsen, M. A. and Chuang, I. L. *Quantum computation and quantum information*. Cambridge University Press (2000).
[77] Weiss, D., Puri, S., and Girvin, S. "Quantum random access memory architectures using 3d superconducting cavities." *PRX Quantum* **5** (2024), 020312. arXiv:2310.08288.
[78] Hann, C. T., Zou, C.-L., Zhang, Y., Chu, Y., Schoelkopf, R. J., Girvin, S. M., and Jiang, L. "Hardware-efficient quantum random access memory with hybrid quantum acoustic systems." *Phys. Rev. Lett.* **123** (2019), 250501. arXiv:1906.11340.
[79] Wang, Z., Qiao, H., Cleland, A. N., and Jiang, L. "Quantum random access memory with transmon-controlled phonon routing." arXiv:2411.00719 (2024).
[80] Sala Cadellans, A. "A transmon based quantum switch for a quantum random access memory." MA thesis: Leiden University (2015).
[81] Chen, K. C., Dai, W., Errando-Herranz, C., Lloyd, S., and Englund, D. "Scalable and high-fidelity quantum random access memory in spin-photon networks." *PRX Quantum* **2** (2021), 030319. arXiv:2103.07623.
[82] Hong, F.-Y., Xiang, Y., Zhu, Z.-Y., Jiang, L.-z., and Wu, L.-n. "Robust quantum random access memory." *Phys. Rev. A* **86** (2012), 010306. arXiv:1201.2250.
[83] Cesa, F., Bernien, H., and Pichler, H. "Fast and error-correctable quantum RAM." arXiv:2503.19172 (2025).
[84] Ji, Z., Liu, Y.-K., and Song, F. "Pseudorandom quantum states." In: *CRYPTO* (2018), 126–152. arXiv:1711.00385.
[85] Brakerski, Z. and Shmueli, O. "(Pseudo) random quantum states with binary phase." In: *Theory of Cryptography* (2019), 229–250. arXiv:1906.10611.
[86] Arunachalam, S., Bravyi, S., Dutt, A., and Yoder, T. J. "Optimal algorithms for learning quantum phase states." In: *TQC* (2023), 3:1–3:24. arXiv:2208.07851.
[87] Aharonov, D. and Ben-Or, M. "Fault-tolerant quantum computation with constant error rate." *SIAM J. Comp.* **38** (2008), 1207–1282. Earlier version in *STOC'97*. arXiv:quant-ph/9906129.
[88] Shor, P. W. "Fault-tolerant quantum computation." In: *FOCS* (1996), 56–65. arXiv:quant-ph/9605011.
[89] Gottesman, D. "Fault-tolerant quantum computation with constant overhead." *Quantum Inf. Comput.* **14** (2014), 1338–1372. arXiv:1310.2984.
[90] Gottesman, D. "An introduction to quantum error correction and fault-tolerant quantum computation." In: *Proceedings of Symposia in Applied Mathematics* (2010), 13–58. arXiv:0904.2557.
[91] Fawzi, O., Grospellier, A., and Leverrier, A. "Constant overhead quantum fault tolerance with quantum expander codes." *Commun. ACM* **64** (2020), 106–114. Earlier version in *FOCS'18*. arXiv:1808.03821.

[92] Yamasaki, H. and Koashi, M. "Time-efficient constant-space-overhead fault-tolerant quantum computation." *Nat. Phys.* **20** (2024), 247–253. arXiv:2207.08826.
[93] Nguyen, Q. T. and Pattison, C. A. "Quantum fault tolerance with constant-space and logarithmic-time overheads." arXiv:2411.03632 (2024).
[94] Dankert, C. "Efficient simulation of random quantum states and operators." arXiv:quant-ph/0512217 (2005).
[95] Flammia, S. T. and Wallman, J. J. "Efficient estimation of Pauli channels." *ACM Trans. Quantum Comput.* **1** (2020). arXiv:1907.12976.
[96] Wallman, J. J. and Emerson, J. "Noise tailoring for scalable quantum computation via randomized compiling." *Phys. Rev. A* **94** (2016), 052325. arXiv:1512.01098.
[97] Mehta, R., Lee, G., and Jiang, L. "Analysis and suppression of errors in quantum random access memory errors under extended noise models." arXiv:2412.10318 (2024).
[98] Wood, C. J., Biamonte, J. D., and Cory, D. G. "Tensor networks and graphical calculus for open quantum systems." *Quantum Inf. Comput.* **15** (2015), 759–811. arXiv:1111.6950.
[99] Kitaev, A. Y. "Quantum measurements and the abelian stabilizer problem." arXiv:quant-ph/9511026 (1995).
[100] Watrous, J. "Notes on super-operator norms induced by Schatten norms." *Quantum Inf. Comput.* **5** (2005), 58–68. arXiv:quant-ph/0411077.
[101] Chen, Y., Gilyén, A., and de Wolf, R. "A quantum speed-up for approximating the top eigenvectors of a matrix." In: *SODA* (2025), 994–1036. arXiv:2405.14765.
[102] Mande, N. S. and de Wolf, R. "Tight bounds for quantum phase estimation and related problems." In: *ESA* (2023), 81:1–81:16. arXiv:2305.04908.
[103] Wiebe, N. and Granade, C. "Efficient Bayesian phase estimation." *Phys. Rev. Lett.* **117** (2016), 010503. arXiv:1508.00869.
[104] Gilyén, A., Su, Y., Low, G. H., and Wiebe, N. "Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics." In: *STOC* (2019), 193–204. Full version in arXiv:1806.01838.
[105] Dong, Y., Lin, L., and Tong, Y. "Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices." *PRX Quantum* **3** (2022), 040305. arXiv:2204.05955.
[106] van Dam, W. "On quantum computation theory." PhD thesis: Universiteit van Amsterdam (2002).
[107] Ajtai, M., Komlós, J., and Szemerédi, E. "An $O(n \log n)$ sorting network." In: *STOC* (1983), 1–9.
[108] Beals, R., Brierley, S., Gray, O., Harrow, A. W., Kutin, S., Linden, N., Shepherd, D., and Stather, M. "Efficient distributed quantum computing." *Proc. R. Soc. A* **469** (2013), 20120686. arXiv:1207.2307.
[109] Clader, B. D., Dalzell, A. M., Stamatopoulos, N., Salton, G., Berta, M., and Zeng, W. J. "Quantum resources required to block-encode a matrix of classical data." *IEEE Trans. Quantum Eng.* **3** (2022), 1–23. arXiv:2206.03505.
[110] Lin, L. "Lecture notes on quantum algorithms for scientific computation." arXiv:2201.08309 (2022).
[111] Tang, E. "A quantum-inspired classical algorithm for recommendation systems." In: *STOC* (2019), 217–228. arXiv:1807.04271.
[112] Tang, E. "Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions." *Phys. Rev. Lett.* **127** (2021), 060503. arXiv:1811.00414.
[113] Chia, N.-H., Gilyén, A. P., Li, T., Lin, H.-H., Tang, E., and Wang, C. "Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning." *J. ACM* **69** (2022), 1–72. Earlier version in *STOC'20*. arXiv:1910.06151.
[114] Shao, C. and Montanaro, A. "Faster quantum-inspired algorithms for solving linear systems." *ACM Trans. Quantum Comput.* **3** (2022). arXiv:2103.10309.
[115] Tang, E. "Dequantizing algorithms to understand quantum advantage in machine learning." *Nat. Rev. Phys.* **4** (2022), 692–693.
[116] Tang, E. "Quantum machine learning without any quantum." PhD thesis: University of Washington (2023).
[117] Yamasaki, H., Subramanian, S., Sonoda, S., and Koashi, M. "Learning with optimized random features: Exponential speedup by quantum machine learning without sparsity and low-rank assumptions." In: *NeurIPS* (2020), 13674–13687. arXiv:2004.10756.
[118] Hestenes, M. R. and Stiefel, E. "Methods of conjugate gradients for solving linear systems." *J. Res. Natl. Bur. Stand.* **49** (1952), 409–435.
[119] Hackbusch, W. *Iterative solution of large sparse systems of equations*. Springer (2016).
[120] Orsucci, D. and Dunjko, V. "On solving classes of positive-definite quantum linear systems with quadratically improved runtime in the condition number." *Quantum* **5** (2021), 573. arXiv:2101.11868.

[121] Gidney, C. "Windowed quantum arithmetic." arXiv:1905.07682 (2019).
[122] Gidney, C. and Ekerå, M. "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits." *Quantum* **5** (2021), 433. arXiv:1905.09749.
[123] Häner, T., Jaques, S., Naehrig, M., Roetteler, M., and Soeken, M. "Improved quantum circuits for elliptic curve discrete logarithms." In: *PQCrypto* (2020), 425–444. arXiv:2001.09580.
[124] Litinski, D. "How to compute a 256-bit elliptic curve private key with only 50 million Toffoli gates." arXiv:2306.08585 (2023).
[125] Berry, D. W., Gidney, C., Motta, M., McClean, J. R., and Babbush, R. "Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization." *Quantum* **3** (2019), 208. arXiv:1902.02134.
[126] von Burg, V., Low, G. H., Häner, T., Steiger, D. S., Reiher, M., Roetteler, M., and Troyer, M. "Quantum computing enhanced computational catalysis." *Phys. Rev. Res.* **3** (2021), 033055. arXiv:2007.14460.
[127] Lee, J., Berry, D. W., Gidney, C., Huggins, W. J., McClean, J. R., Wiebe, N., and Babbush, R. "Even more efficient quantum computations of chemistry through tensor hypercontraction." *PRX Quantum* **2** (2021), 030305. arXiv:2011.03494.
[128] Kim, I. H., Liu, Y.-H., Pallister, S., Pol, W., Roberts, S., and Lee, E. "Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules." *Phys. Rev. Res.* **4** (2022), 023019. arXiv:2104.10653.
[129] Mele, A. A. "Introduction to Haar measure tools in quantum information: A beginner's tutorial." *Quantum* **8** (2024), 1340. arXiv:2307.08956.
[130] Kretschmann, D., Schlingemann, D., and Werner, R. F. "The information-disturbance tradeoff and the continuity of Stinespring's representation." *IEEE Trans. Inf. Theory* **54** (2008), 1708–1717. arXiv:quant-ph/0605009.