

Universal quantum computing with twist-free and temporally encoded lattice surgery

Christopher Chamberland^{1,2} and Earl T. Campbell³

¹*AWS Center for Quantum Computing, Pasadena, CA 91125, USA*

²*IQIM, California Institute of Technology, Pasadena, CA 91125, USA*

³*AWS Center for Quantum Computing, Cambridge, UK*

Lattice surgery protocols allow for the efficient implementation of universal gate sets with two-dimensional topological codes where qubits are constrained to interact with one another locally. In this work, we first introduce a decoder capable of correcting spacelike and timelike errors during lattice surgery protocols. Afterwards, we compute logical failure rates of a lattice surgery protocol for a biased circuit-level noise model. We then provide a new protocol for performing twist-free lattice surgery. Our twist-free protocol eliminates the extra circuit components and gate scheduling complexities associated with the measurement of higher weight stabilizers when using twists. We also provide a protocol for temporally encoded lattice surgery that can be used to reduce both runtimes and the total space-time costs of quantum algorithms. Lastly, we propose a layout for a quantum processor that is more efficient for rectangular surface codes exploiting noise bias, and which is compatible with the other techniques mentioned above.

I. INTRODUCTION

Fault-tolerant quantum computing architectures enable the protection of logical qubits from errors by encoding them in error correcting codes, while simultaneously allowing for gates to be performed on such qubits. Importantly, failures arising during the implementation of logical gates do not result in uncorrectable errors as long as the total number of such failures remains below a certain fraction of the code distance [1–5]. In most practical settings, quantum logic gates are split into two categories. The first category corresponds to Clifford operations, which can be efficiently simulated by classical computers. The second category corresponds to non-Clifford operations, which cannot be efficiently simulated using purely classical resources. Early proposals for fault-tolerant quantum computation used transversal gates to perform logical Clifford operations [6]. Later, it was shown that by braiding defects in a surface code, some Clifford operations could be realised fault-tolerantly in a 2D local architecture with a high-threshold [7]. Recently, lattice surgery [8] has replaced the braiding approach due to its ability to retain locality constraints and high thresholds (features which are required by many hardware architectures), while additionally offering a much lower resource cost [9–11]. These approaches all perform non-Clifford gates by teleportation [12, 13] of magic states prepared by some distillation procedure [14–20]. Alternative ideas have been proposed for circumventing the need for magic states [21–24], but detailed studies [25–27] have not found any of these alternatives to be competitive for a wide range of failure rates below the surface code threshold.

Our work introduces the following key results. After briefly reviewing the model for Pauli based computation and its implementation via lattice surgery in Section II. Next, Section III explicitly provides a decoder compatible with lattice surgery. In particular, our decoder is capable of correcting both spacelike and timelike errors that occur during lattice surgery protocols. We then perform simulations of an $X \otimes X$ Pauli measurement using a biased circuit-level noise model.

In Section IV, we introduce a twist-free approach for measuring arbitrary Pauli operators using the surface code. Our

approach avoids the extra circuit and gate scheduling complexities that arise when using twists. We show that the approximate cost of avoiding twists is a $2\times$ slowdown in the algorithm runtime and a negligibly small additive cost to the number of logical qubits. We expect that twist-based lattice surgery has its own associated costs, which may exceed those of our twist-free approach, but twist performance has never been fully quantified and so represents a currently unknown factor in quantum computing design.

In Section V, we show how to reduce algorithm runtimes using a new technique that we call temporal encoding of lattice surgery. By using fast lattice surgery operations (which are inevitably noisier), errors arising from the extra noise can be corrected by encoding the sequence of measured Pauli operators within a classical error correcting code. The resulting runtime improvement grows (as a multiplicative factor) with the parallelizability of the algorithm and total algorithm runtime. We find in a regime of interest to quantum algorithms of a practical scale, we can achieve a $2.2\times$ runtime improvement. The encoding does not directly lead to additional qubit overhead costs since it occurs in the time domain, and so the overall spacetime complexity is improved.

Lastly, in Section VI, we describe our core-cache architecture. We show that by using thin rectangular strips of surface codes for settings where a large noise bias is present, routing overhead costs in our proposed architecture add a factor of $1.5\times$ to the total resource costs for performing lattice surgery. This can be compared with the $2\times$ cost of Litinski’s fast data access structures [11]. Furthermore, we provide a layout that compactly stores surface code patches in a cache to further reduce the extra overhead arising from routing costs, at the cost of some additional time needed for reading/writing to the cache. Using the numerical results obtained in Section III, in Appendix C we provide resource cost estimates for simulating the Hubbard model using our core-cache architecture.

II. BRIEF REVIEW OF UNIVERSAL QUANTUM COMPUTING VIA LATTICE SURGERY

In Section II A, we briefly review the principles of Pauli based computation used throughout this work. We then review in Section II B how multi-qubit Pauli operators are measured using lattice surgery.

A. Overview of Pauli based computation

In the model of Pauli-based computation (PBC), we have a reserve of magic states and drive the computation by performing a sequence of multi-qubit Pauli measurements $\{P_1, P_2, \dots, P_\mu\}$ where later Pauli measurements depend on measurement outcomes of earlier measurements. In this notation, P_2 does not denote a specific Pauli, but one conditional on the outcome of P_1 . This conditionality occurs because (in the circuit picture) each Pauli measurement would be followed by a conditional Clifford operation. However, in a PBC, these Cliffords are conjugated to the end of the computation, thereby changing subsequent Pauli measurements. Since in a PBC all Cliffords are performed “in software”, it is clear the algorithm runtime will be independent of the Clifford complexity. The idea of PBC appears throughout the literature, but the phrase Pauli-based computation was first coined in Ref. [28]. In Fig. 1, we present several computationally equivalent circuit diagrams for performing 2 T gates, with the last diagram representing the PBC approach.

In Section II B, we review how multi-qubit Pauli measurements can be performed using lattice surgery. Crucially, even when Pauli operators commute, it might not be possible to measure such Pauli operators simultaneously due to the extra space required to performing lattice surgery (known as the routing space). The reason is that the required lattice surgery operators often need access to the same routing space. Therefore, it is appropriate to consider sequentially measuring each Pauli operator, which we call a sequential Pauli based computation (seqPBC). In seqPBC, the time required to execute all the Pauli measurements will then be proportional to $T_{\text{PBC}} = (d_m + 1)\mu$ where we budget +1 for resetting qubits between lattice surgery operations. Here d_m corresponds to the number of rounds of stabilizer measurements during lattice surgery, and μ is the number of sequential Pauli operators being measured. The proportionality factor will depend on the time required to measure the surface code stabilizers during one syndrome measurement round.

There are several contributions to the number of Pauli measurements μ . At a high level of the stack, we may think of a quantum algorithm as consisting of a series of unitaries with some Pauli measurements for readout, and let N_A denote the number of such algorithmic readout measurements. However, as we saw in Fig. 1, non-Clifford unitaries are performed by measurements. If an algorithm has N_T T -gates, then we also need an additional N_T Pauli measurements. The Clifford plus T gate set is universal. However, it is advantageous to use an overcomplete gate set such as Clifford plus T and Toffoli. While Toffoli can be synthesized using 4 T gates [29], it is

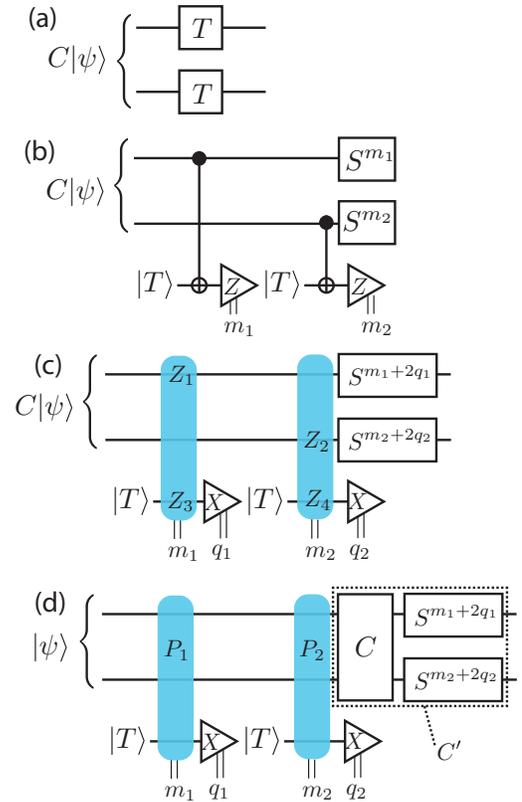


FIG. 1. Equivalent approaches to implementing two T gates. Blue rounded rectangles show multi-qubit Pauli measurements. (a) A simple unitary circuit approach. (b) Using gate teleportation with $|T\rangle := (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$ magic states, CNOT gates, Pauli Z measurements with outcomes m_1 and m_2 and classical conditioned Cliffords based on these outcomes. The conditional S gates are $S = |0\rangle\langle 0| + i|1\rangle\langle 1|$. (c) Using gate teleportation with the CNOT gates replaced by two-qubit Pauli measurements. (d) Given an input state $|\psi\rangle$ which carries a Clifford frame correction C , we have conjugated C through the circuit so that the multi-qubit Pauli-measurements are now $P_1 = CZ_1Z_3C^\dagger$ and $P_2 = CZ_2Z_4C^\dagger$ (which commute) and the output state carries a new Clifford frame correction $C' = S_1^{m_1+2q_1} S_2^{m_2+2q_2} C$. This last circuit represents the standard PBC approach for computing $T^{\otimes 2}$.

often more efficient to directly prepare Toffoli magic states [29–32]. Furthermore, it only takes 3 Pauli measurements to teleport a Toffoli state rather than the 4 measurements needed to teleport T states and then synthesize a Toffoli. As such, if an algorithm can be executed with N_{TOF} Toffoli gates and N_T T -gates, then we need $N_T + 3N_{\text{TOF}}$ measurements to perform these teleportations. Further refinements are possible by using an even richer gate set and preparing more exotic states [33–35], but we will not explicitly discuss those schemes here. Lastly, we can replace some non-Clifford gates with Pauli measurements and feedforward (with no magic state needed). For instance, such a measurement appears in Gidney’s circuits for adders [36] and more generally any uncomputation subroutine of an algorithm where the Toffoli gates are replaced with Pauli measurements and feedforward. If we use N_{unTOF} to denote the

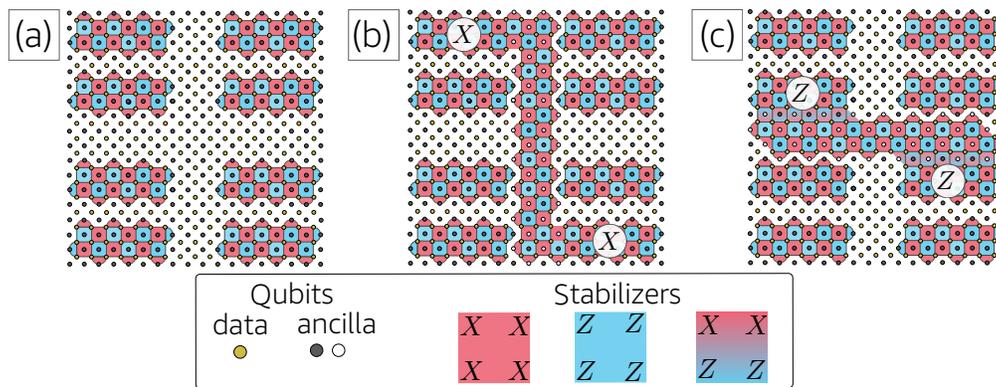


FIG. 2. A simple example of lattice surgery illustrating the space overhead needed for routing. In (a), we show 8 rectangular surface code patches with some surrounding idle qubits that form the routing space. In (b) we show the stabilizers needed to measure a logical $X \otimes X$ operator between the separated surface code patches using the routing space between them. Similarly, in (c) we show the stabilizers needed to measure a logical $Z \otimes Z$ operator between the separated surface code patches using the available routing space. White ancillas mark the stabilizers

that directly contribute towards computing the parity of the $X \otimes X$ and $Z \otimes Z$ measurement outcomes.

number of Toffoli uncomputations performed in this manner, then for a Clifford plus T and Toffoli gate set, the total number of Pauli measurements will be $\mu = N_A + N_T + 3N_{\text{TOF}} + N_{\text{unTOF}}$. Note that in many algorithms Toffolis exclusively appear in compute/uncompute pairs, and then we have $N_{\text{unTOF}} = N_{\text{TOF}}$. Furthermore, algorithms often only have a small number of qubit readouts, so $N_A \ll N_T, N_{\text{TOF}}$. As such, we commonly have $\mu \approx N_T + 4N_{\text{TOF}}$.

An architecture also requires time T_{magic} to produce the required magic states, say N_T T -states and N_{TOF} Toffoli states. If an architecture produces all the required magic states in a shorter amount of time than is required to teleport them, so that $T_{\text{magic}} < T_{\text{PBC}}$, then we say the seqPBC is Clifford-bottle necked. On the other hand, if $T_{\text{magic}} \geq T_{\text{PBC}}$, we say it is magic-state bottle necked. The running time of the algorithm will be determined by $\max\{T_{\text{magic}}, T_{\text{PBC}}\}$.

It is informative to briefly review the impact of these bottle necks on the history of algorithm resource analysis. The time T_{magic} can be made arbitrarily small, by simply increasing the number of magic state factories, though this comes at an increased qubit cost. Some early algorithm overhead estimates [5, 37, 38] minimized T_{magic} by having a large number of factories leading to widespread claims that magic state factories could be $\sim 99\%$ of the whole device. However, such resource estimates ignore the question of how quickly these states can be teleported and ignored the T_{PBC} bottleneck. Accounting for this bottleneck, there is no benefit from pushing T_{magic} to be small $T_{\text{magic}} \ll T_{\text{PBC}}$. As such, more recent and careful overhead analyses [32, 39–41] have assumed only a few factories, which is enough to achieve $T_{\text{magic}} \approx T_{\text{PBC}}$ and results in only a small percentage of device footprint (often less than 1%) being used as a magic state factory. While these analyses have minimal qubit cost, the runtime now is T_{PBC} bottlenecked, motivating rigorous approaches to beating the T_{PBC} bottleneck. Later, in Section V A, we will review prior lattice surgery methods to speed up algorithms by using addi-

tional teleportation gadgets, though this comes at a high qubit cost. We will then introduce our own approach that instead uses temporal encoding of lattice surgery (TELS).

B. Overview of lattice surgery

For quantum hardware where physical qubits can only interact with one another locally, lattice surgery is a fault-tolerant protocol that can be used to measure arbitrary multi-qubit Pauli operators. The main idea is to encode the logical qubits in some topological code arranged in a two-dimensional layout (in this work, all logical qubits are encoded in the rotated surface code [42]). The layout contains extra routing space between the surface code patches which consists of additional qubits. By applying the appropriate gauge-fixing (see for instance Ref.[43]) in the routing space, which involves measuring surface code stabilizers, the surface code patches involved in the Pauli measurement are merged into one larger surface code patch. After gauge fixing, the parity of the measurement outcome of the multi-qubit Pauli operator being measured is obtained by taking the products of the appropriate stabilizers in the routing space. Lastly, the surface code patches for each logical qubit can be detached from the merged patch by measuring the qubits in the routing space in the appropriate basis. An illustration of $X \otimes X$ and $Z \otimes Z$ Pauli measurements is shown in Fig. 2. Products of the surface code stabilizers marked by white ancilla give the parity of the $X \otimes X$ and $Z \otimes Z$ measurement outcomes.

When performing a lattice surgery measurement of a logical Pauli operator, there will be some probability that we obtain the wrong outcome. Even with large code distances, the lattice surgery measurement could still fail due to time-like errors occurring during the finite time allowed for lattice surgery. The probability of these failure events is exponentially suppressed in the number of rounds d_m for which we repeat the stabilizer

measurements during lattice surgery. Therefore, we call d_m the measurement distance which quantifies the protection against repeated measurement failures during lattice surgery, and hence explains our choice of subscript. This exponential suppression will hold until $d_m \gg O(d_z, d_x)$ when logical Pauli errors become the dominate mechanism again. Let us assume that code distances d_x and d_z are chosen so that even a single logical Z and X error is very unlikely over the course of the whole computation. We expect, and numerically find, that these timelike errors occur with a probability \mathbb{P} for which we have a bound of the form

$$\mathbb{P} \leq La(pb)^{c(d_m+1)}, \quad (1)$$

where p quantifies the physical gate failure probabilities, $\{a, b, c\}$ are constants and L is the area of the patch used for lattice surgery. The value of L will vary for different measurements and different layouts, but it will be convenient to think of it as a constant representing the worst (or average) case area of lattice surgery patches.

In general, if we want to sequentially perform μ Pauli measurements in the algorithm and we want them to fail with probability no more than δ , then we choose d_m to be large enough that

$$\mu La(pb)^{c(d_m+1)} \approx 1 - (1 - \mathbb{P})^\mu \leq \delta \quad (2)$$

where the approximation holds for small \mathbb{P} . In Section III, after introducing a decoder compatible with lattice surgery protocols, we will compute timelike failure probabilities in addition to probabilities for other noise processes given a biased circuit-level noise model. Such results will allow us to obtain accurate resource overhead estimates for implementing quantum algorithms and are discussed further in Section VI.

III. DECODING TIMELIKE ERRORS DURING LATTICE SURGERY

In this section, we explicitly provide a decoding protocol for correcting both spacelike and timelike errors during lattice surgery. In particular, our protocol protects data qubits (the logical qubits of an algorithm) encoded in surface code patches while at the same time corrects logical multi-qubit Pauli measurement failures which can occur during lattice surgery. We then provide numerical results for performing $X \otimes X$ measurements showing both the logical parity failure rate as a function of the number of syndrome measurement rounds and the data qubit failure rates. No previous work has performed such realistic, circuit-level simulations of lattice surgery. For comparison, in Ref. [32] timelike errors were simulated, but Ref. [32] used a toy model with idealized boundaries to exclude logical spacelike errors and thereby simplifying both the simulations and required decoding algorithm.

In Fig. 3, we provide an example of an $X \otimes X$ Pauli measurement performed between two $d_x = 5$, $d_z = 7$ surface code patches. Measuring the surface code stabilizers of the ancilla qubits prepared in the $|+\rangle$ state and taking the products of stabilizers marked by white vertices gives the measurement

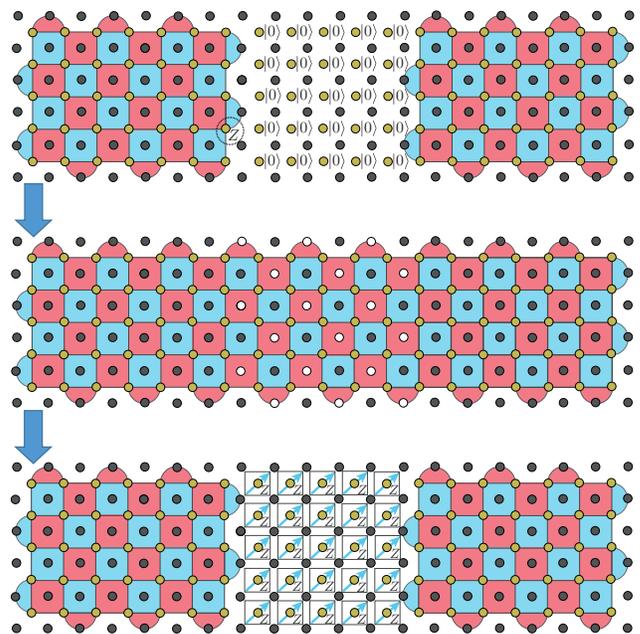


FIG. 3. Example of an $X \otimes X$ measurement of two $d_x = 5$, $d_z = 7$ surface code patches. The data qubits in the routing space are first prepared in the $|0\rangle$ state. In the first round of stabilizer measurements where the surface code patches are merged into one large surface code patch, the product of all stabilizers marked by white vertices (which we call parity vertices) gives the result of the $X \otimes X$ measurement. After measuring the stabilizers of the merged patch for d_m rounds, the patches are split by measuring the data qubits in the routing region in the Z basis. In the first round of the merge, measurement errors occurring on parity vertices can result in wrong $X \otimes X$ measurement outcome. Additionally, an odd number of data qubit Z errors along the boundary of the surface code lattices prior to the merge, such as the one circled in the top row of the figure, can also give the wrong measurement outcome.

outcome of $X \otimes X$. In what follows, white vertices whose product gives the result of a $P_1 \otimes P_2 \otimes \dots \otimes P_k$ Pauli measurement will be referred to as *parity vertices*. We say that a logical timelike failure occurs if a set of errors result in the wrong parity measurement of $P_1 \otimes P_2 \otimes \dots \otimes P_k$. Further, we assume that split surface code patches are measured for r rounds before being merged by the Pauli measurement in round $r + 1$.

When measuring $P_1 \otimes P_2 \otimes \dots \otimes P_k$ using lattice surgery, in addition to an odd number of measurement errors during round $r + 1$, an odd number of errors along the boundaries of the patches before the merge can also result in a logical timelike failure. An example is provided in Fig. 3 where a single Z error along the left boundary of the surface patch prior to the merge gives the wrong parity of $X \otimes X$. We also note that during round $r + 1$, an odd number data-qubit errors which anticommute with $P_1 \otimes P_2 \otimes \dots \otimes P_k$ will also result in a logical timelike failure.

The above examples show that in order to obtain the correct parity measurement of a $P_1 \otimes P_2 \otimes \dots \otimes P_k$ operator in the presence of full circuit level noise, one must have a decoding

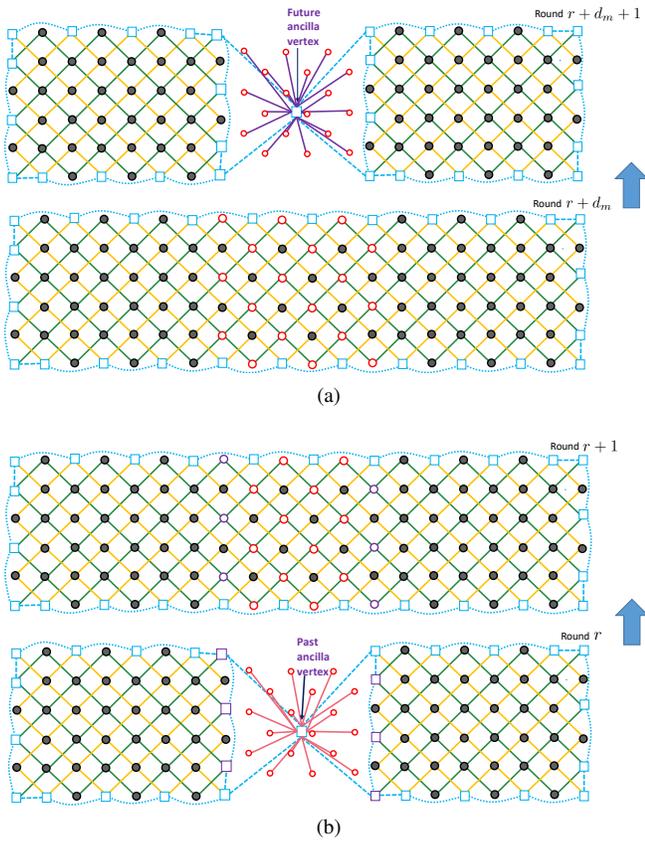


FIG. 4. (a) Two-dimensional slices of the surface code matching graphs for syndrome measurement rounds $r + d_m$ (bottom) and $r + d_m + 1$ (top). The graph in round $r + d_m + 1$ includes the future ancilla vertex with future vertical edges (solid purple edges) connecting to the parity vertices (white vertices circled in red) of the matching graph in round $r + d_m$. (b) Two-dimensional slices of the surface code matching graphs for syndrome measurement rounds r (bottom) and $r + 1$ (top). The graph in round $r + 1$ includes the past ancilla vertex with past vertical edges (solid pink edges) connecting to the parity vertices (white vertices circled in red) of the matching graph in round $r + 1$. Transition vertices are the purple boundary vertices on the graph prior to the merge (note that transition vertices appear in rounds 1 to r), in addition to the purple parity vertices in round $r + 1$.

scheme which, while constantly correcting errors on the data qubit patches, also corrects spacelike and timelike errors which can flip the parity of the measurement outcome.

In order to correct logical timelike failures using a minimum-weight-perfect-matching (MWPM) decoder [44], we must add timelike boundaries to the matching graphs of the surface code as shown in Fig. 4. In particular, we divide the measurement of an operator $P_1 \otimes P_2 \otimes \dots \otimes P_k$ into three steps. In the first step, the surface code patches encoding data qubits are measured for r rounds. In round $r + 1$, the patches are merged by measuring the appropriate operators in the ancilla space (see for instance Fig. 3) and the parity of the measurement outcome is given by the product of all parity vertices. The merged patches are measured for d_m rounds, and then in round $d_m + 1$, the qubits in the ancilla space are measured in the appropriate basis to split the patches back to their original

configuration. In round r , we add extra virtual vertices to the matching graphs of the split surface codes with vertical edges that connect to the parity vertices of the merged matching graph in round $r + 1$ (see the pink edges in Fig. 4b for the $X \otimes X$ measurement). We call such vertices *past ancilla vertices* and the pink edges incident to them *past vertical edges*. Similarly, in round $r + d_m + 1$ (i.e. right after the split), we add virtual vertices to the matching graphs of the split surface codes with vertical edges which connect to the parity vertices of the merged matching graph in round $r + d_m$ (see the purple edges in Fig. 4a for the $X \otimes X$ measurement). We call such vertices *future ancilla vertices* and the purple edges incident to them *future vertical edges*. Importantly, the pink vertical edges that are incident to the past ancilla vertices and to the parity vertices in the round $r + 1$ have zero weight, while the purple vertical edges incident to the parity vertices in the round $r + d_m$ and the future ancilla vertices have non-zero weights. These weights are computed from all timelike failure processes which can result in measurement errors in the round $r + d_m$. When performing MWPM over the full syndrome history, the parity of the measurement outcome of $P_1 \otimes P_2 \otimes \dots \otimes P_k$ is flipped if there are an odd number of highlighted vertical edges incident to parity vertices in rounds $r + 1$ and $r + 2$. In such a setting one would require a sequence of consecutive measurement errors that is greater than $(d_m - 1)/2$ in order to cause a timelike logical failure. Note that since the measurement outcomes of the parity vertices in round $r + 1$ are random, such vertices are never highlighted. If a change in the measurement outcomes of a subset of the parity vertices are observed between rounds $r + 1$ and $r + 2$, then vertices corresponding to such parity vertices in round $r + 2$ would be highlighted. Furthermore, green horizontal edges incident to the parity vertices in round $r + 1$ are taken to have zero weight (or can be omitted).

As mentioned above, a lattice surgery decoder also needs to correct logical timelike failures arising from sets of data qubit errors along surface code boundaries of the data qubit patches prior to merging them (recall the example shown in Fig. 3). The decoder also needs to correct wrong parity measurements arising from data qubit errors in round $r + 1$ which anticommute with the Pauli operator being measured by lattice surgery. In constructing such a decoder, note that prior to merging the surface code patches for the $X \otimes X$ measurement, a single Z error along the relevant boundaries would result in a highlighted edge (after implementing MWPM over the full syndrome history) incident to one of the purple boundary vertices shown in Fig. 4b. Note that such boundary vertices become parity vertices after merging the surface code patches. For the measurement of a general $P_1 \otimes P_2 \otimes \dots \otimes P_k$ Pauli operator, we define *transition vertices* to be vertices in the set $V_{Bd}^{(s)} = \{v_{b_1}^{(s)}, \dots, v_{b_m}^{(s)}\}$, where $1 \leq s \leq r + 1$. When $s < r + 1$, $\{v_{b_1}^{(s)}, \dots, v_{b_m}^{(s)}\}$ are labels for boundary vertices of the graphs of split surface code patches which become parity vertices in round $r + 1$. If $s = r + 1$, then $V_{Bd}^{(r+1)}$ is the set of the parity vertices along the boundaries of the data and ancilla patch regions as shown for instance in Fig. 5. Based on previous observations, after implementing MWPM over the full syndrome history of a multi-qubit Pauli measurement

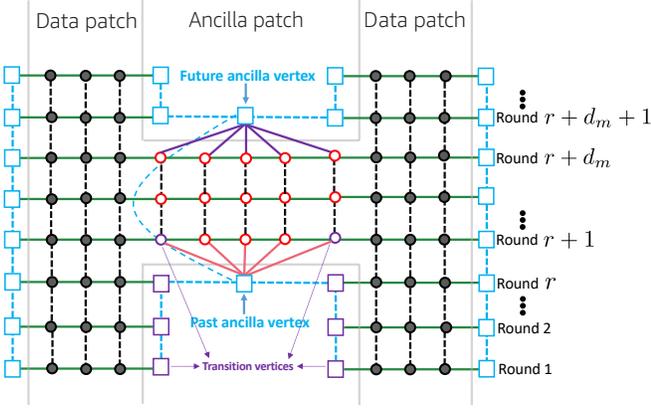


FIG. 5. A two-dimensional slice of the surface code matching graphs in the timelike direction for syndrome measurement rounds performed during an $X \otimes X$ measurement protocol using lattice surgery. The vertical axis corresponds to the timelike direction. We show a subset of the vertices and edges of the matching graph for correcting Z errors for an $X \otimes X$ measurement using lattice surgery. Parity vertices are the white and red vertices in the ancilla patch region. Transition vertices are boundary and parity vertices colored in purple. Pink edges are past vertical edges incident to the past ancilla vertex and parity vertices in round $r + 1$, whereas purple edges are future ancilla edges incident to the parity vertices in round $r + d_m$ and the future ancilla vertex. We also add a dashed-blue weightless edge connecting the past and future ancilla vertices.

via lattice surgery, if there are an odd number of highlighted edges on the data qubit regions incident to transition vertices, the parity of the Pauli measurement needs to be flipped. An illustration of a two-dimensional slice of the matching graphs of Fig. 4 in the timelike direction (which contains a subset of the spacelike edges and vertices) is shown in Fig. 5. In particular, the figure illustrates transition vertices, the past and future ancilla vertices in addition to the past and future vertical edges.

Combining all the notions introduced in this section, the decoding algorithm for implementing a multi-qubit Pauli measurement via lattice surgery is described in Algorithm 1. Each highlighted edge in step 7) of Algorithm 1 encodes a particular data qubit correction. Writing such corrections as a binary row vector, where each column corresponds to a data qubit, we add all corrections arising from each highlighted edge using modulo-two arithmetic. Further, note that space-time correlated edges incident to parity vertices in round $r + 1$ need to be treated with care in order to correct errors up to the full code distance. In particular, a subset of the space-time correlated edges incident to transition vertices can also contribute to v_2 (defined in Algorithm 1). A more careful treatment of such edges is provided in Appendix B.

We point out a simplification that can be made in the implementation of Algorithm 1. Notice that vertices in $V_{\text{par}}^{(r+1)}$ are never highlighted during MWPM due to the random outcomes of the ancilla patch stabilizers marked by white vertices in round $r + 1$. As such, one could remove all vertices in $V_{\text{par}}^{(r+1)}$ and instead have the past vertical edges incident to the

Algorithm 1: Decoding algorithm for measuring $P_1 \otimes P_2 \otimes \dots \otimes P_k$ via lattice surgery.

Result: Data qubit and parity measurement corrections.
initialize: $v_1 = v_2 = 0$. Let G_r be the graph for the surface code patches before, during and after the merge.;

Measurement: Measure the stabilizers of the split surface code patches for r rounds. Merge the surface code patches in round $r + 1$ via lattice surgery to perform the $P_1 \otimes P_2 \otimes \dots \otimes P_k$ measurement and let s_{par} be the parity of the measurement outcome. Repeat the stabilizer measurements of the merged patch for $d_m - 1$ rounds.;

- 1) Add the past ancilla vertex v_{past} to G_r for the round r (round before the merge). Let $V_{\text{par}}^{(r+1)} = \{v_{\text{par}}^{(1)}, \dots, v_{\text{par}}^{(k)}\}$ be the set of parity vertices for the syndrome measurement round $r + 1$. Add weightless past vertical edges to G_r which are incident to v_{past} and all vertices $v \in V_{\text{par}}^{(r+1)}$. Add weightless edges to G_r which are between v_{past} and virtual boundary edges of all surface code patches;
 - 2) Add the future ancilla vertex v_{future} to G_r for the round $r + d_m + 1$ (round after the merge). Let $V_{\text{par}}^{(r+d_m)} = \{\tilde{v}_{\text{par}}^{(1)}, \dots, \tilde{v}_{\text{par}}^{(k)}\}$ be the set of parity vertices for the syndrome measurement round $r + d_m$. Add future vertical edges (of non-zero weight) to G_r which are incident to v_{future} and all vertices $v \in V_{\text{par}}^{(r+d_m)}$;
 - 3) Add a weightless edge to G_r which is incident to v_{past} and v_{future} ;
 - 4) Set all edges incident to any two vertices $v_i, v_j \in V_{\text{par}}^{(r+1)}$ to have zero weight;
 - 5) Given the full syndrome measurement history, if the total number of highlighted vertices (obtained by taking the difference between any two consecutive syndrome measurement rounds modulo 2) is odd, highlight v_{future} ;
 - 6) Implement MWPM on G_r . Set v_1 to be the number of highlighted edges incident to vertices in $V_{\text{par}}^{(r+1)}$ and $V_{\text{par}}^{(r+2)}$, and v_2 to be the number of highlighted edges incident to transition vertices in the data-qubit patch regions. If $v_1 + v_2$ is odd, set $s_{\text{par}} \rightarrow s_{\text{par}} + 1 \pmod{2}$;
 - 7) Apply data qubit corrections based on all highlighted two-dimensional and space-time correlated edges.
-

vertices in $V_{\text{par}}^{(r+2)}$. In such a setting, edges incident to $V_{\text{par}}^{(r+1)}$ and $V_{\text{par}}^{(r+2)}$ would be removed, and their weights would be assigned to the past vertical edges.

Using the decoding algorithm given in Algorithm 1, below in Eqs. (3) to (5) and Fig. 7, we describe simulation results of a full circuit-level noise simulation of various code distances and syndrome measurement rounds to estimate the parameters in Eq. (1) for a $X \otimes X$ measurement. We chose the following biased circuit-level noise model:

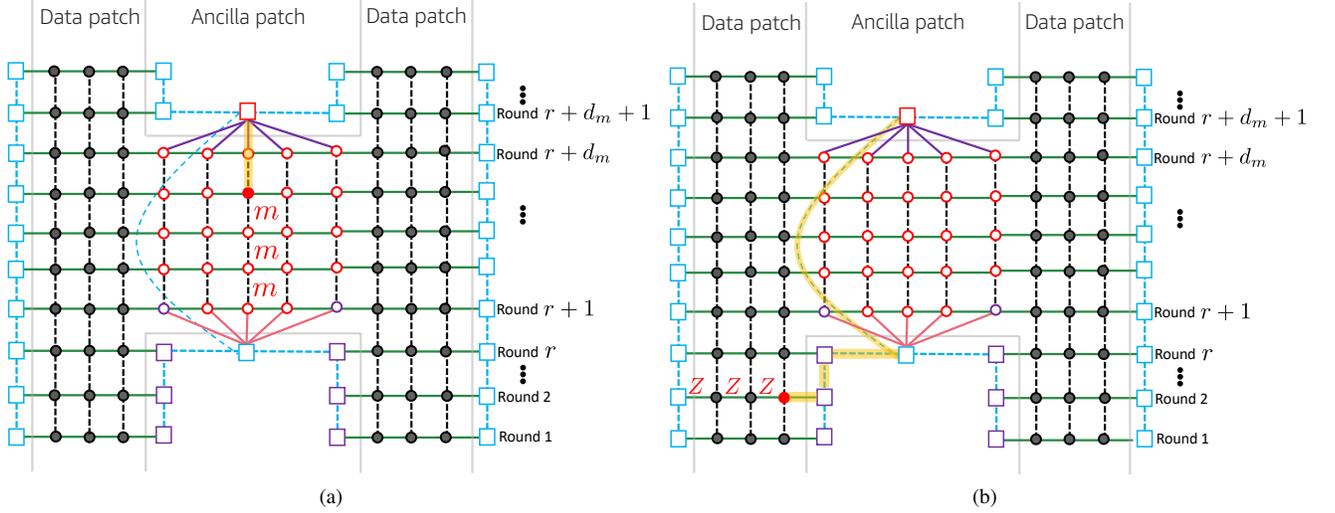


FIG. 6. (a) Example of a series of consecutive measurement errors on the same parity vertex resulting in a $(0, 1, 0, 0)$ logical failure for an $X \otimes X$ measurement. (b) Example of a horizontal string of spacelike Z errors on the left data qubit patch resulting in a $(1, 1, 0, 0)$ logical failure.

1. Each single-qubit gate location is followed by a Pauli Z error with probability $\frac{p}{3}$ and Pauli X and Y errors each with probability $\frac{p}{3\eta}$.
2. Each two-qubit gate is followed by a $\{Z \otimes I, I \otimes Z, Z \otimes Z\}$ error with probability $p/15$ each, and a $\{X \otimes I, I \otimes X, X \otimes X, Z \otimes X, Y \otimes I, Y \otimes X, I \otimes Y, Y \otimes Z, X \otimes Z, Z \otimes Y, X \otimes Y, Y \otimes Y\}$ each with probability $\frac{p}{15\eta}$.
3. With probability $\frac{2p}{3\eta}$, the preparation of the $|0\rangle$ state is replaced by $|1\rangle = X|0\rangle$. Similarly, with probability $\frac{2p}{3}$, the preparation of the $|+\rangle$ state is replaced by $|-\rangle = Z|+\rangle$.
4. With probability $\frac{2p}{3\eta}$, a single-qubit Z basis measurement outcome is flipped. With probability $\frac{2p}{3}$, a single-qubit X -basis measurement outcome is flipped.
5. Lastly, each idle gate location is followed by a Pauli Z with probability $\frac{p}{3}$, and a $\{X, Y\}$ error each with probability $\frac{p}{3\eta}$.

In our simulations, we chose $\eta = 100$ and for simplicity added a single idle location during ancilla qubits measurement and re-initialization. Note that in the limit $\eta \rightarrow 1$, the above noise model reduces to the depolarizing noise model used in Refs.[45, 46] with the exception that two idle locations were included during measurement and re-initialization of the ancillas. Furthermore, the above noise model assumes that the duration of a CNOT gate is identical to the duration of an ancilla measurement and re-initialization.

For a biased circuit-level noise model (such as the one described above) and an $X \otimes X$ Pauli measurement implemented via lattice surgery, there are fifteen different types of failure mechanisms that can arise during the protocol. We label such failure mechanisms using the binary string $(b_{ZL}, b_{tL}, b_{ZR}, b_X)$.

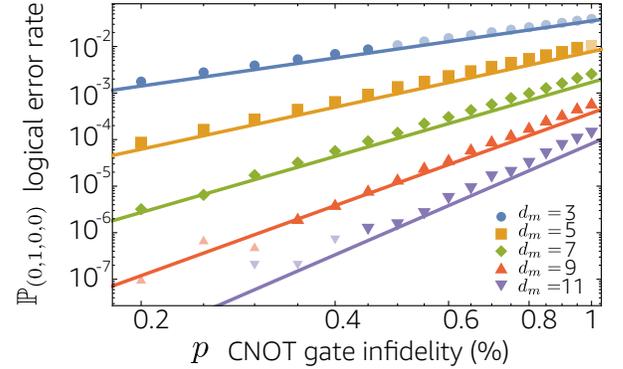


FIG. 7. Comparison between the best fit polynomial $\mathbb{P}_{(0,1,0,0)}$ for various values of d_m given in Eq. (3) with the data obtained from our Monte Carlo simulations. We chose parameters where $d_x = 9$, $d_z = 11$, $r = d_z$ and $\ell = 5$. Translucent data points were omitted when obtaining the fitting polynomial in Eq. (3).

The bit $b_{ZL} = 1$ corresponds to a logical Z error on the left patch, whereas $b_{ZR} = 1$ corresponds to a logical Z error on the right patch. The bit $b_{tL} = 1$ indicates a logical timelike failure. Finally, the bit b_X indicates if a logical X error occurred during the lattice surgery protocol. Examples of such failure mechanisms are shown in Fig. 6. For instance, in Fig. 6a, a series of consecutive measurement errors occur on the same parity vertex, starting in round $r + 1$. Since a single measurement error occurs in round $r + 1$, the wrong parity of $X \otimes X$ is measured. Due to the series of measurement errors, a single parity vertex near the top boundary is highlighted in G_r . The shortest path correction (highlighted in yellow) matches the highlighted parity vertex to the future ancilla vertex. Hence no parity corrections are applied and a logical $(0, 1, 0, 0)$ error occurs.

Another example is shown in Fig. 6b, where a high-weight string of Z errors results in the highlighted vertices shown in the figure. Prior to the merge, there are no Z errors at the boundary between the ancilla and data-qubit patches and thus the correct parity of $X \otimes X$ is measured. However, the minimum-weight path (highlighted in yellow) connecting the highlighted vertex to the future ancilla vertex goes through a transition vertex, so that $v_2 = 1$. The correction thus results in a logical Z error on the left data-qubit patch, in addition to a logical parity measurement failure (since the decoder incorrectly flips the parity) leaving the code with a logical $(1, 1, 0, 0)$ error.

We note that for syndrome measurement rounds up to $r + d_m$, higher order failure mechanisms are required to produce logical failures corresponding to $(1, 0, 0, 0)$, $(0, 0, 1, 0)$, $(1, 0, 1, 0)$ and $(1, 1, 1, 0)$ bit strings. Failures corresponding to such strings are thus much less likely compared to $(1, 1, 0, 0)$, $(0, 1, 0, 0)$ and $(0, 1, 1, 0)$. For instance, to obtain a logical failure of the type $(1, 0, 0, 0)$, a logical error on the left surface code patch would need to occur without flipping the parity of the $X \otimes X$ measurement. As such, in addition to a logical Z error occurring before the merge, a second failure mechanism would need to occur to undo the wrong parity flip (such as a string of measurement errors like the one shown in Fig. 6a). Given these observations, we only present the logical failure rate polynomials $\mathbb{P}_{(0,1,0,0)}$, $\mathbb{P}_{(1,1,0,0)}$ and $\mathbb{P}_{(0,1,1,0)}$. Note that due to the high noise bias, we choose a d_x distance such that $\mu \mathbb{P}_{(0,0,0,1)} \leq \delta$, where μ is the number of lattice surgery operation in the algorithm.

Our simulations are performed for syndrome measurement rounds 1 to $r + d_m$, based on the biased circuit-level noise model described above. The last round is a round of perfect error correction to guarantee projection to the code space. We used Algorithm 1 to correct both spacelike and timelike errors, where each edge in step 7) of the algorithm encodes a particular correction on a subset of the data qubits. Note that we do not perform a round of perfect error correction between rounds r and $r + 1$ as was done in Ref. [43]. Instead, we performed MWPM using the full syndrome measurement history from rounds 1 to $r + d_m$, and use Algorithm 1 to determine which corrections are applied.

For each of the dominant failure mechanisms in $(b_{ZL}, b_u, b_{ZR}, b_X)$, we fit all our data to an ansatz with two free parameters to generate the failure rate polynomials given by

$$\mathbb{P}_{(0,1,0,0)} = 0.01634 d_x \ell (21.93p)^{(d_m+1)/2}, \quad (3)$$

$$\mathbb{P}_{(1,1,0,0)} = 0.03148 d_x (28.91p)^{(d_z+1)/2}, \quad (4)$$

$$\mathbb{P}_{(0,1,1,0)} = 0.03 d_x (28.95p)^{(d_z+1)/2}, \quad (5)$$

$$\mathbb{P}_{(0,0,0,1)} = 0.0148 d_z (0.762p)^{(d_x+1)/2}. \quad (6)$$

In Eq. (3), ℓ corresponds to the width of the spacing between two data-qubit surface code patches. All logical error rate polynomials in Eqs. (4) to (6) provide error rates per syndrome measurement round. Per round error rates are computed by varying the number of syndrome measurement rounds for fixed d_x and d_z distance, repeating such procedures for different

d_x and d_z distances, and fitting all the obtained data to an ansatz. Note that the d_z distance in Eq. (6) is taken to be the d_z distance of the full merged surface code patches, whereas d_z in Eqs. (4) and (5) is the d_z distance of the individual data qubit patches (since logical Z errors are much less likely to occur when surface code patches are merged due to the increased d_z distance). In Fig. 7, we compare the best-fit polynomial $\mathbb{P}_{(0,1,0,0)}$ with a representative subset of our data obtained from our Monte Carlo simulations for various values of d_m , where the chosen parameters are described in the caption. The plot shows the exponential suppression in purely timelike error probabilities as a function of d_m , and that the data is in good agreement with our best-fit polynomials. In Section VI, after introducing our protocol for minimizing routing costs, we use the logical failure rate polynomials in Eqs. (3) to (6) to estimate the overhead costs for implementing quantum algorithms.

We conclude this section by pointing out that the above labels used in the logical error rate polynomials, which represent different failure mechanisms that can occur during lattice surgery, can be generalized using $k + 2$ bits for an arbitrary $P_1 \otimes P_2 \otimes \dots \otimes P_k$ Pauli measurement. Out of the $k + 2$ bits, k bits are used to represent a logical Pauli error on the data qubit patches which can also flip the parity of the measured Pauli (such as logical Z errors for $X \otimes X$). Another bit represents a logical parity measurement failure. The last bit encodes the logical Pauli error that affects all logical qubits in the merged patch (such as a logical X error during an $X \otimes X$ measurement).

IV. PROTOCOL FOR TWIST-FREE LATTICE SURGERY

As discussed in earlier sections, lattice surgery provides a fault-tolerant way to measure Pauli operators and is well suited for topological codes. However, not all Pauli operators are equally easy to measure. We say an operator is a XZ -Pauli when it is a tensor product of $\{\mathbb{I}, X, Z\}$. For XZ Pauli operators, a standard lattice surgery approach will suffice and a surface code architecture would need only weight 4 stabilizer measurements. However, for some topological codes such the surface code, measuring Pauli operators containing any Y terms is more difficult. It has been shown that this can be achieved by introducing a twist defect for each Y in the Pauli operator [10, 11, 47]. For an example of twist defect lattice surgery, see Fig. 40(d) of Ref. [11]. However, each surface code twist defect requires a stabilizer measurement on 5 physical qubits. This can be very challenging to implement in a 2D architecture with limited connectivity and could require multiple ancilla qubits. The additional ancilla qubits and gates will thus increase the total measurement failure probabilities for weight-five checks. Furthermore, even a single isolated weight-five check will have an impact on the gate scheduling over the whole surface code patch which can introduce additional types of correlated errors. Lastly, twist-based surface codes coupled with a MWPM decoder have been shown to have a reduced effective code distance [48]. As such, we expect twist-based Pauli measurements will suffer a performance loss relative to twist-free Pauli measurements. Any increases in measurement

error probabilities during twist-based lattice surgery can be suppressed by extending d_m . In other words, we expect use of twists to increase the runtime of lattice surgery computations. The exact magnitude of this runtime cost is currently unknown and will depend on the precise twist implementation details and the noise model [49].

Here we outline an alternative twist-free approach to measuring operators containing Pauli Y terms. The additional cost of supporting Pauli Y terms relative to measuring XZ -Pauli operators is roughly a $2\times$ slowdown in algorithm runtime and a $+2$ additive cost in the number of logical qubits (though we show later that one of these logical qubits can be borrowed from space allocated to routing). Whether this $2\times$ slowdown is preferable to the slowdown incurred by using twists is an open question, due to a lack of data on twist performance.

To explain our protocol, we make use of the notation

$$X[\mathbf{u}] := \prod_{j=1}^N X_j^{u_j} \quad (7)$$

$$Z[\mathbf{v}] := \prod_{j=1}^N Z_j^{v_j} \quad (8)$$

for any binary vectors $\mathbf{u} = (u_1, u_2, \dots, u_N)$ and $\mathbf{v} = (v_1, v_2, \dots, v_N)$. It is well-known that any Hermitian Pauli operator can (up to a ± 1 phase) be decomposed as

$$P = i^{\mathbf{u} \cdot \mathbf{v}} Z[\mathbf{v}] X[\mathbf{u}]. \quad (9)$$

Then $\mathbf{u} \cdot \mathbf{v} = \sum_j u_j v_j$ counts the number of locations where $X[\mathbf{u}]$ and $Z[\mathbf{v}]$ have overlapping support. Therefore, using $X_j Z_j \propto Y_j$, we see that $\mathbf{u} \cdot \mathbf{v}$ gives the number of Y terms in P . Furthermore, $X[\mathbf{u}]$ and $Z[\mathbf{v}]$ commute whenever $\mathbf{u} \cdot \mathbf{v}$ is even. Equivalently, whenever P contains an even number of Y terms, it can be decomposed (up to a phase) into a product of two commuting operators $X[\mathbf{u}]$ and $Z[\mathbf{v}]$.

Let assume for now that $\mathbf{u} \cdot \mathbf{v}$ is even, returning to the odd case later. This suggests that we could measure P by using twist-free lattice surgery to measure $X[\mathbf{u}]$ and $Z[\mathbf{v}]$. However, this would reveal additional unwanted information about $X[\mathbf{u}]$ and $Z[\mathbf{v}]$. To obfuscate this unwanted information, we perform the protocol as illustrated in Fig. 8 and described below

1. Prepare an ancilla (qubit A) in the state $|0\rangle$;
2. Measure $X[\mathbf{u}] \otimes X_A$ with outcome $m_x \in \{0, 1\}$;
3. Measure $Z[\mathbf{v}] \otimes X_A$ with outcome $m_z \in \{0, 1\}$;
4. Return $m_x \oplus m_z \oplus c$ as outcome of $P = i^{\mathbf{u} \cdot \mathbf{v}} X[\mathbf{u}] Z[\mathbf{v}]$;
5. Measure qubit A in the Z basis with outcome $q \in \{0, 1\}$;
6. If $q = 1$, then apply a $Z[\mathbf{v}]$ correction (to the Pauli frame).

In step 4, we use a constant c that we define as follows

$$c = \begin{cases} 0 & \text{if } \mathbf{u} \cdot \mathbf{v} = 0 \pmod{4} \\ 0 & \text{if } \mathbf{u} \cdot \mathbf{v} = 1 \pmod{4} \\ 1 & \text{if } \mathbf{u} \cdot \mathbf{v} = 2 \pmod{4} \\ 1 & \text{if } \mathbf{u} \cdot \mathbf{v} = 3 \pmod{4} \end{cases} \quad (10)$$

Clearly, the product of measurement outcomes in steps 2 and 3 gives $X[\mathbf{u}] Z[\mathbf{v}]$ up to some constant. However, at no point do we learn the value $Z[\mathbf{v}]$ or $X[\mathbf{u}]$, therefore the protocol works as claimed. In Appendix A we provide a more formal proof of the correctness of our protocol and the derivation of the constant c .

We assumed earlier that the Pauli operator P contains an even number of Y terms. To handle odd numbers of Y terms, we prepare an additional ancilla in the Y basis eigenstate $|Y\rangle = (|0\rangle + i|1\rangle)/\sqrt{2}$. Then, by measuring $Y \otimes P$ we effectively measure P . Furthermore, if P contains an odd number of Y terms, then $Y \otimes P$ contains an even number and can be measured using the above construction. This modified variant of the twist-free lattice surgery measurement is also illustrated in Fig. 8. Note that the $Y \otimes P$ measurement does not affect the $|Y\rangle$ state and so it can be reused many times and its preparation cost (e.g. through state distillation [7]) only needs to be paid once per algorithm and is therefore negligible.

Our twist-free approach uses up to two logical ancillas, a $|0\rangle$ ancilla that is repeatedly reset and sometimes a $|Y\rangle$ ancilla that is reused. Therefore, we have an additive $+2$ logical qubit cost. The runtime cost is dominated by steps 2 and 3. All other steps use only single-qubit operations that effectively take zero time in lattice surgery. Therefore, the runtime has doubled compared to the runtime of measuring a Pauli operator free from Y terms. If steps 2 and 3 each fail with probability \mathbb{P} (e.g as in Eq. (1)), then whole protocol fails with probability $\mathbb{P}' = 2\mathbb{P}(1 - \mathbb{P}) \approx 2\mathbb{P}$. This is a minor effect since failure probabilities are exponentially suppressed by increasing the runtime d_m and using large enough d_x and d_z code distances.

In Fig. 9, we show an example of how the circuit picture of Pauli measurements in Fig. 8 can be explicitly mapped into a 2D lattice surgery protocol consisting only of XZ -Pauli operator measurements. We present our protocol for thin rectangular surface codes, though our protocol would also work for square surface code patches. First, we notice how the temporary $|0\rangle$ ancilla is prepared in the spare routing space provided for performing lattice surgery, so that it does not actually contribute to the space overhead. Notice also that to accomplish the $Z[\mathbf{v}] \otimes X_A$ measurement in panel 4, there is one region (highlighted in pink) where we measure elongated stabilizers. Here we have assumed that the hardware is permanently deformed in this region. Alternatively, these elongated stabilizers could also be measured in a homogeneous hardware layout but with a modified procedure for performing the measurement. For instance, one could use two ancilla qubits prepared in a GHZ state to measure the elongated stabilizers. However, since the result of the stabilizer measurement would be given by the product of the measurement outcomes of both ancillas, and due to the extra fault locations, using GHZ states would increase the

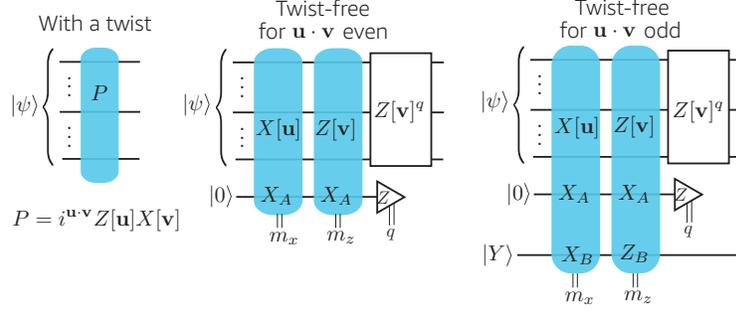


FIG. 8. Implementations of a Pauli measurement of $P = i^{\mathbf{u} \cdot \mathbf{v}} X[\mathbf{u}]Z[\mathbf{v}]$. Whenever P contains any Y terms, simple lattice surgery operations can not be used. The standard solution is to use twist based lattice surgery. However, we show the same outcome can be achieved twist-free with an extra $|Y\rangle$ ancilla used to handle cases where $X[\mathbf{u}]$ and $Z[\mathbf{v}]$ do not commute. The twist-free approaches report $m_x \oplus m_z \oplus c$ as the outcome for the measurement of P where c is a constant determined by \mathbf{u} and \mathbf{v} .

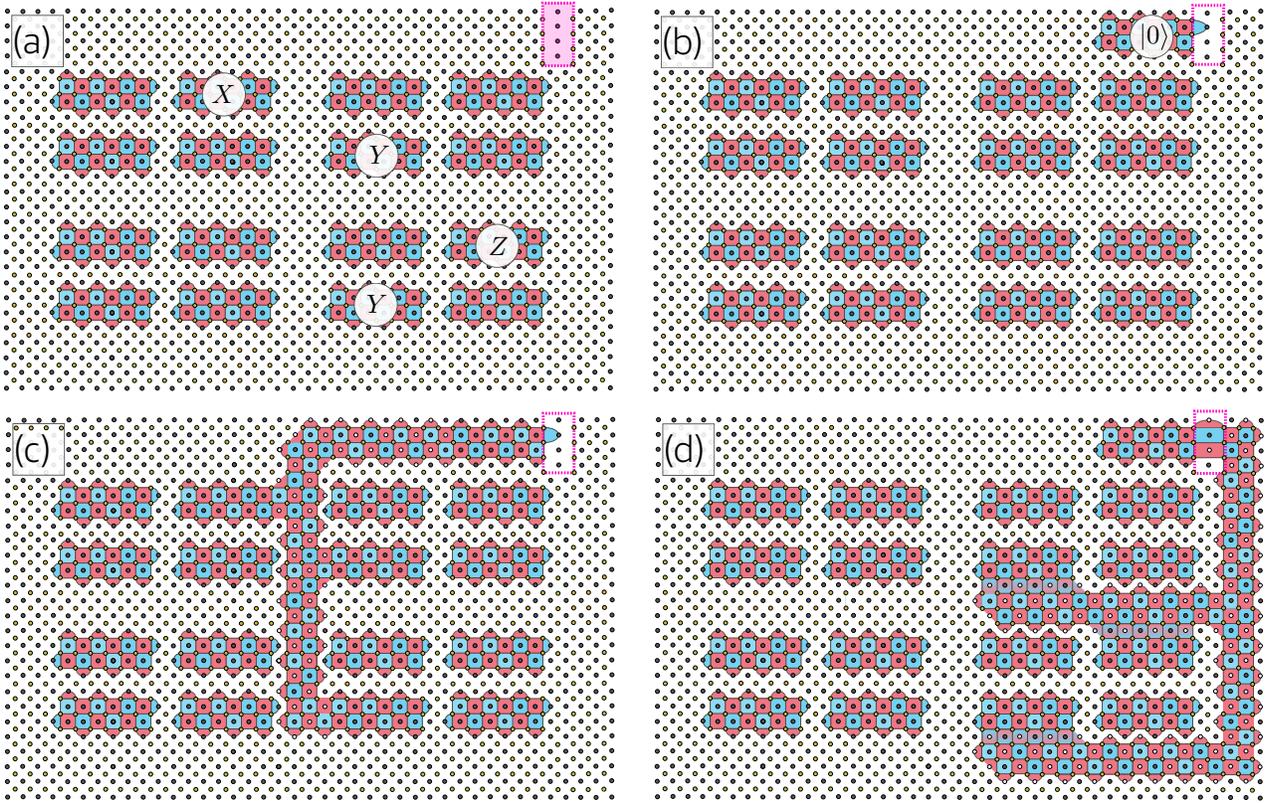


FIG. 9. Example of a 2D layout for a quantum computer and the implementation of twist-free lattice surgery operations to realise a $P = Y \otimes Y \otimes X \otimes Z$ Pauli measurement, for which the circuit diagram is given in Fig. 8. In (a), we show an initial layout of thin, rectangular surface codes ($d_x = 3$ and $d_z = 7$) with labels showing the Pauli operators to be measured. Note the pink rectangle where the hardware layout is slightly adapted to enable elongated stabilizer measurements just within this region. The space between surface code patches is referred to as *routing space* and will be used in the subsequent steps. In (b), we show the preparation of a logical $|0\rangle$ state in the routing space. In (c), we show a lattice surgery measurement of $X \otimes X \otimes X \otimes \mathbb{1} \otimes X$. In (d), we show a lattice surgery measurement of $Z \otimes Z \otimes \mathbb{1} \otimes Z \otimes X$ where at every Z logical boundary we use a dislocation and at the single X logical boundary use elongated stabilizers within the pink region. In both steps (c) and (d), the parity of the logical Pauli measurement is determined by the product of the stabilizers marked with a white vertex, with corrections applied by the decoder.

total measurement failure probability of the elongated checks. Another possibility would be to use the second ancilla qubit as a flag qubit [19, 20, 45, 46, 50–57]. However, by doing so, one might require an additional time step to perform all two-qubit gates for the stabilizer measurements while avoiding scheduling conflicts.

V. TEMPORAL ENCODING FOR FAST LATTICE SURGERY

In Sections II A and II B, we discussed the standard approach to Pauli-based computation using lattice surgery and related algorithm runtime bottlenecks. In this section, we show how to exceed this bottleneck and run algorithms at faster clockspeeds using temporal encoding of lattice surgery (TELS). The key idea behind TELS is to use fast, noisy lattice surgery operations, with this noise corrected by encoding the sequence of Pauli measurements within a classical error correction code. This encoding can be thought of as taking place in the time domain, so the encoding does not directly lead to additional qubit overhead costs. There can be a small additive qubit cost when TELS is used for magic state injection, with magic states needing to be stored for slightly longer times.

A. Parallelizable Pauli measurements

Here we review parallelization where the sequence of Pauli measurements can be grouped into sets of *parallelizable* Pauli measurements. Let $P_{[t,t+k]} := \{P_t, P_{t+1}, \dots, P_{t+k}\}$ be a sub-sequence of our Pauli operators. We say $P_{[t,t+k]}$ is a parallelizable set if: they all commute; any Clifford corrections can be commuted to the end of the sub-sequence. For example, we get a parallelizable set whenever we use magic states to perform a $T^{\otimes k}$ gate. We show in Fig. 1, several ways to implement $T^{\otimes 2}$ with the PBC approach requiring two parallelizable measurements $\{P_1, P_2\} = \{CZ_1Z_3C^\dagger, CZ_2Z_4C^\dagger\}$. Therefore, given a circuit with μ T -gates and T -depth γ , the Pauli measurement sequence can always be split into a sequence of γ parallelizable sets of average size $k := \mu/\gamma$.

Fowler introduced the notion of time-optimal quantum computation [58] and Litinski (see Sec 5.1 of Ref. [11]) showed how this can be realised using lattice surgery in a two-dimensional layout. In time-optimal PBC, an n -qubit computation of T -depth γ can be reduced to runtime $O(n + \gamma)$. However, the space-time volume is never compressed by using the time-optimal approach, so that reducing the algorithm runtime to 10% of a seqPBC runtime would require at least a $10\times$ increase in qubit cost. Litinski worked through some highly-parallelizable examples in greater detail, showing: a reduction to 56.5% of seqPBC runtime would need over $6\times$ the qubit costs; and a reduction to 11% of seqPBC runtime would need over $20\times$ the qubit cost. The qualifier “over” in these estimates reflects that an increase in space-time volume will also increase the code distance needed, further increasing the overhead of the time optimal approach by a polylogarithmic factor on top of Litinski’s estimates. While this is a powerful

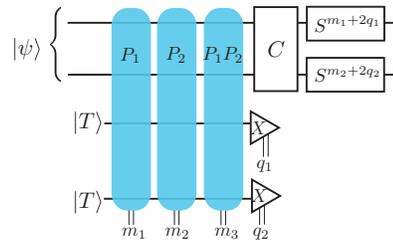


FIG. 10. A simple TELS protocol: it is an error detecting version of the PBC used in Fig. 1 using the measurement code given in Eq. (14). While this approach uses 3 multi-qubit Pauli measurements, the capability to detect an error means that lattice surgery can be executed over a shorter time d_m . If there is no error, we will have that the measurement outcomes obey $m_1 \oplus m_2 = m_3$. When an error is detected we simply remeasure the Pauli operators.

approach to exploring space-time tradeoffs, early fault-tolerant quantum computers will be qubit limited. Kim *et al.* [59] also proposed another way to exploit large parallelizable sets, but they used long-range gates not possible in 2D hardware and also made some strong assumptions regarding the speed at which transversal gates can be fault-tolerantly applied.

From the above, we see that it is crucial to understand the extent to which algorithms possess potential for parallelization. Fortunately Kim *et al.* [59] have already studied quantum algorithms for chemistry and found k to vary between 9 and 14 depending on the orbitals used, so we will regard this as a practically reasonable range.

B. Encodings and code parameter proofs

Here we introduce our own approach to exploiting a parallelizable Pauli sets. Unlike, previous time-optimal approaches, it does not incur a multiplicative qubit overhead cost and can reduce the overall space-time cost.

Due to the properties of a parallelizable Pauli set, all Pauli operators within the set can be measured in any order. Furthermore, we can measure any set \mathcal{S} that generates the group $\langle P_t, P_{t+1}, \dots, P_{t+k} \rangle$. If the set \mathcal{S} is overcomplete, there will be some linear dependencies between the measurements that can be used to detect (and correct) for any errors in the lattice surgery measurements. For example, consider the simplest parallelizable set $\{P_1, P_2\}$ as in Fig. 1 and let d_m be the required lattice surgery time, so performing both measurements takes $2(d_m + 1)$ error correction cycles. We could instead measure $\{P_1, P_2, P_1P_2\}$. If the 3rd measurement outcome is not equal to the product of the first two measurements, then we know something has gone wrong and can repeat the measurements to gain more certainty of the true values. By measuring the overcomplete set $\{P_1, P_2, P_1P_2\}$ we have performed an extra lattice surgery measurement, but we have gained that we can tolerate a single lattice surgery failure. This means that we could instead use $d'_m \ll d_m$ and still achieve the same overall success probability. If $3d'_m \ll 2d_m$ then the computation has been sped-up. This is the key insight behind temporal encoding

of lattice surgery (TELS) and next we dive deeper into more general encoding schemes and their performance.

In general, given a parallelizable Pauli set

$$\mathcal{P} = \{P_t, P_{t+1}, \dots, P_{t+k}\} \quad (11)$$

we can define operators generated from this set as follows

$$Q[\mathbf{x}] := \prod_{j=0}^{k-1} P_{t+j}^{x_j} \quad (12)$$

where \mathbf{x} is a length k binary column vector. Given a set that generates all the required Pauli operators, so $\langle \mathcal{S} \rangle = \langle \mathcal{P} \rangle$ we can write the elements as

$$\mathcal{S} = \{Q[\mathbf{x}^1], Q[\mathbf{x}^2], \dots, Q[\mathbf{x}^n]\} \quad (13)$$

with superscripts denoting different vectors. Since this is a generating set, the vectors $\{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n\}$ must span the relevant space. Furthermore, we can define a matrix G with these vectors as columns and this matrix will specify the TELS protocol. In the simple $k = 2$ example where $\mathcal{S} = \{P_1, P_2, P_1 P_2\}$, we would have that

$$G = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix} = (\mathbf{x}^1 \quad \mathbf{x}^2 \quad \mathbf{x}^3) \quad (14)$$

Notice that the rows of this matrix generate the codewords of the $[3, 2, 2]$ classical code. In general, we will consider G as the generator matrix for the codewords of an $[n, k, d]$ classical code and we call this the *measurement code* for the protocol. Note that k is the number of (unencoded) Pauli operators in the generating set. We only consider full-rank G where k equals the number of rows in G . The number n represents how many Pauli measurements we physically perform in the encoded scheme and corresponds to the number of columns in G . The distance d is the lowest weight vector in the row-span of G .

Next we show that the code distance d does indeed capture the ability of TELS to correct errors. First, we formalise the redundancy in the set of lattice-surgery measurements. For any length n binary vector $\mathbf{u} = (u_1, u_2, \dots, u_n)$, we have that

$$\prod_{j:u_j=1} Q[\mathbf{x}^j] = Q\left[\sum_l u_l \mathbf{x}^l\right] \quad (15)$$

Since the matrix G is full-rank and has more columns than rows, there will exist \mathbf{u} such that $\sum_j u_j \mathbf{x}^j = 0$. For these \mathbf{u} , we have that

$$\prod_{j:u_j=1} Q[\mathbf{x}^j] = \mathbb{1} \quad (16)$$

Therefore, these \mathbf{u} vectors describe redundancy in the measurements. The condition $\sum_j u_j \mathbf{x}^j = 0$ can be rewritten compactly as $G\mathbf{u} = 0$. Following the convention in coding theory, this set of \mathbf{u} is called the dual of G and denoted

$$G^\perp := \{\mathbf{u} : G\mathbf{u} = 0 \pmod{2}\} \quad (17)$$

Next, we consider how this redundancy is used to detect time-like lattice surgery errors. We let $\mathbf{m} = \{m_1, m_2, \dots, m_n\}$ be a binary vector denoting the outcomes of the lattice surgery Pauli measurements in the set \mathcal{S} . That is, if a measurement of $Q[\mathbf{x}^j]$ gives outcome “+1” we set $m_j = 0$ and when the measurement of $Q[\mathbf{x}^j]$ gives “-1” we set $m_j = 1$. Given a $\mathbf{u} \in G^\perp$, we know the Pauli operators product to the identity (recall Eq. (16)) so when there are no time-like lattice surgery errors we have

$$\prod_{j:u_j=1} m_j = \mathbf{u} \cdot \mathbf{m} = 0 \pmod{2}. \quad (18)$$

Conversely, if we observe

$$\prod_{j:u_j=1} m_j = \mathbf{u} \cdot \mathbf{m} = 1 \pmod{2} \quad (19)$$

then we know a time-like lattice surgery error must have occurred. Let us write $\mathbf{m} = \mathbf{s} + \mathbf{e}$ where \mathbf{s} is the ideal measurement outcome and \mathbf{e} is the measurement error. The ideal measurement outcomes are self-consistent and so always satisfy $\mathbf{u} \cdot \mathbf{s} = 0$ for all $\mathbf{u} \in G^\perp$. Therefore, we see that an error \mathbf{e} is undetected if and only if $\mathbf{u} \cdot \mathbf{e} = 0$ for some $\mathbf{u} \in G^\perp$. This is equivalent to undetected errors \mathbf{e} being in the row-span of G (since the dual of the dual is always the original space). Recall, the distance d denotes the lowest (non-zero) weight vector in the row-span of G . Therefore, d also denotes the smallest number of time-like lattice surgery errors needed for them to be undetected by TELS. Consequently, if \mathbb{P} is the probability of a single timelike error, TELS error-detection will fail with probability $O(\mathbb{P}^d)$.

Matrices such as G also appear in the literature in the context of measuring overcomplete sets of stabilizers for some quantum error correction code. In the error correction setting, these codes have been called measurement codes [60], meta-checks [61, 62], symmetries [63] and syndrome measurement codes [64]. However to the best of our knowledge, this is the first time that this idea has been deployed in the context of lattice surgery as a strategy to improve algorithm runtimes.

C. Examples and numerics

The simplest examples of TELS will detect a single error. Given a Pauli set $\{P_t, P_{t+1}, \dots, P_{t+k}\}$ we measure each of these observables separately, and then the product of them so that the measurement code has generator matrix

$$G = \begin{pmatrix} 1 & 0 & \dots & 0 & 1 \\ 0 & 1 & \dots & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 1 \end{pmatrix} \quad (20)$$

which is an identity matrix padded with an extra column that is an all 1 vector. Therefore, this corresponds to a $[\alpha + 1, \alpha, 2]$ classical code that detects a single error. Concatenating such a code m times gives a code with parameters $[(\alpha + 1)^m, \alpha^m, 2^m]$.

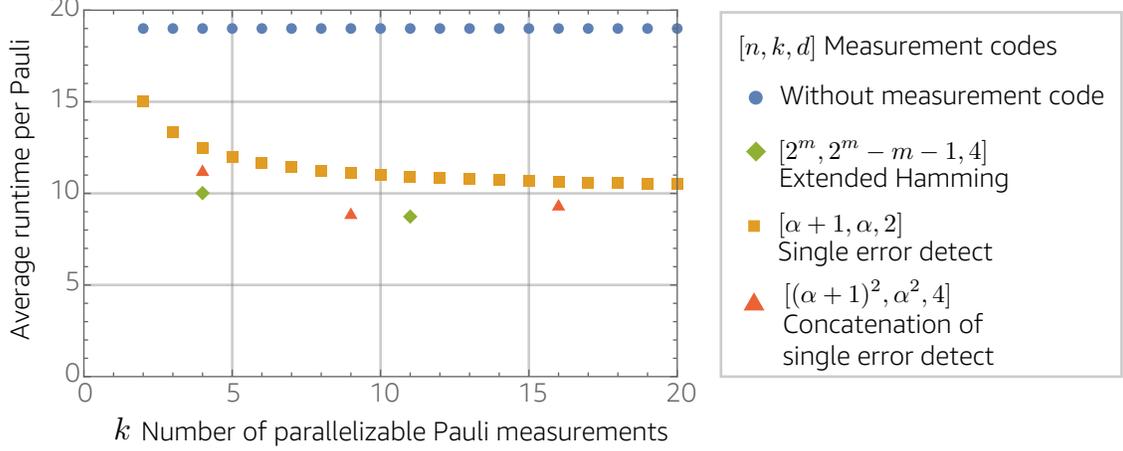


FIG. 11. A comparison of the standard lattice surgery approach (without the measurement code) with 3 different temporally encoded lattice surgery (TELS) schemes for performing a set of k parallelizable Pauli measurements. We assume that Pauli measurements give an incorrect outcome due to a time-like failure with a probability determined by Eq. (3) with $d_x \ell = 100$ and $p = 10^{-3}$. We set the allowed error per Pauli at $\delta = 10^{-15}$. The $[2^m, 2^m - m - 1, 4]$ are well-known Hamming codes. The $[\alpha + 1, \alpha, 2]$ are single error detection codes, and $[(\alpha + 1)^2, \alpha^2, 4]$ are concatenated single error detection codes. While there are big jumps in k between the best performing codes, these jumps could be partially smoothed out by considering other codes such as concatenated codes with different inner/outer code sizes, such as the $[(\alpha + 1)(\beta + 1), \alpha\beta, 4]$ codes. We also considered the triply concatenated codes with parameters $[(\alpha + 1)^3, \alpha^3, 8]$ but they performed poorly in the parameter regime shown here and so have been omitted for clarity.

We can also consider using a simple $[8, 4, 4]$ extended Hamming code as the measurement code, with generator matrix

$$G = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad (21)$$

This corresponds to replacing $\{P_1, P_2, P_3, P_4\}$ with \mathcal{S} containing the 8 operators

$$\mathcal{S} = \{P_2 P_3 P_4, P_2 P_3, P_2 P_4, P_2, P_1 P_3 P_4, P_1 P_3, P_1 P_4, P_1\} \quad (22)$$

Because the generator matrix has distance 4, this scheme will detect up to 3 errors. This Hamming code is the $m = 3$ member of a family of $[2^m, 2^m - m - 1, 4]$ extended Hamming codes.

There are several viable strategies to handle a detected error. Here we consider the following detect/remeasure strategy: if a distance d measurement code is used with lattice surgery performed for time d_m , then whenever an error is detected we “remeasure” but this time using the original Pauli set \mathcal{P} instead of using the overcomplete set \mathcal{S} . For the remeasure round, we perform lattice surgery using an amount of time $[q d_m]$ where q is some constant scaling factor whose value we discuss shortly. The expected runtime to execute the protocol is then

$$T = n(d_m + 1) + p_d k d d_m \quad (23)$$

where p_d is the probability of detecting an error. When we do not detect an error, the probability of a failure is

$O(p^{d(\frac{d_m+1}{2})}) \approx O(p^{d d_m/2})$. When we do detect an error, the remeasure round will fail with probability $O(p^{(q d_m + 1)/2}) \approx O(p^{q d_m/2})$. The total failure probability will then be $O(p^{d d_m/2} + p_d p^{q d_m/2})$. Therefore, we can ensure the total failure probability is $O(p^{d d_m/2})$ by setting $q = d$. However, due to constant factors the optimal choice of q maybe slightly different from $q = d$ and so we numerically optimize from this initial guess. When an error detection occurs, this leads to a long delay of time $k d d_m$ to implement the remeasure round, but in practice p_d is so small that this has minimal impact on the expected runtime.

We could alternatively just measure the overcomplete set \mathcal{S} and run the measurement code in error correction mode with lattice surgery repeated for time d'_m . Then the protocol would fail with probability $\approx O(p^{d d'_m/4})$. Compared to detect/remeasure strategy, we need $d'_m \approx 2 d_m$ to achieve the same failure probability. The runtime of an error correction scheme will then be

$$T' = n(d'_m + 1) \approx n(2 d_m + 1). \quad (24)$$

Compared to Eq. (23), in T' we have dropped the second term at the price of roughly doubling the first term. However, the second term was small because p_d is small, so overall error correction is not favourable compared to our detect/remeasure scheme.

Fig. 11 shows some example numerical results using distance 2 and 4 codes. For example, when performing $k = 11$ parallelizable gates the TELS scheme using extended Hamming codes will have a runtime of 46% that of a standard fast seqPBC approach that measures the original parallelizable

Pauli set \mathcal{P} . Since Kim *et al.* [59] found interesting quantum algorithms can have average k between 9 and 14, this suggests around a $2.2\times$ speedup due to TELS on practical problems. To obtain a similar speedup, Litinski estimated the time-optimal approach would cost over $6\times$ in qubit overhead. The TELS scheme has no multiplicative qubit overhead, though it does have a small additive qubit overhead as all k magic states must be stored for the full duration of the protocol. However, fault-tolerant algorithms typically have $N \gg 100$ logical qubits, and so the increase in logical qubits $N \rightarrow N + k$ is small. Indeed, overall the spacetime volume will decrease, which is impossible using the time optimal approach.

We did not find any examples of higher distance codes ($d > 4$) that performed better in this parameter regime (e.g. $\delta = 10^{-15}$). Going to even lower error rates ($\delta \ll 10^{-15}$) or changing the noise model, then higher distance codes will become useful and the advantage will improve further. Indeed, because of the existence of good classical codes with $n/k = O(1)$ and $d = \Omega(k)$, we know that TELS will asymptotically (for large k) be able to execute k parallelizable Pauli measurements in $O(k)$ time and with error $\delta = O(p^{\alpha k})$ for some constant α . In contrast, the standard approach of seqPBC with unencoded lattice surgery would take runtime $O(k \text{polylog}(k))$ to achieve the same error.

VI. THE CORE-CACHE ARCHITECTURE AND ROUTING OVERHEADS

In this section, we discuss a layout and data access structure for a quantum computer that extends on the layout given in panel 1 of Fig. 9. We will consider patches of (possibly rectangular) surface codes of size d_z by d_x . Between these patches we will have some qubits dedicated as a lattice surgery “bus” or routing space. We say the routing space supports fast access if logical X and Z operators of every patch are adjacent to the routing space. Litinski proposed several data access structures [11], with his fast data structures using two-tile two-qubit patches (surface code patches that each encode two logical qubits) that are sometimes called hexon surface codes.

In Section VIA, we show that the hexon approach is not necessary, and give a layout for a quantum core (what Litinski calls a fast data access structure). Furthermore, we show that a lower routing overhead is possible when the surface code patches are thin rectangles (e.g. $d_z \gg d_x$) as is the optimal choice for highly biased noise. In Section VIB, we will also discuss the idea of a core-cache model where logical qubits are temporally moved out of the fast data access structure to reduce the routing overhead.

A. A quantum computer core

We count resource costs in terms of square tiles as defined in Fig. 12a. Each tile contains a single data qubit and four quarter ancilla qubits. Therefore, a device with T tiles will require roughly $2T$ qubits. However, we cannot cut qubits into quarters and so a precise counting will include these. For instance, a

rectangular device with a height of h tiles and width of w tiles would have a total of $T = wh$ tiles and $2T + w + h + 1$ qubits. When the device is roughly square, then h and w are of size $O(\sqrt{T})$ and so a negligible additive cost compared to $2T$. We can realise a surface code patch using $d_x d_z$ tiles as in Fig. 12b and therefore $2d_x d_z$ qubits. The number of data and ancilla qubits actively used in the surface code patch is $2d_x d_z - 1$, and so when we try to pack them in a 2D arrangement, the tightest possible packing will contain one idling qubit per patch. However, in practice, we also use qubits for routing.

We collect surface code patches into groups of four, which we call a unit cell (see Fig. 12c). These unit cells are then repeated as shown in Fig. 12d to get the required number of logical qubits. In addition we arrange the unit cells to form a quantum “core”, and assume some additional padding shown in Fig. 12e. Notice that in Fig. 12, every patch has logical X and Z boundary operators connected to the routing space, which enables us to quickly perform multi-qubit Pauli measurements between any subset of qubits within the core. Additionally, there are unused qubits between some of the surface code patches. The spacing of the qubits ensures that lattice surgery can be performed (as we saw in Fig. 9) without using lattice twists that incur additional practical difficulties to implement in fixed and low connectivity hardware. In contrast, the data access structures proposed by Litinski [11] assumed liberal use of twists.

The routing overhead for unit cells is then the ratio of the number of tiles divided by the cost without any routing space (e.g. $4d_z d_x$).

$$O_{(d_z, d_x)}^{\text{(unit cell)}} = \frac{(2d_z + d_x + 1)(3d_x + 1)}{4d_z d_x}. \quad (25)$$

The overhead for the entire core will include a contribution from the additional padding shown in Fig. 12e. However, in the limit of many unit cells, the total overhead is dominated by the unit cell overhead. In the limit of large distances $d_z, d_x \gg 1$, we have

$$O_{(d_z, d_x)}^{\text{(unit cell)}} \approx \frac{3}{2} + \frac{3}{4} \frac{d_x}{d_z}. \quad (26)$$

Therefore, in the limit of large noise bias, $d_z \gg d_x$, the routing overhead factor is $1.5\times$. We can compare this with the fast data blocks of Litinski (see e.g. Figure 23 of [11]) where the overhead factor is $2\times$ (in large device limit) and so is more expensive. In contrast, for unbiased noise and $d_z = d_x$ our scheme has an asymptotic routing overhead factor of $9/4$, which is slightly worse than Litinski’s $2\times$ routing overhead. Indeed, solving Eq. (26) equal to 2 shows that $d_x < (2/3)d_z$ is the condition for our approach to have a routing overhead advantage. A more general analysis of the routing overhead which includes the green padded regions shown in Fig. 12e and contributions from the cache (see Section VIB) is given in Appendix C. We also point out that routing overhead is not the only important figure of merit. Our proposed design avoids twists and twist defects, and therefore may be useful even without noise bias.

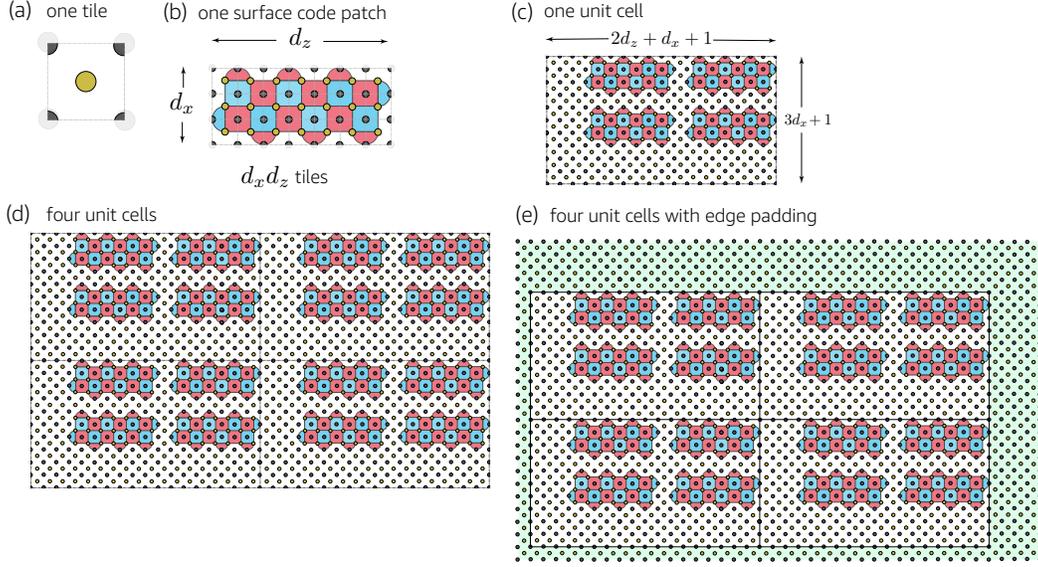


FIG. 12. Building blocks for a quantum computer. (a) A single tile containing 2 qubits, ancilla qubits (cut at the four corners of the tile) are shared with neighbouring tiles. (b) A single (asymmetric) surface code patch. The number of tiles needed is the height times width, and so $d_x d_z$. (c) A single unit cell containing four surface code patches and some routing space. Notice all surface code patches have a X and Z logical boundary adjacent to either routing space within the bulk or the boundary of the unit cell. (d) Four unit cells tiled together. (e) Same as (d) but with the inclusion of extra padding highlighted in green at the edges to provide access to the remaining X and Z logical boundaries and ample routing space. There is more padding on the top and right since we need room the access the remaining X and Z logical operators.

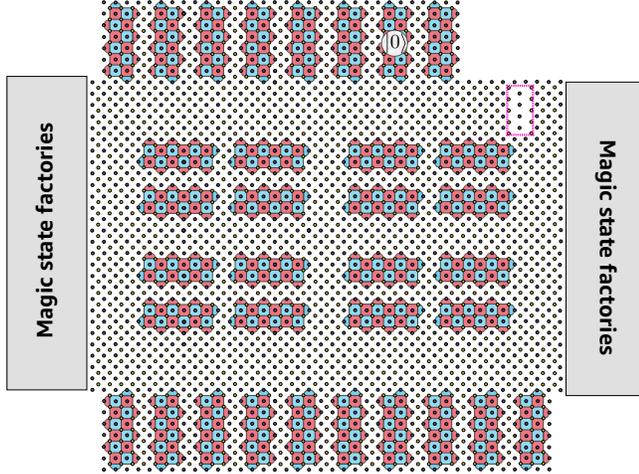


FIG. 13. A small example of a quantum computer layout, with a core composed of 4 unit cells (and therefore 16 surface code patches); cache on the top/bottom edge of the core; and magic state factories on the left/right of the core. A pink, dashed rectangle in the top-right corner indicates a hardwired lattice defect so to enable twist-free lattice surgery.

B. A quantum computer cache

We now proceed to build additional structure around the core. Using state distillation to prepare magic states, we will need factories that supply the core. The purely fast data access

approach prioritises speed over qubit cost. Here we also discuss the idea of a core and cache architecture, where some logical qubits are temporarily stored in a quantum analog of cache. However, with some time cost, logical qubits can be quickly swapped in and out of the cache. A small-scale sketch of a device comprising core, cache and magic state factories is illustrated in Fig. 13.

Packing qubits more compactly in the cache will clearly reduce the overhead costs. However, such a layout comes at a price since only the X logical operators of these qubits can be accessed when it is in the cache. To access the logical Z operators, it must be swapped out of the cache and into the core. For a qubit stored in the cache, we can perform the following operations

1. Perform single qubit X or Z measurements for a qubit in the cache (time cost: zero).
2. Measure multi-qubit operators of the form $A \otimes B$ where B acts on the cache qubits and is a tensor product of X operators only, and A acts on the core qubits and can be an arbitrary Pauli operator. (time cost: d_m).
3. Perform Pauli updates to the Pauli frame (in software).
4. Perform Cliffords to qubits in the core, by updating the Clifford frame (in software).

For algorithms where swapping in/out of the cache can be made infrequently (compared to other time costs), our approach will reduce the routing overhead with a mild impact on the algorithm runtime. Note also that our core-cache architecture can

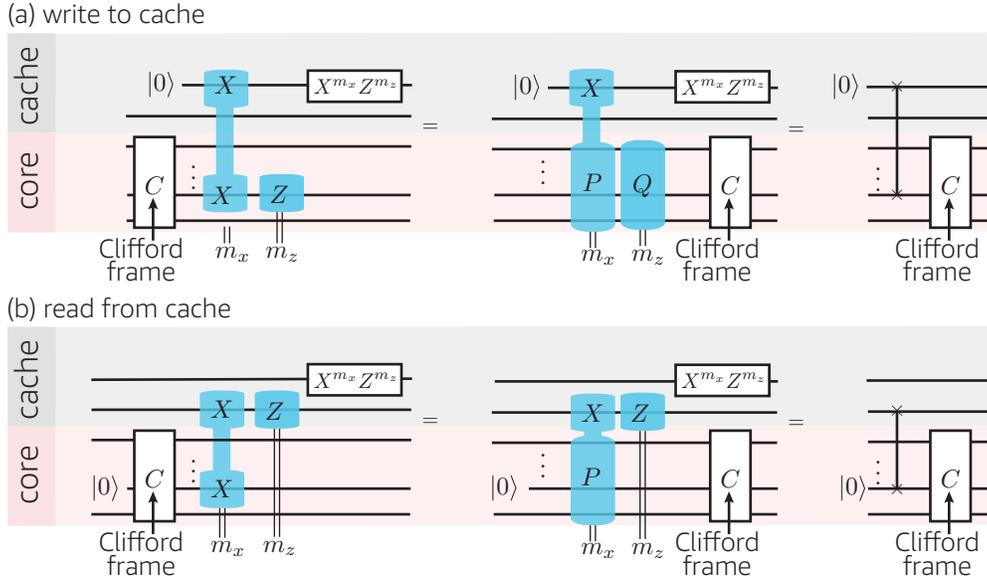


FIG. 14. Circuit diagrams for the operations (a) write to cache (WTC); and (b) read from cache (RFC).

be used in combination with the twist-free or TELS schemes already proposed.

This leaves the question of how to swap the location of qubits from the cache to the core. We cannot directly implement the Clifford SWAP operation, since Clifford operations can only be performed on core qubits. Furthermore, surface code patches cannot be moved around to swap their positions since such operations would require the Clifford frame C to be relabelled. Such a relabelling might make C act non-trivially on qubits in the cache. Rather, when performing a SWAP from a qubit in the core to the cache, we need to clean the Clifford frame so it only acts on core qubits.

We now define two elementary operations: write to cache (WTC) and read from cache (RFC), which when combined will enable a Clifford-cleaned swap. We first present the WTC and RFC protocols using circuit diagrams in Fig. 14. They are both essentially 2-qubit teleportation protocols, but with the Clifford frame adapting the Pauli measurements performed. The WTC operation uses two multi-qubit Pauli measurements, whereas RFC can be performed faster as it uses only one multi-qubit Pauli measurement and a single-qubit Pauli measurement (that takes zero time). The WTC operation requires a logical $|0\rangle$ ancilla in the cache, which through the protocol swaps place with a logical qubit in the core. The RFC operation requires a logical $|0\rangle$ ancilla in the core, which through the protocol swaps place with a logical qubit in the cache.

For a pair of logical qubits, one in the core and one in the cache, we can cleanly SWAP their positions, by executing WTC followed by RFC. The $|0\rangle$ ancilla initially in the cache, will move to the core, then back to the cache, but to a different location than it started. The whole swap procedure requires three multi-qubit Pauli measurements. We show an example of how this can be realised using lattice surgery in Fig. 15. This figure shows a simple scenario where the Clifford frame conveniently allows these three multi-qubit Pauli measurements to be

performed using three lattice surgery operations and no twists. However, more generally, we need to either: use twists (and benchmark their performance); or use our twist-free protocol and realise the swap with up to six lattice surgery operations.

In Appendix C, we perform a resource cost analysis for simulating the Hubbard model using the full core-cache architecture described in this section. In particular, we provide a rigorous analysis of routing overhead costs including contributions from the cache and green padded region in Fig. 12e.

VII. CONCLUSION

In Section III we introduced a decoding algorithm compatible with lattice surgery protocols and numerically computed failure rate polynomials for the dominant failure mechanisms of an $X \otimes X$ Pauli measurement. Our analysis allows one to compute appropriate d_x and d_z code distances, as well as the number of syndrome measurement rounds d_m during lattice surgery for successfully implementing algorithms.

In Section IV, we introduced a twist free protocol for measuring arbitrary Pauli operators using lattice surgery. The protocol incurs a $2\times$ slowdown in algorithm runtime. However surface codes with twists require higher weight stabilizer measurements, increased gate scheduling complexities and have a reduced effective code distance when using a MWPM decoder. Such features will inevitably cause a reduction in performance compared to lattice surgery protocols involving only X and Z Pauli measurements. Consequently, a careful numerical analysis with twists is needed in order to determine whether using twists can beat the $2\times$ cost of the twist-free approach.

In Section V we introduced a technique which we call temporal encoding of lattice surgery. By encoding lattice surgery measurements in the time domain, we showed that a $2\times$ reduction in algorithm runtimes can be achieved for quantum

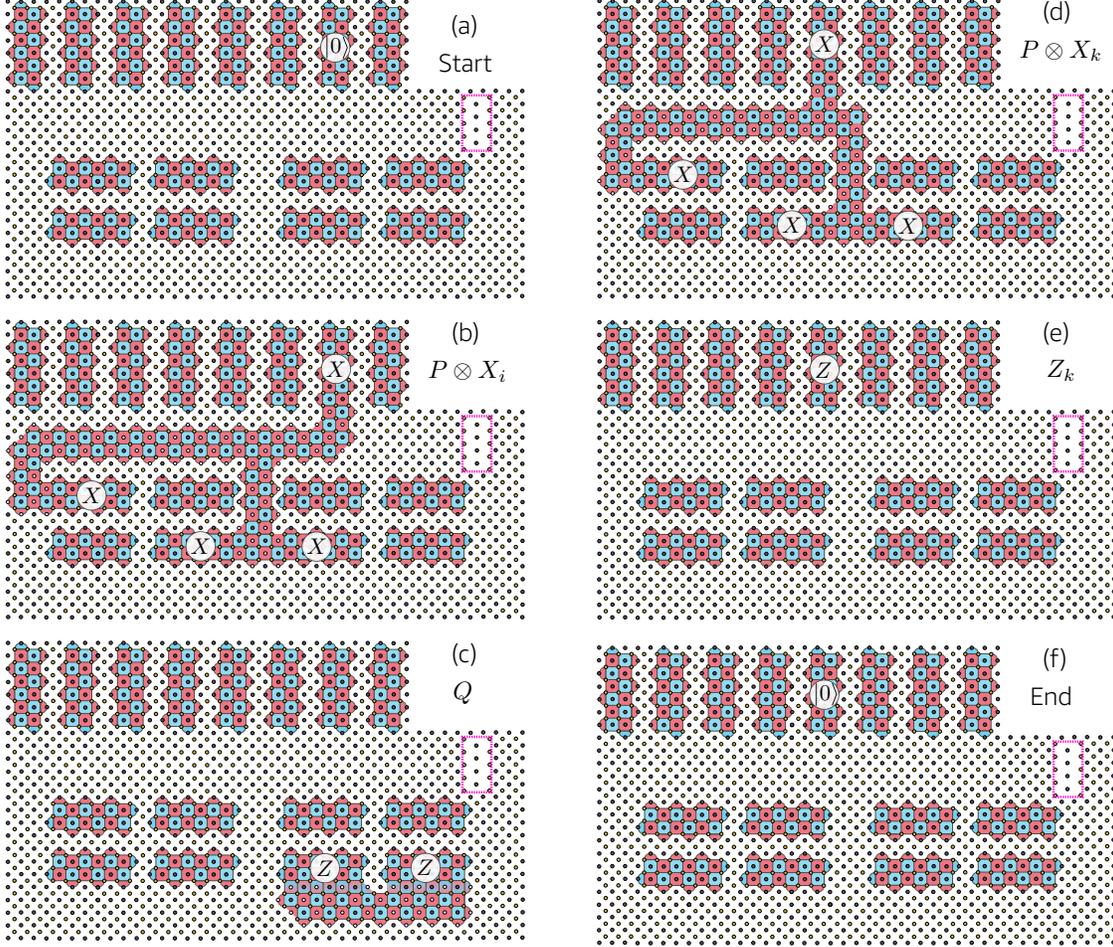


FIG. 15. Example lattice surgery diagrams for: using write to cache and read from cache to exchange qubits j and k between core and cache. (a) Illustration of the initial configuration with a $|0\rangle$ ancilla (index label i) in the cache. Qubit j is in the core and qubit k is in the cache and we wish to swap their locations. (b) We measure $P \otimes X_i$ where $P = CX_jC^\dagger$ where C is the Clifford frame. For simplicity, we assume P is composed of only X operators. (c) We measure $Q = CZ_jC^\dagger$ and this time assume it is composed of only Z operators. (d) We measure $P \otimes X_k$ where $P = CX_jC^\dagger$ is the same operator as in (b). (e) We measure the single qubit Pauli Z_k . If either of the simplifying assumptions we made for P or Q do not hold, then we will need to use the twist-free protocol (or use twists). (f) At the end of the protocol, qubit j is in the cache, qubit k is in the core, and there is a $|0\rangle$ ancilla in the cache ready for further swaps.

algorithms of practical scale without incurring additional qubit overhead costs. For more highly parallelizable algorithms or larger algorithms, using larger classical code distances and exploring other code families will lead to even greater improvements in algorithm runtimes.

Lastly, in Section VI, we provided a core-cache architecture compatible with our new lattice surgery protocols. A subset of the data qubits are stored in a cache, which reduces the footprint of the routing space, and can be quickly accessed when Z measurements need to be performed. We found that for such an architecture, routing overhead costs add a factor of $1.5\times$ to the total resource costs for performing lattice surgery. A clear direction for future work would be to analyze such architectures in the presence of twists in order to better understand the tradeoffs with using a twist-free lattice surgery protocol.

Apart from considering twists, a direction of future work

would be to apply our methods using other error correcting codes to potentially achieve lower resource costs. Promising code candidates include codes tailored for biased noise such as the XZZX surface codes [63, 65–67], subsystem codes with high thresholds [68] and other codes families with high encoding rates such as hyperbolic surface codes [69, 70].

VIII. ACKNOWLEDGMENTS

We thank Giacomo Torlai for his help with using the AWS clusters. We thank Oscar Higgott, Fernando Brandao and Noah Shutty for their comments and suggestions on our manuscript.

Appendix A: Twist-free proof

Here we give a formal proof that the twist-free lattice surgery protocol works as claimed. Consider the case when step 5 yields a $q = 0$ outcome so we project onto the $|0\rangle$ state. Then steps 1-5 implement

$$M_{+1} := |0\rangle\langle 0|_A \Pi_{ZX}(m_z) \Pi_{XX}(m_x) |0\rangle\langle 0|_A, \quad (\text{A1})$$

where

$$\Pi_{ZX}(m_z) := \frac{1}{2} (\mathbb{1} \otimes \mathbb{1} + (-1)^{m_z} Z[\mathbf{v}] \otimes X_A) \quad (\text{A2})$$

$$\Pi_{XX}(m_x) := \frac{1}{2} (\mathbb{1} \otimes \mathbb{1} + (-1)^{m_x} X[\mathbf{u}] \otimes X_A) \quad (\text{A3})$$

Using that for arbitrary Q

$$\begin{aligned} |0\rangle\langle 0|_A (Q \otimes \mathbb{1}) |0\rangle\langle 0|_A &= Q \otimes |0\rangle\langle 0|_A, \quad (\text{A4}) \\ |0\rangle\langle 0|_A (Q \otimes X_A) |0\rangle\langle 0|_A &= 0, \end{aligned}$$

we deduce

$$M_{+1} = \frac{1}{4} [\mathbb{1} + (-1)^{m_x + m_z} (Z[\mathbf{v}] X[\mathbf{u}])] \otimes |0\rangle\langle 0|_A, \quad (\text{A5})$$

which is proportional to the projector for a $Z[\mathbf{v}]X[\mathbf{u}]$ measurement with outcome $m_x \oplus m_z$. When $\mathbf{u} \cdot \mathbf{v} = 0 \pmod{4}$, we have $P = Z[\mathbf{v}]X[\mathbf{u}]$ and so $m_x \oplus m_z$ is the outcome of measuring P (justifying $c = 0$ in this case). On the other hand, if $\mathbf{u} \cdot \mathbf{v} = 2 \pmod{4}$, we have $P = -Z[\mathbf{v}]X[\mathbf{u}]$ and so $m_x \oplus m_z \oplus 1$ is the outcome of measuring P (justifying $c = 1$ in this case). Recall that we are currently assuming $\mathbf{u} \cdot \mathbf{v}$ is even.

In the event that step 5 yields a $q = 1$ outcome, we have that

$$M_{-1} := |1\rangle\langle 1|_A \Pi_{ZX}(m_z) \Pi_{XX}(m_x) |0\rangle\langle 0|_A, \quad (\text{A6})$$

Using that for arbitrary Q

$$\begin{aligned} |1\rangle\langle 1| (Q \otimes \mathbb{1}) |0\rangle\langle 0| &= 0, \quad (\text{A7}) \\ |1\rangle\langle 1| (Q \otimes X_A) |0\rangle\langle 0| &= Q \otimes |1\rangle\langle 0|. \end{aligned}$$

We have

$$M_{-1} = \frac{1}{4} [(-1)^{m_z} Z[\mathbf{v}] + (-1)^{m_x} X[\mathbf{u}]] \otimes |1\rangle\langle 0|_A, \quad (\text{A8})$$

$$= ((-1)^{m_z} Z[\mathbf{v}]) \otimes X_A M_{+1}. \quad (\text{A9})$$

Therefore, we see that M_{-1} differs from M_{+1} by a $Z[\mathbf{v}]$ correction that we perform in step 6. There is also a global phase $(-1)^{m_z}$ but this is unimportant.

Appendix B: Treating space-time correlated edges incident to parity vertices

As mentioned in Section III, the space-time correlated edges incident to the parity vertices during the first round of the merged surface code patches (i.e. vertices $v \in V_{\text{par}}^{(r+1)}$) need to

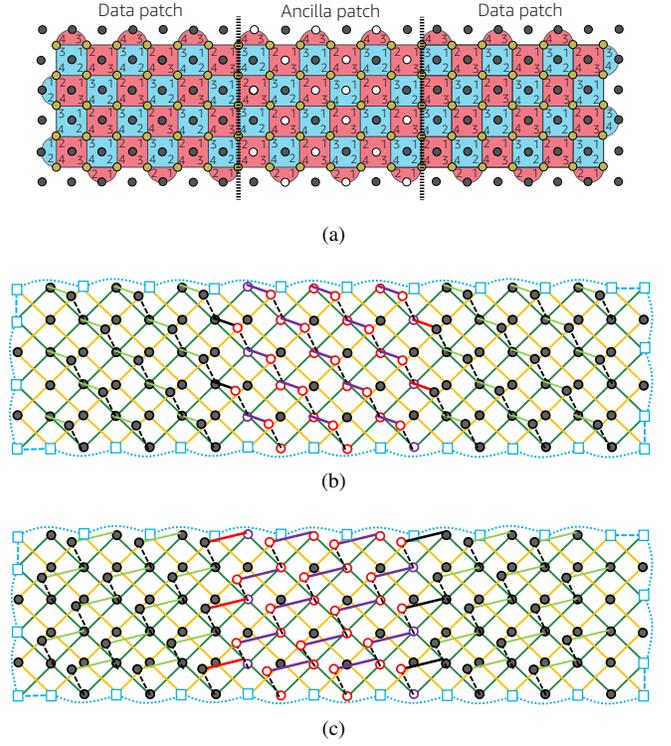


FIG. 16. (a) Surface code lattice with $d_x = 5$, $d_z = 7$. The numbers in each check indicate the gate scheduling of the two-qubit gates (which in our simulations are CNOT gates). (b) First type of space-time correlated edges for X -stabilizers. (c) Second type of space-time correlated edges for X -stabilizers. Space-time correlated edges incident to vertices on the ancilla patch region are assigned different colors than green for the reasons described in the text.

be treated with care. Such edges are highlighted in red, black and purple in Figs. 16b and 16c for an $X \otimes X$ measurement.

Firstly, in round $r + 1$, edges highlighted in purple must have infinite weight and can thus be removed. This is due to the fact that when the data surface code patches are merged with the ancilla patch in round $r + 1$, individual X -stabilizer measurements on the ancilla patch (whose ancilla's are marked by white vertices) will have a random outcomes and thus cannot be highlighted. As such, failures arising from two-qubit gates on the ancilla patch, and which introduce errors that do not interact with data patch stabilizers, cannot generate a non-trivial measurement outcome between two consecutive rounds of stabilizer measurements. Although the purple edges discussed above must be removed when they are incident to vertices in rounds $r + 1$ and $r + 2$, such edges incident to vertices v_i and v_j belonging rounds greater than $r + 1$ must be included since they will have finite weights.

Secondly, space-time edges incident to a single parity vertex $v \in V_{\text{par}}^{(r+1)}$ are highlighted in red in Figs. 16b and 16c. Note that in round $r + 1$, such edges are incident to parity vertices colored in purple, which are also transition vertices. Such edges arise from two-qubit gate failures in round $j \geq r + 1$ with the property that the errors introduced by the failure only flip the parity vertices in that round. In round $j + 1$, the error

is detected by X -stabilizers on both the ancilla and data patch regions. Errors introduced by such failures can flip the parity of the $X \otimes X$ measurement outcome. As such, v_2 (defined in Algorithm 1) should also include the number of highlighted red space-time correlated edges incident to transition vertices.

Thirdly, space-time correlated edges highlighted in black have the same effect for correcting errors as all other space-time correlated edges belonging to the data qubit patches (highlighted in green). The reason is that they are incident to parity vertices in rounds $j \geq r + 2$ and thus failure mechanisms leading to such edges cannot flip the parity of the $X \otimes X$ measurement outcome.

Lastly, we note that a two-qubit gate failure arising in round $r + 1$ and which results in a red highlighted space-time correlated edge will only highlight a single vertex (belonging to the data-qubit patch in round $r + 2$) throughout the entire syndrome measurement history (assuming no other failures occur). This is due to the random outcomes of X -stabilizers in round $r + 1$ (so that vertices for such stabilizers cannot be highlighted in round $r + 1$). Since there is an asymmetry between the number of red space-time correlated edges incident to the left data qubit patch and those incident to the right data qubit patch, an asymmetry in the logical failure rate polynomials $\mathbb{P}_{(1,0,0,0)}$ and $\mathbb{P}_{(0,0,1,0)}$ (defined in Section III) will also arise.

Appendix C: Resource cost analysis of the Hubbard model

Following [71], the total number of logical qubits used for simulating a Hubbard model of lattice size L is $N_Q = 2L^2 + \frac{L^2}{2} + 2$. If T gates are performed by catalysis, then an extra logical qubit is needed. We also add another logical qubit for the logical $|0\rangle$ required in the WTC/RFC protocol described in Section VI B. Using the core-cache model shown in Fig. 13, let N_1 be the number of logical qubits in the core, and N_2 the number of logical qubits in the cache. Adding the cost of the green padding shown in Fig. 12e, we now compute the total routing overhead cost in the core with h unit cells stacked in the vertical direction, and w unit cells stacked in the horizontal direction ($h = w = 2$ in Fig. 13d). We begin by defining the following functions which count the number of tiles in the green padded region of Fig. 12e, where we have separated the region into four sections

$$\begin{aligned} s_1(d_x, h) &= h(3d_x + 1), \\ s_2(d_x, d_z, w) &= d_x + 2 + w(2d_z + d_x + 1), \\ s_3(d_x, h) &= h(3d_x + 1)(d_x + 1), \\ s_4(d_x, d_z, w) &= (d_x + 1) \left(w(2d_z + d_x + 1) + d_x + 2 \right). \end{aligned}$$

The total number of tiles in the green padded region is then given by

$$\begin{aligned} S_{\text{TRGP}}(d_x, d_z, h, w) &= s_1(d_x, h) + s_2(d_x, d_z, w) \\ &\quad + s_3(d_x, h) + s_4(d_x, d_z, w). \end{aligned} \quad (\text{C1})$$

The total routing overhead in the core is then

$$O_{(d_x, d_z, h, w)}^{(\text{core})} = \frac{wh(2d_z + d_x + 1)(3d_x + 1) + S_{\text{TRGP}}}{4whd_zd_x}. \quad (\text{C2})$$

In the cache, the routing cost adds an additional $1 + \frac{N_2 - 1}{N_2 d_x} \sim 1 + \frac{1}{d_x}$ multiplier when using surface code patches of distance d_x and d_z . Hence, the total routing costs including both the core and the cache is given by

$$O_{(d_x, d_z, h, w)}^{(\text{total})} = \frac{\tilde{O}_{(d_x, d_z, h, w)}^{(\text{core})} + d_z(N_2(d_x + 1) - 1)}{d_x d_z (4wh + N_2)}, \quad (\text{C3})$$

where we defined $\tilde{O}_{(d_x, d_z, h, w)}^{(\text{core})} = 4whd_x d_z O_{(d_x, d_z, h, w)}^{(\text{core})}$.

For the Hubbard model, the total number of logical qubits N_{TLQ} which excludes those used in the magic state factory is

$$N_{\text{TLQ}} = 4wh + N_2, \quad (\text{C4})$$

with

$$N_2 = 2L^2 + \frac{L^2}{2} + 3 - 4wh. \quad (\text{C5})$$

since there are $4wh$ logical qubits in the core. The total number of physical qubits used in the algorithm is then

$$N_{\text{phys}} = 2d_x d_z N_{\text{TLQ}} O_{(d_x, d_z, h, w)}^{(\text{total})}, \quad (\text{C6})$$

where we used the fact that a surface code patch can be realized using a rectangular region of size $2d_x d_z$.

We now compute the algorithm runtime. Let μ be the total number of injected magic states in the core and Pauli measurements in the algorithm. Recall that in Section II A, μ was shown to be given by $\mu = 4N_{\text{TOF}} + N_T$. The time T_b required to inject magic states via lattice surgery is thus $T_b = \mu(d_m + 1)T_{\text{surf}}$ where T_{surf} is the time required to measure the stabilizers and reset ancillas during one surface code syndrome measurement cycle. Using Eqs. (4) and (6), the parameters d_x , d_z and d_m are chosen such that

$$0.01634\mu d_x \ell (21.93p)^{(d_m + 1)/2} < \delta/3 \quad (\text{C7})$$

$$0.03148\mu N_{\text{TLQ}} d_m d_x (28.91p)^{(d_z + 1)/2} < \delta/3 \quad (\text{C8})$$

$$0.0148\mu d_m \frac{FA}{d_x} (0.762p)^{(d_x + 1)/2} < \delta/3, \quad (\text{C9})$$

In Eq. (C7), we pessimistically take $d_x \ell = (d_x + 2 + h(3d_x + 1))(d_x + 2 + w(2d_z + d_x + 1)) - 4whd_x d_z$, which is the full area of the routing space in the core. By doing so, we consider a worst case scenario where the full routing space is used to perform lattice surgery for each injected magic states. In Eq. (C8), we used $\mathbb{P}_{Z_L} = \mathbb{P}_{(1,1,0,0)}$ since the difference with $\mathbb{P}_{(0,1,1,0)}$ is negligible. We also ignore higher order contributions arising from lattice surgery for the reasons explained in Section III. Lastly, in Eq. (C9) we pessimistically set $FA = (d_x + 2 + h(3d_x + 1))(d_x + 2 + w(2d_z + d_x + 1))$

size						# of physical	# of logical	# of logical	$O_{(d_x, d_z, h, w)}^{(\text{total})}$
L	h	w	d_x	d_z	d_m	qubits	qubits (core)	qubits (cache)	
$u/\tau = 8$									
8	2	6	7	13	12	46472	48	115	1.57
8	6	6	7	13	12	63992	144	19	2.16
32	6	8	7	15	12	657276	192	2371	1.23
32	14	18	7	15	12	812532	1008	1555	1.51

TABLE I. For a Hubbard model simulation with lattice size L and unit cells of height h and width w in the core, we provide the minimum values of d_x , d_z and d_m such that Eqs. (C7) to (C9) are satisfied with $\delta \sim 1\%$. For the given parameters, we also provide the total number of physical qubits using Eq. (C6), and the number of logical qubits in the core and in the cache. The last column includes the multiplicative factor (defined in Eq. (C3)) that is added to the physical qubit overhead which takes into account routing costs. The values for N_{TOF} and N_T used in computing μ are obtained from Ref.[71]. All resource costs exclude contributions from the magic state factory.

to be the full area in the core. Such an assignment is done to take into account the possibility that the full routing space can be used when performing lattice surgery after injecting a magic state. Such a scenario would lead to a large d_z distance, which also includes contributions from the logical qubits.

In Table I we provide overhead costs associated with performing a Hubbard model simulation of lattice size L . Given the chosen values of h and w for a unit cells in the core, we first compute the minimum required values of d_x , d_z and d_m by solving Eqs. (C7) to (C9) with $\delta \sim 1\%$. We then compute the required number of physical qubits using Eq. (C6) and give the number of logical qubits in the core and in the cache. The last column includes the multiplicative factor resulting from the routing overhead costs. As can be seen, having more logical qubits in the core relative to those in the cache can substantially increase routing overhead costs.

Apart from the chosen value of d_m which satisfies Eq. (C7), the algorithm runtime will depend on several factors. The first factor is the ratio of logical qubits used in the cache and in the core. Such a ratio will affect how many times one needs to read from, and write to the cache during the algorithm runtime. The second factor involves whether multi-qubit Pauli operators with Y terms are measured using our twist-free approach or with twists. Lastly, the third factor includes runtime savings which can be achieved using our temporal encoding scheme for fast lattice surgery. As such, a more careful analysis of the algorithm runtimes is left for future work. We also leave the inclusion of resource costs associated with the magic state factories to future work. However, from the results of Refs.[32, 72], we expect contributions from the magic state factories to only have a mild effect on the total resource overhead costs shown in Table I.

-
- [1] J. Preskill, Reliable quantum computers, *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, 385 (1998).
 - [2] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [3] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature* **549**, 172 (2017).
 - [4] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2 (2003).
 - [5] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [6] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science* (IEEE, 1996) pp. 56–65.
 - [7] R. Raussendorf, J. Harrington, and K. Goyal, Topological fault-tolerance in cluster state quantum computation, *New Journal of Physics* **9**, 199 (2007).
 - [8] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
 - [9] A. G. Fowler and C. Gidney, Low overhead quantum computation using lattice surgery, [arXiv preprint arXiv:1808.06709](https://arxiv.org/abs/1808.06709) (2018).
 - [10] D. Litinski and F. v. Oppen, Lattice Surgery with a Twist: Simplifying Clifford Gates of Surface Codes, *Quantum* **2**, 62 (2018).
 - [11] D. Litinski, A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery, *Quantum* **3**, 128 (2019).
 - [12] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, *Nature* **402**, 390 (1999).
 - [13] X. Zhou, D. W. Leung, and I. L. Chuang, Methodology for quantum logic gate construction, *Phys. Rev. A* **62**, 052316 (2000).
 - [14] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
 - [15] S. Bravyi and J. Haah, Magic-state distillation with low overhead, *Phys. Rev. A* **86**, 052329 (2012).
 - [16] A. M. Meier, B. Eastin, and E. Knill, Magic-state distillation with the four-qubit code, *Quant. Inf. and Comp.* **13**, 195 (2013).
 - [17] C. Jones, Multilevel distillation of magic states for quantum computing, *Phys. Rev. A* **87**, 042305 (2013).
 - [18] E. T. Campbell and M. Howard, Magic state parity-checker with pre-distilled components, *Quantum* **2**, 56 (2018).
 - [19] C. Chamberland and A. W. Cross, Fault-tolerant magic state preparation with flag qubits, *Quantum* **3**, 143 (2019).
 - [20] C. Chamberland and K. Noh, Very low overhead fault-tolerant

- magic state preparation using redundant ancilla encoding and flag qubits, *npj Quantum Information* **6**, 91 (2020).
- [21] H. Bombin and M. A. Martin-Delgado, Topological quantum distillation, *Phys. Rev. Lett.* **97**, 180501 (2006).
- [22] T. Jochym-O'Connor and R. Laflamme, Using concatenated quantum codes for universal fault-tolerant quantum gates, *Phys. Rev. Lett.* **112**, 010505 (2014).
- [23] H. Bombin, Dimensional jump in quantum error correction, *New Journal of Physics* **18**, 043038 (2016).
- [24] T. J. Yoder, R. Takagi, and I. L. Chuang, Universal fault-tolerant gates on concatenated stabilizer codes, *Phys. Rev. X* **6**, 031039 (2016).
- [25] C. Chamberland, T. Jochym-O'Connor, and R. Laflamme, Thresholds for universal concatenated quantum codes, *Phys. Rev. Lett.* **117**, 010501 (2016).
- [26] C. Chamberland, T. Jochym-O'Connor, and R. Laflamme, Overhead analysis of universal concatenated quantum codes, *Phys. Rev. A* **95**, 022313 (2017).
- [27] M. E. Beverland, A. Kubica, and K. M. Svore, Cost of universality: A comparative study of the overhead of state distillation and code switching with color codes, *PRX Quantum* **2**, 020341 (2021).
- [28] S. Bravyi, G. Smith, and J. A. Smolin, Trading classical and quantum computational resources, *Phys. Rev. X* **6**, 021043 (2016).
- [29] C. Jones, Low-overhead constructions for the fault-tolerant toffoli gate, *Phys. Rev. A* **87**, 022328 (2013).
- [30] B. Eastin, Distilling one-qubit magic states into toffoli states, *Phys. Rev. A* **87**, 032321 (2013).
- [31] C. Gidney and A. G. Fowler, Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation, *Quantum* **3**, 135 (2019).
- [32] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, A. Keller, G. Refael, J. Preskill, L. Jiang, A. H. Safavi-Naeini, O. Painter, and F. G.S.L. Brandao, Building a fault-tolerant quantum computer using concatenated cat codes, *arXiv preprint arXiv:2012.04108* (2020).
- [33] E. T. Campbell and M. Howard, Unified framework for magic state distillation and multiqubit gate synthesis with reduced resource cost, *Phys. Rev. A* **95**, 022316 (2017).
- [34] E. T. Campbell and M. Howard, Unifying gate synthesis and magic state distillation, *Phys. Rev. Lett.* **118**, 060501 (2017).
- [35] J. Haah and M. B. Hastings, Codes and Protocols for Distilling T , controlled- S , and Toffoli Gates, *Quantum* **2**, 71 (2018).
- [36] C. Gidney, Halving the cost of quantum addition, *Quantum* **2**, 74 (2018).
- [37] J. O'Gorman and E. T. Campbell, Quantum computation with realistic magic-state factories, *Phys. Rev. A* **95**, 032338 (2017).
- [38] E. Campbell, A. Khurana, and A. Montanaro, Applying quantum algorithms to constraint satisfaction problems, *Quantum* **3**, 167 (2019).
- [39] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of Arbitrary Basis Quantum Chemistry Leveraging Sparsity and Low Rank Factorization, *Quantum* **3**, 208 (2019).
- [40] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, Improved Fault-Tolerant Quantum Simulation of Condensed-Phase Correlated Electrons via Trotterization, *Quantum* **4**, 296 (2020).
- [41] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2**, 030305 (2021).
- [42] Y. Tomita and K. M. Svore, Low-distance surface codes under realistic quantum noise, *Phys. Rev. A* **90**, 062320 (2014).
- [43] C. Vuillot, L. Lao, B. Criger, C. G. Almudéver, K. Bertels, and B. M. Terhal, Code deformation and lattice surgery are gauge fixing, *New J. Phys.* **21**, 033028 (2019).
- [44] J. Edmonds, Paths, trees, and flowers, *Canadian Journal of mathematics* **17**, 449 (1965).
- [45] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and subsystem codes on low-degree graphs with flag qubits, *Phys. Rev. X* **10**, 011022 (2020).
- [46] C. Chamberland, A. Kubica, T. J. Yoder, and G. Zhu, Triangular color codes on trivalent graphs with flag qubits, *New Journal of Physics* **22**, 023019 (2020).
- [47] H. Bombin, Topological order with a twist: Ising anyons from an abelian model, *Phys. Rev. Lett.* **105**, 030403 (2010).
- [48] T. J. Yoder and I. H. Kim, The surface code with a twist, *Quantum* **1**, 2 (2017).
- [49] An analysis for the impact of twists on the runtime and performance of lattice surgery will appear in future work.
- [50] R. Chao and B. W. Reichardt, Quantum error correction with only two extra qubits, *Phys. Rev. Lett.* **121**, 050502 (2018).
- [51] R. Chao and B. W. Reichardt, Fault-tolerant quantum computation with few qubits, *npj Quantum Information* **4**, 2056 (2018).
- [52] C. Chamberland and M. E. Beverland, Flag fault-tolerant error correction with arbitrary distance codes, *Quantum* **2**, 53 (2018).
- [53] Y. Shi, C. Chamberland, and A. Cross, *New J. Phys.* **21**, 093007 (2019).
- [54] T. Tansuwannont, C. Chamberland, and D. Leung, Flag fault-tolerant error correction, measurement, and quantum computation for cyclic calderbank-shor-steane codes, *Phys. Rev. A* **101**, 012342 (2020).
- [55] R. Chao and B. W. Reichardt, Flag fault-tolerant error correction for any stabilizer code, *PRX Quantum* **1**, 010302 (2020).
- [56] B. W. Reichardt, *Quantum Sci. Technol.* **6**, 015007 (2020).
- [57] T. Tansuwannont and D. Leung, Achieving fault tolerance on capped color codes with few ancillas, *arXiv e-prints*, arXiv:2106.02649 (2021), arXiv:2106.02649.
- [58] A. G. Fowler, Time-optimal quantum computation, *arXiv preprint arXiv:1509.03239* (2012).
- [59] I. H. Kim, E. Lee, Y.-H. Liu, S. Pallister, W. Pol, and S. Roberts, Fault-tolerant resource estimate for quantum chemical simulations: Case study on li-ion battery electrolyte molecules, *arXiv preprint arXiv:2104.10653* (2021).
- [60] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, *et al.*, Bias-preserving gates with stabilized cat qubits, *Science advances* **6**, eaay5901 (2020).
- [61] E. T. Campbell, A theory of single-shot error correction for adversarial noise, *Quantum Science and Technology* **4**, 025006 (2019).
- [62] A. O. Quintavalle, M. Vasmer, J. Roffe, and E. T. Campbell, Single-shot error correction of three-dimensional homological product codes, *PRX Quantum* **2**, 020340 (2021).
- [63] D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, Fault-tolerant thresholds for the surface code in excess of 5% under biased noise, *Phys. Rev. Lett.* **124**, 130501 (2020).
- [64] A. Ashikhmin, C.-Y. Lai, and T. A. Brun, Quantum data-syndrome codes, *IEEE Journal on Selected Areas in Communications* **38**, 449 (2020).
- [65] D. K. Tuckett, A. S. Darmawan, C. T. Chubb, S. Bravyi, S. D. Bartlett, and S. T. Flammia, Tailoring surface codes for highly biased noise, *Phys. Rev. X* **9**, 041031 (2019).
- [66] J. P. Bonilla Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia,

- and B. J. Brown, The XZZX surface code, *Nature Communications* **12**, 2172 (2021).
- [67] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, Practical quantum error correction with the XZZX code and Kerr-cat qubits, *arXiv e-prints* (2021), 2104.09539 [quant-ph].
- [68] O. Higgott and N. P. Breuckmann, Subsystem codes with high thresholds by gauge fixing and reduced qubit overhead, *Phys. Rev. X* **11**, 031039 (2021).
- [69] N. P. Breuckmann, C. Vuillot, E. Campbell, A. Krishna, and B. M. Terhal, Hyperbolic and semi-hyperbolic surface codes for quantum storage, *Quantum Science and Technology* **2**, 035007 (2017).
- [70] J. Conrad, C. Chamberland, N. P. Breuckmann, and B. M. Terhal, The small stellated dodecahedron code and friends, *Phil. Trans. R. Soc. A*, **376**, 20170323 (2018).
- [71] E. T. Campbell, Early fault-tolerant simulations of the Hubbard model, *arXiv*, 2012.09238 (2020).
- [72] D. Litinski, Magic State Distillation: Not as Costly as You Think, *Quantum* **3**, 205 (2019).