

Optimizing compilation of error correction codes for $2 \times n$ quantum dot arrays and its NP hardness

Anthony Micciche^{1,*}, Feroz Ahmed Mian,¹ Anasua Chatterjee², Andrew McGregor¹, and Stefan Krastanov^{1,†}

¹Manning College of Information and Computer Sciences, *University of Massachusetts Amherst*,
140 Governors Drive, Amherst, Massachusetts 01003, USA

²QuTech and Kavli Institute of Nanoscience, *Delft University of Technology*, Delft, The Netherlands



(Received 27 February 2025; accepted 12 January 2026; published 25 February 2026)

The ability to physically move qubits within a register allows the design of hardware-specific error correction codes, which can achieve fault tolerance while respecting other constraints. In particular, recent advancements have demonstrated the shuttling of electron and hole spin qubits through a quantum dot array with high fidelity. It is therefore timely to explore error correction architectures consisting merely of two parallel quantum dot arrays, an experimentally validated architecture compatible with classical wiring and control constraints. Upon such an architecture, we develop a suite of heuristic methods for compiling any Calderbank-Shor-Steane (CSS) error-correcting code's syndrome-extraction circuit to run with a reduced number of shuttling operations. We demonstrate how column-regular qLDPC codes can be compiled in a provably minimal number of shuttles that is exactly equal to the column weight of the code when Shor-style syndrome extraction is used. We provide tables stating the number of required shuttles for many contemporary codes of interest. In addition, we provide a proof of the NP hardness of minimizing the number of shuttle operations for general codes, even when using Shor-syndrome extraction. We also discuss how one could get around this by placing blanks in the ancilla array to achieve minimal shuttles with Shor-syndrome extraction on any CSS code, at the cost of longer ancilla arrays.

DOI: [10.1103/PhysRevResearch.8.013212](https://doi.org/10.1103/PhysRevResearch.8.013212)

I. INTRODUCTION

Motivated by the potential of far easier scalability than other quantum computing architectures due to industrial fabrication compatibility [1] as well as long coherence times [2], in recent years there have been numerous advancements in experimentally realizing spin qubits in electrically defined quantum dots. In particular, the repeated shuttling of an electron or hole up and down an array of quantum dots with a per-hop error rate less than 0.01% [3,4] has been demonstrated. On this experimentally validated foundation of being able to physically move qubits along arrays of quantum dots, we explore quantum computing architectures consisting merely of two of these quantum dot arrays, as shown in Fig. 1. These hardware layouts have been referred to as “ $2 \times n$ ” architectures [5,6]. Current fabrication restrictions on the fan-out, wiring, and wirebonding of these semiconductor quantum dot arrays make it necessary to explore the possibility of fault-tolerant computation in this restricted configuration of two parallel arrays. Natural implementations of error correction codes for such architectures were recently proposed in Ref. [7]; however, without yet the capability of

optimizing the syndrome-extraction circuits under the significant connectivity constraints of the hardware for codes beyond the surface code. In this work, we provide techniques for efficiently compiling *any* Calderbank-Shor-Steane (CSS) quantum error correction code to such a $2 \times n$ architecture. In addition, we show that “column-regular” qLDPC codes can always be compiled to run with a small constant number of shuttles, as their property guarantees the existence of a trivial compilation solution under a certain style of fault-tolerant syndrome extraction.

More explicitly, we propose techniques that convert one quantum circuit to an equivalent circuit, which will then require fewer shuttling operations than the original to execute, thereby causing less error to accumulate. We refer to this process of modifying the initial syndrome-extraction circuit as “compilation.”

While portions of our approach could potentially be applied to running arbitrary circuits on this architecture, we focus our discussion and assumptions on syndrome-extraction circuits used in quantum error correction, which is a natural fit for a two-row architecture as we can simply assign data qubits to the top row, and ancillary qubits to the bottom row. Shuttling these rows past each other enables us to perform gates between distant qubits, enabling the efficient implementation of codes that even require long-range connectivity such as qLDPC codes, a very promising class of codes for achieving large-scale quantum error correction [8,9]. This work aims to provide insight on which error correction codes would be most effective for this hardware. Furthermore, we present tools for evaluating how much shuttling error an arbitrary error-correcting stabilizer code would incur. More general

*Contact author: amicciche@umass.edu

†Contact author: skrastanov@umass.edu




FIG. 1. Schematic of the quantum dot array hardware. On the left is a possible configuration of this quantum memory, which we parametrize as $\delta^\tau = 5$, where δ^τ represents the offset between the indices of the two rows of qubits (data and ancillary qubits in the case of quantum error correction) at time τ . This $\delta^\tau = 5$ implies that cross-array interaction is only permitted between qubit pairs (1, 6), (2, 7), and (3, 8). If we shuttle the bottom row to the left, we would then have the right figure which has $\delta^{\tau+1} = 6$ and interaction permitted only between (1, 7) and (2, 8).

layouts, connecting separate $2 \times n$ blocks are now also being suggested [10] and our work will be of significant benefit to such large-scale deployments as well.

This paper is structured as follows: In Sec. II, we present a way to formalize the general minimization of qubit shuttles. In Sec. III, we then present a simple way of visualizing the problem formalized in the preceding section. This visualization naturally lends itself to simple heuristic methods that, while not always optimal, yield surprisingly effective results for certain circuits. Following that, in Sec. IV, we discuss a special case of our problem corresponding to fault-tolerant Shor-style syndrome-extraction circuits and how this special case gives us enough structure to substantially improve our compilation. After these discussions on how to reduce the number of shuttling operations, in Sec. V we provide many relevant examples of both column-regular and column-irregular codes, and provide the number of shuttles our methods were able to compile their syndrome-extraction circuits to. This then culminates in Sec. VI, where we discuss how regular column weight LDPC codes compile with a provably minimal number of shuttles under fault-tolerant syndrome extraction. Further elaborating on ways to construct an error-correcting scheme with minimum number of shuttle operations of the ancilla array, Sec. VI also discusses another way of viewing the compilation problem, trading additional space in the ancilla array for the guarantee of minimum shuttling operations on any code. Finally, in the concluding remarks, we discuss some of the limitations and potential future work, such as the importance of investigating postselected cat state preparation for use with Shor-syndrome extraction.

In the Appendix, we provide a proof of the NP hardness of compilation via reduction from 3-partition, a brief discussion of how the larger error correction pipeline would look including the cat state preparation required for fault-tolerant syndrome extraction on $2 \times n$ architecture, as well as a gate-by-gate simulation of the error correction pipeline, verifying that our techniques of $2 \times n$ compilation do indeed result in functional syndrome circuits with better error rates corresponding to fewer shuttles. The Appendix also includes a few additional figures and examples to help illustrate the methods presented here.

II. FOUNDATIONS OF $2 \times n$ CIRCUIT COMPILATION

We formalize the problem as follows: We are given a circuit as an array of m two-qubit gates, $\mathcal{C} = A_1, A_2, \dots, A_m$, and we

say that $\delta(A)$ represents the difference in the indices of qubits involved in gate A . For example, if A is a CNOT from a qubit at index i to a qubit at index j , then $\delta(A) = |i - j|$. We assume that we can lay out our qubits such that each gate always involves one qubit from the top row and one qubit from the bottom row (e.g., in a syndrome-extraction circuit, the top row would be data qubits and the bottom row would be ancillary qubits). We also let δ^t represent the state of our system at a given time step t , as enumerated and shown in Fig. 1. That is, δ^t denotes the difference in indices between cross-row pairs of qubits. Because δ^t is thus a discrete measure of the difference of positions of the two rows and we consider one-dimensional dot arrays, all values of delta are reachable by shuttling only one of the two rows. Therefore, as the way we define δ^t depends only on the positions of the two rows relative to each other, our abstraction holds when only one of the two rows is allowed to move, and so we assume a stationary row of data qubits throughout this work. A more detailed sketch of our envisioned hardware is presented in Fig. 12.

Notice that a gate A can execute only when $\delta(A) = \delta^t$, i.e., when the two rows are aligned such that the two qubits involved in gate A are facing each other. It will often be convenient to track the δ values of the gates in the circuit $\mathcal{C}$, which we denote by $\mathcal{G} := [\delta(A_1), \delta(A_2), \dots, \delta(A_m)]$. We now discuss the two general operations that we can perform on a given syndrome-extraction circuit $\mathcal{C}$ in order to compile it into a $\mathcal{C}'$ that extracts the same syndrome, only now requiring fewer shuttle operations:

- (1) Change the order of gates in $\mathcal{C}$. We refer to this as “gate shuffling.”
- (2) Apply a permutation to the indices of the qubits. In other words, define a bijection from $\{1, \dots, n + s\}$ to $\{1, \dots, n + s\}$, where n is the number of data qubits and s is the number of ancillary qubits. Then, for each gate $A \in \mathcal{C}$, apply this bijection to both indices. We refer to this as “qubit reindexing.” For example, see Fig. 2.

For qubit reindexing, there are no restrictions on how to define the bijection, as qubit indices are arbitrary in general, while in the case of gate shuffling, gates cannot arbitrarily be reordered. However, within a syndrome-extracting circuit for a CSS code, all gates measuring the X stabilizers may be permuted arbitrarily, and the same holds true for gates measuring Z stabilizers. Similar to the idea of performing syndrome extraction on one type of stabilizers while shuttling in one direction and then extract the other stabilizers while shuttling back in the other as proposed in Ref. [7], we also will




FIG. 2. Example of how qubit reindexing reduces shuttles. On the left is an unoptimized circuit. $\mathcal{G}$ corresponds to the values of δ for each gate in the circuit, $\delta(A)$, $\delta(B)$, and $\delta(C)$. $|\mathcal{G}|_0$ denotes the number of unique elements in $\mathcal{G}$. Notice that to run this circuit on a $2 \times n$ architecture, two shuttles (or $|\mathcal{G}|_0$ positions) are necessary. On the right is the same circuit, except the indexing of the qubits has changed, which then allows the same circuit as before to run, this time only requiring a single shuttle.

assume our code is CSS, and compile the syndrome-extraction circuits for X and Z checks separately (e.g., by reindexing only ancillary qubits and keeping the data qubit indices fixed).

We refer to a circuit that measures the parity of a weight w stabilizer by applying w gates to a single ancillary qubit in a non-fault-tolerant way as “naive syndrome measurement” [see Fig. 3(a)]. For the purposes of compilation, this corresponds to the general problem, as each ancilla may have multiple gates attached to it. We consider only the two-qubit gates in the circuits to be compiled (i.e., single-qubit gates are folded into the neighboring two-qubit gates and not explicitly tracked).

A. Gate shuffling

Before considering the more interesting problem of qubit reindexing, we first describe how for a given ordering of the qubits, changing the gate order can reduce the number of required shuttles. This will be an implicit primitive behind the




FIG. 3. Measuring the parity of $Z_1Z_2Z_3Z_4$ via two types of syndrome extraction. (a) Naive syndrome extraction: Parity measurement is given by coupling all pertinent data qubits to single ancillary qubit, and then performing a measurement. (b) Shor-style syndrome extraction [11]: Each data qubit indicated by a nonidentity element in a weight w stabilizer is coupled to its own ancillary qubit. For that single parity check, the participating ancilla are beforehand prepared in a w -qubit GHZ state. Each ancilla is then measured, and the classical XOR is taken over the resulting values to yield the parity of the stabilizer measurement. Double lines here represent the flow of classical information.

definition of the objective over which we want to minimize when solving the qubit reindexing problem.

Starting from some arbitrary order of a circuit $\mathcal{C} = [A_1, A_2, \dots, A_m]$, referring to Fig. 1, it should become obvious that if two distinct (acting on different qubits) gates A_i and A_j have the same δ value, then both gates can run concurrently at time step t if the relative position of the two rows is $\delta^t = \delta(A_i) = \delta(A_j)$.

Now, because gates with the same δ can run concurrently, by merely permuting the gates of $\mathcal{C}$ so that gates of the same δ are adjacent to one another (sorting $\mathcal{C}$ by δ is a straightforward way to achieve this), we can partition $\mathcal{C}$ into a collection of subarrays for each corresponding different value of δ . We call the number of such subarrays $|\mathcal{G}|_0$, where $|\mathcal{G}|_0$ is defined as the number of unique values within $\mathcal{G}$ —the array that contains the δ values for the current order of gates. All gates in a given subarray can be executed concurrently, and a single shuttling operation is required to set the architecture in the correct configuration for the corresponding subarray to be executed. We can therefore state that $|\mathcal{G}|_0$ represents the minimum number of required shuttles obtainable by gate shuffling, as we can always trivially sort the gates by their corresponding δ , assuming all gates in $\mathcal{C}$ commute.

Building on this result of how to perform gate shuffling to achieve $|\mathcal{G}|_0$ shuttles, we now discuss qubit reindexing with the goal of finding the bijection of qubit indices that would best minimize $|\mathcal{G}|_0$, further reducing the required number of shuttle operations. To solve this, we fix the data qubits, and only think about reindexing the ancillary qubits. While the main motivation for fixing the data qubits is to make the problem more tractable, as mentioned before, fixing the data qubits also causes the case of having different X and Z extraction circuits for CSS codes to be more manageable as well.

B. Reducing shuttles via ancillary qubit reindexing

For a given naive syndrome-extraction circuit $\mathcal{C}$, we would like to find a bijection of qubit indices that minimizes the




FIG. 4. Illustration of primitive sets under different reorderings. On the left is an arbitrary circuit, and below are the δ values for each of its gates, $\mathcal{G}$, as well as the number of unique values, $|\mathcal{G}|_0$. Below that, we partition $\mathcal{G}$ into three sets, each corresponding to the gates connected to each ancillary qubit. Below each set is a “primitive set” with some integer being added elementwise. This addition encodes the current indexing of the ancillary qubits. On the right is a different way of indexing these ancillary qubits. Notice that although the values in $\mathcal{G}$ have changed, the primitive sets remain the same.

resulting $|\mathcal{G}|_0$ (i.e., the number of necessary shuttles). As an example, we refer to Fig. 4, where we would like to reorder qubits 6, 7, and 8, noting that when the index of an ancillary qubit is changed, the δ values of all gates connected to it will also change. To account for this, we partition $\mathcal{G}$ into s sets, where s is the number of ancillary qubits, and each set will contain the δ values for the gates touching the corresponding ancillary qubit. In the example present in the left-hand portion of Fig. 4, $\mathcal{G} = [3, 4, 5, 2, 3, 5]$ is partitioned into sets R_6, R_7, R_8 , equal to, respectively, $\{3, 4, 5\}, \{2, 3\}, \{5\}$, where $\{3, 4, 5\}$ are the δ values for all gates connected to qubit 6, $\{2, 3\}$ are the δ values for all gates connected to qubit 7, and $\{5\}$ is the δ of the one gate connected to qubit 8. As we are interested in reordering these ancillary qubits, we need to see how their current position affected the values of $\mathcal{G}$. For this reason, we label the ancillary qubits $+1, +2, \dots, +s$, and notice that each position augments the values in corresponding set elementwise. For example, qubit 6 corresponds to the set $R_6 = \{3, 4, 5\}$, which we can then write as $R_6 = \{2, 3, 4\} + 1$, where the $+1$ encodes that this is the first ancillary qubit. We refer to the set P_i of $R_i = P_i + k_i$ as the *primitive set* of the i th ancillary qubit (k_i is the integer corresponding to the current position of P_i). The primitive set encodes the relative positions between all top-row (data) qubits that interact with a given bottom-row (ancillary qubit), while k_i encodes the current position of that ancillary qubit in the bottom row. On the right-hand side of Fig. 4, the ancillary qubits have been reordered to minimize $|\mathcal{G}|_0$. Notice how the primitive sets remain the same under this reordering (how to derive the primitive sets from a given error-correcting code is discussed in more detail in Sec. VI). Furthermore, we can concretely define an admittedly not very friendly combinatorial optimization problem as follows:

Problem formulation. Given a collection of s sets $\mathcal{S} = \{P_1, P_2, \dots, P_s\}$, find a bijection $\pi : [s] \rightarrow [s]$, such that $|\mathcal{G}|_0$ is minimized ($\mathcal{G}$ depends on the R sets, which each depend on their corresponding P set and the index reordering π). More

formally,

$$\min_{\pi} |\mathcal{G}|_0 = \min_{\pi} \left| \bigcup_{i=1}^s \bigcup_{j \in P_i} \{j + \pi(i)\} \right|. \quad (1)$$

As $|\mathcal{G}|_0$ represents the number of required shuttles, this procedure of minimizing $|\mathcal{G}|_0$ by reordering the ancillary qubits results in a circuit that is more efficient by requiring fewer shuttles and would then have a lower execution time due to increased gate parallelism.

Although we prove in Appendix A2 that, even in the special case where all primitive sets are cardinality 1, finding the optimal reindexing is an NP-hard problem, in the following section, we provide some simple heuristics to act as a baseline for solving this problem. Furthermore, in Sec. IV, we provide a more substantial heuristic algorithm for minimizing $|\mathcal{G}|_0$ when we are promised that the input circuit is in the style of Shor-syndrome extraction (primitive sets are cardinality 1).

III. HEURISTICS FOR GENERAL COMPILATION

In this section, we discuss how to develop heuristics for the naive syndrome qubit reindexing problem stated in the previous section. The first step of which is converting the problem to something that is easier to visualize. Again, referring to the example present in Fig. 4, we now recast it to stacking blocks on a “staircase.” We do so by thinking of each primitive set as a chain of blocks, of which there are two varieties and the staircase to discuss (Fig. 5 for a visual reference):

(1) *Solid blocks (blue).* These blocks represent the presence of an element, or what used to be a target ($\oplus$ for CNOT gates) in the quantum circuit diagram. As we are fixing the data qubits, we can also think of solid blocks as gates.

(2) *Air blocks (green).* These blocks represent the absence of an element and therefore are referred to as “air.” The reason for them is to preserve the structure and constraints of the original problem.




FIG. 5. Visualization of block placement. (a) Primitive sets of the circuit in Fig. 4 represented in our so-called block form. The green blocks represent the absence of an element in the set, while the blue blocks represent the presence of the corresponding element. Placing these blocks on the staircase represent assignments. For example, placing P_6 on +1 would cause all its blocks to shift to the right by one. After placing all blocks, all columns that contain at least one blue block would correspond to a unique element in $\mathcal{G}$, and recall unique values in $\mathcal{G}$ correspond to different configurations of the $2 \times n$ architecture and therefore require shuttles. (b) Optimal placement for this problem.

(3) *Staircase (gray)*. The staircase represents the set $A = \{1, \dots, s\}$ that we are assigning primitive sets to. By placing P_i on the first step of the staircase, we thereby are shifting all elements of P_i to the right by one, visually adding 1 to the set elementwise. If we were to place P_i on the j th step, then we would be adding j to P_i elementwise, and the resulting set corresponds to what we call R_i . The goal then is to place each chain of blocks on this staircase such that the number of columns containing at least one blue block is minimized, as each column in this staircase abstraction corresponds to a particular configuration of the $2 \times n$ architecture. Moreover, minimizing the number of required configurations of the two rows of qubits is exactly the same as minimizing the number of required shuttling operations.

In order to create the corresponding chain of blocks for a given primitive set P , we start with a number of blocks equal to the largest element in P plus one. Then, enumerating those blocks from 0 to the largest element in P , blocks that correspond to elements in P are solid blocks, and all other blocks are air blocks. The example provided in Fig. 5 has only contiguous solid blocks; however, in general this will not be the case and is a source of difficulty in finding good heuristics (Fig. 9 has an example of noncontiguous solid blocks). The main benefit of this stacking block interpretation is that now it becomes much more simple to reason about heuristics and the problem as a whole. For instance, two rather decent initial heuristics for finding good orders are as follows:

- (1) Sort these chains of blocks by total length, placing the longest ones at the top (+1 is the top of the staircase).
- (2) Sort these chains by the number of air blocks preceding the first solid block.

Because each of these will only take $O(n \log(n))$ time, it is quite efficient to run both and simply take the better of the two choices, breaking ties with the heuristic not being primarily sorted by. Thereupon, as the map from the old order to the new order discovered for packing blocks on the staircase directly corresponds to the bijection π mentioned in the problem formulation, we can apply this mapping to the ancillary qubits in the original circuit to create a compiled circuit. We call this method ancillary heuristic reindexing (AHR). Notice that both sorting metrics would work in the simple problem of Fig. 5. To see these heuristics finding optimal solutions on real (although pedagogical) codes, such as Shor's 9 qubit code, and Steane's 7 qubit code, refer to Appendix A 1. These techniques are also sufficient to find optimal ancilla orders on column-regular LDPC codes when using Shor-syndrome extraction (see Table I). A third heuristic which is also checked when we mention AHR (the best of three is chosen as the output) is to sort all of the primitive sets by the first element in their set in descending order. Values that have occurred already are deprioritized according to the number of times they have occurred. For example, this would sort $\{6, 6, 6, 5, 5, 4, 3, 3, 2, 1, 1, 0\}$ as $\{6, 5, 4, 3, 2, 1, 0, 6, 5, 3, 1, 6\}$. This was originally introduced as first guess to solving the case when all primitive sets are size 1 (this corresponds to Shor-syndrome extraction and is discussed in the following section); however, we find that occasionally this method finds a decent solution when the first two heuristics fail. For example, this third heuristic is the best of the three for the toric code's naive syndrome-extraction circuit.

IV. SHOR-STYLE FAULT-TOLERANT CIRCUIT COMPILATION

In the previous section, we discussed the general problem of assigning primitive sets to integers. However, the circuits that the sets were derived from were not quite representative of fault-tolerant quantum error-correcting circuits, as they do not provide any protection against the propagation of errors to the ancillary qubits. To remedy this and to assess a simplified version of the problem, in this section we consider a special case where all primitive sets are of size 1 (or equivalently, this is the special case where no ancillary qubit will participate in more than a single two-qubit gate), which corresponds directly to fault-tolerant Shor-style syndrome extraction [11]. Whereas before we were trying to optimize circuits that use one ancillary qubit per stabilizer check, we now model the circuit that uses w ancillary qubits for a weight w stabilizer check [see Fig. 3(b)]. There are some state-preparation concerns for the ancilla involved in this technique, as discussed in Appendix A 3, but they are of no consequence to what we discuss here.

This leaves us with just the optimization of two-qubit gates between the data and ancillary qubits as in the previous sections. This cost of extra qubit overhead is balanced not only by guarantees of fault tolerance, but also by additional structure that will allow us to compile the code to run in some cases

TABLE I. Compilation numbers for Shor-style syndrome circuits for extracting X , Z stabilizers separately on CSS codes. As described in Sec. VB, the different columns correspond to the required numbers of shuttles corresponding to different levels of compilation (the number of shuttles is listed for X circuits first and then Z). However, “NumChains” corresponds to the smallest number of shuttles possible, obtainable by introducing “blanks” into the ancilla array. A very naive upper bound on the number of blanks required for perfect compilation is presented in the “Blanks” column, rather than a number of shuttles. Refer to Sec. VI for more on blanks and perfect compilation. Quantum error-correcting code parameters: N is the number of physical qubits, K is the number of logical qubits, and true code distance $D = \min(d_X, d_Z)$ is defined as the minimum weight of a nontrivial logical Pauli operator. Distances were calculated using the mixed-integer programming method of Ref. [19].

Code	N	K	D	Uncompiled	GateShuffled	AHR	SSSC	NumChains	Blanks
2DTORIC	18	2	3	(27, 26)	(21, 21)	(4, 5)	(4, 4)	(2, 2)	(2, 3)
2DTORIC	32	2	4	(49, 48)	(35, 35)	(4, 5)	(4, 5)	(2, 2)	(2, 3)
2DTORIC	50	2	5	(77, 76)	(62, 61)	(4, 5)	(4, 5)	(2, 2)	(2, 3)
2DSURFACE	13	1	3	(18, 14)	(14, 13)	(3, 5)	(3, 4)	(2, 2)	(3, 5)
2DSURFACE	25	1	4	(36, 30)	(26, 25)	(3, 6)	(3, 4)	(2, 2)	(4, 7)
2DSURFACE	41	1	5	(60, 52)	(48, 34)	(3, 7)	(3, 5)	(2, 2)	(5, 9)
SHOR	9	1	3	(2, 4)	(2, 4)	(2, 4)	(2, 4)	(2, 2)	(0, 4)
STEANE	7	1	3	(7, 7)	(7, 7)	(3, 3)	(3, 3)	(3, 3)	(1, 1)
2DCOLOR4.8.8	17	1	5	(19, 19)	(15, 15)	(7, 7)	(4, 4)	(3, 3)	(7, 7)
2DCOLOR4.8.8	31	1	7	(38, 38)	(27, 27)	(9, 9)	(7, 7)	(3, 3)	(10, 10)
2DCOLOR4.8.8	71	1	11	(94, 94)	(67, 67)	(13, 13)	(9, 9)	(3, 3)	(18, 18)
2DCOLOR6.6.6	19	1	5	(22, 22)	(17, 17)	(7, 7)	(6, 6)	(3, 3)	(5, 5)
2DCOLOR6.6.6	37	1	7	(48, 48)	(32, 32)	(10, 10)	(7, 7)	(3, 3)	(9, 9)
2DCOLOR6.6.6	91	1	11	(127, 127)	(82, 82)	(16, 16)	(12, 12)	(3, 3)	(17, 17)
LIFTEDPRODUCT	882	24	$d \leq 24$	(2212, 2212)	(1820, 1820)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
LIFTEDPRODUCT	175	19	$d \leq 10$	(582, 584)	(395, 377)	(4, 4)	(4, 4)	(4, 4)	(0, 0)
GENERALIZEDBICYCLE	48	6	8	(192, 192)	(139, 139)	(4, 4)	(4, 4)	(4, 4)	(0, 0)
GENERALIZEDBICYCLE	254	28	$14 \leq d \leq 20$	(1144, 1144)	(827, 827)	(5, 5)	(5, 5)	(5, 5)	(0, 0)
TWOBLOCKGROUPALGEBRA	70	8	10	(212, 212)	(145, 145)	(4, 4)	(4, 4)	(4, 4)	(0, 0)
TWOBLOCKGROUPALGEBRA	40	8	5	(121, 121)	(74, 74)	(6, 6)	(6, 6)	(6, 6)	(0, 0)
TWOBLOCKGROUPALGEBRA	56	28	2	(142, 142)	(136, 136)	(6, 6)	(6, 6)	(6, 6)	(0, 0)
TWOBLOCKGROUPALGEBRA	64	8	8	(224, 224)	(159, 168)	(6, 6)	(6, 6)	(6, 6)	(0, 0)
MULTIVARIATEBICYCLE	48	4	6	(144, 144)	(93, 93)	(4, 4)	(4, 4)	(4, 4)	(0, 0)
LA-CROSS	100	4	5	(192, 156)	(135, 99)	(5, 19)	(12, 15)	(3, 3)	(24, 45)
LA-CROSS	98	18	4	(252, 252)	(201, 201)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
LA-CROSS	117	9	4	(240, 216)	(180, 160)	(5, 21)	(12, 5)	(3, 3)	(45, 77)
LA-CROSS	296	16	7	(630, 580)	(461, 458)	(5, 31)	(46, 27)	(3, 3)	(98, 163)
LA-CROSS	400	16	8	(876, 816)	(618, 605)	(5, 35)	(52, 39)	(3, 3)	(112, 187)
BIVARIATEBICYCLE	144	12	12	(372, 372)	(229, 229)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
BIVARIATEBICYCLE	360	12	$d \leq 24$	(930, 930)	(549, 549)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
BIVARIATEBICYCLE	756	16	$d \leq 34$	(2268, 2268)	(1442, 1442)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
TRIVARIATETRICYCLE	72	6	6	(207, 415)	(148, 315)	(3, 6)	(3, 6)	(3, 6)	(0, 0)
TRIVARIATETRICYCLE	432	12	12	(1266, 2532)	(838, 1722)	(3, 6)	(3, 6)	(3, 6)	(0, 0)
TRIVARIATETRICYCLE	840	9	16	(2520, 5040)	(1642, 3214)	(3, 6)	(3, 6)	(3, 6)	(0, 0)
TRIVARIATETRICYCLE	1029	18	16	(3087, 6174)	(1890, 4429)	(3, 6)	(3, 6)	(3, 6)	(0, 0)
3DSURFACE	12	1	2	(9, 23)	(9, 16)	(3, 4)	(3, 4)	(2, 3)	(4, 4)
3DSURFACE	51	1	3	(66, 138)	(46, 90)	(13, 10)	(10, 8)	(2, 4)	(17, 31)
4DSURFACE	33	1	4	(64, 75)	(49, 54)	(8, 8)	(7, 7)	(4, 4)	(16, 16)
4DSURFACE	241	1	9	(656, 715)	(419, 461)	(67, 37)	(44, 27)	(4, 4)	(152, 196)
GENERALIZEDTORIC	292	18	8	(730, 730)	(504, 504)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
GENERALIZEDTORIC	392	6	$d \leq 28$	(280, 280)	(196, 196)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
GENERALIZEDTORIC	396	8	$d \leq 36$	(991, 991)	(583, 583)	(3, 3)	(3, 3)	(3, 3)	(0, 0)
HOMOLOGICALPRODUCT	117	9	4	(225, 225)	(155, 162)	(12, 8)	(22, 7)	(3, 3)	(27, 42)
HOMOLOGICALPRODUCT	225	9	6	(618, 630)	(432, 447)	(27, 14)	(28, 11)	(4, 4)	(72, 120)
DBLEHOMOLOGICALPRODUCT	486	6	9	(1782, 1782)	(1204, 1204)	(4, 4)	(4, 4)	(4, 4)	(0, 0)
QUASICYCLIC_GHP	882	48	16	(3093, 3093)	(2151, 2151)	(5, 5)	(5, 5)	(5, 5)	(0, 0)
QUASICYCLIC_GHP	882	12		(3976, 3976)	(2829, 2829)	(7, 7)	(7, 7)	(7, 7)	(0, 0)




FIG. 6. Feasible solution to reindexing of Shor-style syndrome extraction for the Steane code. On the right is our set $\mathcal{S}$, derived from the Shor-style syndrome-extraction circuit for the Steane code (X and Z checks are compiled together to create a more pedagogical example). On the left, we have a nonoptimal ordering. Here, we are placing the numbers from the circle (each number corresponding to a gate in the syndrome circuit) onto this “ladder,” with the goal being to minimize the number of different output values obtaining by the addition of the placed value with the index it was placed in. Each output value corresponds to the δ of the placed gate after reindexing. The ordering in this figure results in seven different output (or δ) values, although an optimal solution of six values exists, namely, Fig. 7.

far fewer shuttles than the naive case. We show in Sec. VI that column-regular codes always compile optimally using this syndrome-extraction technique.

A. Special case of qubit reindexing

As we were doing in the naive syndrome-extraction circuit case, we again only reindex ancillary qubits, so that the data qubits retain the same labeling between the X and Z circuits. However unlike before, we conjecture that in the case of Shor-syndrome circuits, reindexing the data qubits will never lead to an advantage in shuttle minimization.

We now note how the structure of Shor-syndrome circuits causes our problem definition to change. Because all primitive sets become size 1 (i.e., an ancillary qubit interacts with exactly one data qubit, ensuring fault tolerance), instead of finding an assignment from integers to sets, we now only need to find an assignment between integers and integers. More formally, making use of all sets being singletons, we can simplify the problem formulation from the naive syndrome case to

Shor-circuit problem formulation. Given a multiset of s integers $\mathcal{S} = \{p_1, p_2, \dots, p_s\}$, find a bijection $\pi : [s] \rightarrow [s]$, such that $|\mathcal{G}|_0$ is minimized, which can be written as follows:

$$\min_{\pi} |\mathcal{G}|_0 = \min_{\pi} \left| \bigcup_{i=1}^s \{p_i + \pi(i)\} \right|. \quad (2)$$

Just as in the naive syndrome case, the minimization of $|\mathcal{G}|_0$ corresponds to the minimization of the number of required number of shuttles, and the π we find in this process describes how we should reorder our ancillary qubits. As before, doing so also increases efficiency and decreases the amount of time required per syndrome-extracting cycle. Refer to Fig. 6 for an example.

This formulation allows us to make a few claims about the hypothetical optimal solution structure, which we can leverage in the design of our algorithm. First, we notice that if a

value v occurs f times in $\mathcal{S}$, then $\min_{\pi} |\mathcal{G}|_0 \geq f$. This follows simply from us considering a submultiset $U := \{v, v, \dots, v\} \subseteq \mathcal{S}$, and U contains f elements. Now, if we add to each element of U a value from $\{1, 2, \dots, s\}$ without repetition, and call the resulting multiset U' , then U' will contain exactly f unique elements, regardless of how we choose elements from $\{1, \dots, s\}$.

Therefore, we can extrapolate this result to saying that a lower bound on the minimum number of shuttles will be the maximum number of times any number appears in $\mathcal{S}$, which corresponds to the maximum column weight of a code’s stabilizer tableau (Sec. VI elaborates on this point). Referring to Fig. 6, one might ascertain that an optimal solution would be some collection of sequences that together map to a single value, that is, mapping $\{6, 5, 4, 3, 2, 1, 0\}$ to $\{u, u+1, u+2, u+3, u+4, u+5, u+6\}$, $u \in \mathbb{Z}^+$, would result in a set with only one unique element (and therefore the corresponding gates can then run concurrently on the same $2 \times n$ architecture configuration). Moreover, these sequences need not be strictly consecutive; there may be gaps in them, such as mapping $\{4, 2, 1, 0\}$ to $\{v, v+2, v+3, v+4\}$, $v \in \mathbb{Z}^+$. We call these sets derived from $\mathcal{S}$ “chains,” and discontinuities in their elements “gaps.” In other words, a *chain* is a collection of distinct elements that can all be easily mapped to the same value when placed on the ladder (the “ladder” is functionally the same as the “staircase” in the previous section; see Fig. 11 for further clarification on how to translate between these two notations). When we interpret this back to being a circuit, it will be a collection of gates that will run at the same time step, as they will all have the same δ after reindexing.

Therefore, the first step of our algorithm is to find a minimal set of chains, as shown in Algorithm 1 (a minimal set because there may be multiple minimal sets of chains possible for a given input). We do this by initializing an empty set that will become our first chain and then begin moving values from $\mathcal{S}$ to this set, never adding a value that is already present in the chain. When we can no longer add any more values, we create an empty set to be the second chain and repeat this process. Then, we create a third chain and so on until there are no elements left in $\mathcal{S}$.

Note that this construction will always result in a minimal set of chains with a size equal to the maximum column weight. This is because the most common element in $\mathcal{S}$ will necessitate a chain for each copy of itself. A minimal set of chains for the Steane code is presented in the right-hand portion of Fig. 7.

Once we have this set of chains, we want to place them in a space of size $|\mathcal{S}|$, making as few cuts to the set of chains as possible, as each chain corresponds to a set of gates that can run in parallel. As we show in Appendix A2, this is NP hard. Therefore, we consider a greedy method where we always place the largest chain (not counting the gaps in the chain) in the first possible position. If a chain does not fit, then we break it into two chains, and return those chains to our set of chains, and continue to place the largest chain remaining in our set. We heuristically break chains with gaps at the gap that is closest to either of the two ends of the chain. If a chain has no gaps, then we break off a single link. The pseudocode for this is outlined in Algorithm 2, for an illustration of how this

ALGORITHM 1. Generation of a minimal set of chains.

Input: An array of integers X , corresponding to the primitive sets.
Output: A 2D array Y , corresponding to the minimal set of chains

Sort X from largest to smallest
 Calculate max_value to be the largest value in X
 Calculate num_chains to be the highest number of times any given element appears in X .
 Initialize Y to a 2D array of size $\text{num_chains} \times \text{max_value}$, each element is equal to -1 .

for x in X **do**
 $\text{index} \leftarrow \text{index}(x) \triangleright$ get index in Y that corresponds to the value x
 $\text{placed} \leftarrow \text{False}$
 for i in $1 : \text{num_chains}$ **do**
 if $Y[i][\text{index}] = -1$ AND NOT placed **then**
 $Y[i][\text{index}] \leftarrow x$
 $\text{placed} \leftarrow \text{True}$
 end if
 end for
end for

Depending on implementation, one may have some chains with leading or trailing -1 's. If desired, remove those.
return Y

transforms a circuit refer to Fig. 8, and an implementation in Julia is available on GitHub [12,13]. The left-hand portion of Fig. 7 shows this method resulting in an optimal assignment for the Steane code.

V. COMPILATION RESULTS

Tables I and II host the number of shuttles required for extracting both the X and Z syndromes of various CSS codes on Shor and naive syndrome circuits, respectively. Furthermore, in Appendix A4, we show through gate-by-gate simulation how the logical error rate decreases with compilation on the toric code.




FIG. 7. Optimal solution to Shor reindexing on Steane code. Using Algorithms 1 and 2, we can find an optimal reordering of qubits for the Steane code's Shor-style syndrome-extraction circuit. On the right, we group the elements into six chains, the maximum column weight of the code. On the left, we show how we can pack these without breaking these chains up, giving us an optimal result of six shuttles. Each output value from mapping values to the ladder corresponds to a particular configuration of the $2 \times n$ architecture. Note: as with Fig. 6, the X and Z checks are presented together for pedagogical reasons.




FIG. 8. Steane code's shor-style syndrome-extraction circuit's two-qubit gates, both before and after specialized Shor-syndrome compilation. Consecutive gates that are in a color-coded group together can run concurrently on the $2 \times n$ architecture (repeated colors have no relation to each other). Furthermore, each grouping corresponds to a particular value of δ . (a) The uncompiled (yet gate shuffled) circuit. There are 14 values for delta in this circuit and therefore 14 shuttles are required to execute this. (b) The circuit after our specialized Shor-syndrome compilation algorithm. Notice that now there are only six values of δ , and therefore this circuit will only require six shuttles to execute.

A. Codes tested

All codes tested were generated with QuantumClifford.jl [14]. Specifically, following the QuantumClifford's error correcting code (ECC) application programming interface (API), stabilizer tableau were obtained, and then the appropriate methods were called to isolate the X and Z checks, and then to generate the syndrome-extracting circuit to be compiled.

(1) Two-dimensional toric and surface codes: $[[18, 2, 3]]$, $[[18, 2, 3]]$, $[[50, 2, 4]]$ 2D toric codes [15] and $[[13, 1, 3]]$, $[[25, 1, 4]]$, $[[41, 1, 5]]$ 2D surface codes [16] with lattice sizes ranging from 3×3 to 5×5 were tested.

(2) Shor's $[[9, 1, 3]]$ code [17].

(3) Steane $[[7, 1, 3]]$ code [18].

(4) 2D color codes: The triangular planar codes with 4.8.8 tiling, given by the instances $[[17, 1, 5]]$, $[[31, 1, 7]]$, and $[[71, 1, 11]]$, and those with 6.6.6 tiling, given by $[[19, 1, 5]]$, $[[37, 1, 7]]$, and $[[91, 1, 11]]$, are considered [19].

(5) Lifted product (LP) codes: a $[[882, 24, d \leq 24]]$ code from Appendix B of Ref. [20] and a $[[175, 19, \leq 10]]$ code from Appendix A of Ref. [21].

(6) Generalized bicycle codes: The $[[48, 6, 8]]$ and $[[254, 28, 14 \leq d \leq 20]]$ code instances from Appendix B of Ref. [22].

(7) Two Block Group Algebra (2BGA) codes: The $[[40, 8, 5]]$, $[[56, 28, 2]]$, and $[[70, 8, 10]]$ abelian 2BGA codes from Table I of Ref. [23], as well as $[[64, 8, 8]]$ non-abelian 2BGA code from Table II of the same work [23].

(8) Multivariate bicycle code: a weight-6 $[[48, 4, 6]]$ code from [24].

(9) Bivariate bicycle codes: The $[[144, 12, 12]]$, $[[360, 12, \leq 24]]$, and $[[756, 16, 34]]$, as reported in Table 3 of Ref. [25].

(10) Trivariate tricycle codes: The $[[72, 6, 6]]$, $[[432, 12, 12]]$, $[[840, 9, 16]]$, and $[[1029, 18, 16]]$, as reported in Table V of Ref. [26].

TABLE II. Different compilations of the naive syndrome circuit for extracting all stabilizers. Same as Table I, except for naive syndrome circuits. Naturally, there are fewer columns, as we developed additional methods specifically for Shor-syndrome extraction.

Code	N	K	D	Uncompiled	GateShuffled	AHR
2DTORIC	18	2	3	(32, 32)	(6, 6)	(11, 10)
2DTORIC	32	2	4	(60, 60)	(6, 6)	(12, 11)
2DTORIC	50	2	5	(96, 96)	(6, 6)	(12, 11)
2DSURFACE	13	1	3	(20, 20)	(5, 6)	(5, 6)
2DSURFACE	25	1	4	(42, 42)	(6, 7)	(6, 7)
2DSURFACE	41	1	5	(72, 72)	(7, 8)	(7, 8)
SHOR	9	1	3	(12, 10)	(8, 4)	(8, 4)
STEANE	7	1	3	(12, 12)	(8, 8)	(7, 7)
2DCOLOR4.8.8	17	1	5	(36, 36)	(13, 13)	(13, 13)
2DCOLOR4.8.8	31	1	7	(72, 72)	(21, 21)	(20, 20)
2DCOLOR4.8.8	71	1	11	(180, 180)	(43, 43)	(38, 38)
2DCOLOR6.6.6	19	1	5	(42, 42)	(13, 13)	(11, 11)
2DCOLOR6.6.6	37	1	7	(90, 90)	(23, 23)	(20, 20)
2DCOLOR6.6.6	91	1	11	(240, 240)	(52, 52)	(47, 47)
LIFTEDPRODUCT	882	24	$d \leq 24$	(2646, 2646)	(13, 13)	(81, 32)
LIFTEDPRODUCT	175	19	$d \leq 10$	(588, 588)	(46, 53)	(103, 55)
GENERALIZEDBICYCLE	48	6	8	(192, 192)	(13, 13)	(38, 36)
GENERALIZEDBICYCLE	254	28	$14 \leq d \leq 20$	(1270, 1270)	(18, 18)	(67, 67)
TWOBLOCKGROUPALGEBRA	70	8	10	(280, 280)	(14, 14)	(42, 40)
TWOBLOCKGROUPALGEBRA	40	8	5	(160, 160)	(17, 17)	(41, 31)
TWOBLOCKGROUPALGEBRA	56	28	2	(224, 224)	(15, 15)	(24, 15)
TWOBLOCKGROUPALGEBRA	64	8	8	(256, 256)	(55, 55)	(67, 66)
MULTIVARIATEBICYCLE	48	4	6	(144, 144)	(16, 16)	(35, 23)
LA-CROSS	100	4	5	(252, 252)	(16, 20)	(16, 20)
LA-CROSS	98	18	4	(294, 294)	(10, 10)	(37, 28)
LA-CROSS	117	9	4	(270, 270)	(16, 22)	(16, 22)
LA-CROSS	296	16	7	(720, 720)	(24, 32)	(24, 32)
LA-CROSS	400	16	8	(1008, 1008)	(28, 36)	(28, 36)
BIVARIATEBICYCLE	144	12	12	(432, 432)	(12, 12)	(21, 68)
BIVARIATEBICYCLE	360	12	$d \leq 24$	(1080, 1080)	(12, 12)	(43, 186)
BIVARIATEBICYCLE	756	16	$d \leq 34$	(2268, 2268)	(19, 19)	(39, 58)
TRIVARIATETRICYCLE	72	6	6	(216, 432)	(28, 53)	(55, 73)
TRIVARIATETRICYCLE	432	12	12	(1296, 2592)	(46, 90)	(86, 217)
TRIVARIATETRICYCLE	840	9	16	(2520, 5040)	(43, 84)	(98, 254)
TRIVARIATETRICYCLE	1029	18	16	(3087, 6174)	(39, 76)	(107, 390)
3DSURFACE	12	1	2	(16, 27)	(10, 9)	(9, 9)
3DSURFACE	51	1	3	(84, 150)	(16, 22)	(16, 33)
4DSURFACE	33	1	4	(83, 84)	(23, 18)	(23, 18)
4DSURFACE	241	1	9	(758, 759)	(72, 57)	(79, 83)
GENERALIZEDTORIC	292	18	8	(876, 876)	(230, 230)	(189, 201)
GENERALIZEDTORIC	392	8	$d \leq 28$	(336, 336)	(95, 95)	(95, 93)
GENERALIZEDTORIC	396	8	$d \leq 36$	(1188, 1188)	(312, 312)	(251, 243)
HOMOLOGICALPRODUCT	117	9	4	(225, 225)	(28, 22)	(43, 52)
HOMOLOGICALPRODUCT	225	9	6	(630, 630)	(38, 32)	(49, 114)
DBLEHOMOLOGICALPRODUCT	486	6	9	(1944, 1944)	(83, 83)	(259, 91)
QUASICYCLIC_GHP	882	48	16	(3528, 3528)	(19, 19)	(102, 49)
QUASICYCLIC_GHP	882	12		(4410, 4410)	(31, 31)	(145, 126)

(11) La-cross codes: The instances $[[100,4,5]]$, $[[117,9,4]]$, $[[296,16,7]]$, and $[[400,16,8]]$ from the $[[n-k)^2 + n^2, k^2, d]]$ La-Cross code family with open boundary conditions, together with the instance $[[98,18,4]]$ from the $[[2n^2, 2k^2, d]]$ La-Cross code family with periodic boundary conditions, are reported in Fig. 3 of Ref. [27].

(12) Quasicyclic generalized Hypergraph Product code: We consider the code instance $[[882,48,16]]$, in which half of the columns have weight 3 and the remaining half have weight

5, as reported in Ref. [22]. Building on this construction, we generate an instance with parameters $[[882, 12, d]]$, where half of the columns have weight 3 and the other half have weight 7. This construction employs a sparse quasi-cyclic polynomial matrix A over the quotient ring $R = \mathbb{F}_2[x]/(x^\ell - 1)$, with each entry being either 0 or a monomial x^i for $0 \leq i < \ell$, where ℓ denotes the circulant size. Each polynomial entry corresponds to a binary $\ell \times \ell$ circulant block. To ensure sparsity, each row and column of A contains at most w

ALGORITHM 2. Shor-syndrome specialized compilation.

Input: A minimal set of chains Y , which can be generated via Algorithm 1. This is an array which consists of num_chains rows, which each contains its own array corresponding to the links and gaps in the chain. -1 denotes a gap in the chain, while other numbers correspond to primitive set values.

Output: A sorted list of numbers whose values are the primitive sets, and whose ordering is the order that when $[1, \dots, s]$ is added, results in an array with as ideally as few distinct values as possible.

sort(Y) $\leftarrow$ sorts Y by the total of number of non -1 elements in each chain, from largest to smallest.

ladder $\leftarrow$ array of size equal to the total number of non -1 elements over all chains in Y , which is the size of our original input, X , in Algorithm 1. Initialize all elements to -1 .

pool $\leftarrow Y$

sort(pool)

while pool is not empty **do**

 chain $\leftarrow$ pop from front of pool. $\triangleright$ This should be the chain with the most non -1 values

 Try to find the first position in ladder such that the non-negative elements in chain line up with -1 elements in ladder. That is, all elements in chain can be assigned sequentially to ladder, starting at the found position, without overwriting any non-negative values.

if such a position was found **then**

 Assign the values in chain to ladder starting from that position.

else

 Split chain into two smaller chains, chain_a and chain_b. If the distance from the beginning of chain to its first -1 is shorter than the distance from the end of chain to its last -1 , then split at the first -1 , otherwise split at the last -1 . (removing the trailing or leading -1). If there are no -1 's, simply let chain_a be the first element, and chain_b be the rest.

 push chain_a and chain_b into pool

 sort(pool)

end if

end while

return ladder

nonzero entries, so that the resulting parity-check matrices are $w + \deg b(x)$ limited. A fixed polynomial $b(x) \in R$, referred to as the circulant polynomial, is used to construct the off-diagonal blocks. For the present code instance, we take $\ell = 63$ and construct a matrix $A \in M_{n \times n}(R)$ with $n = 7$; i.e., A is a 7×7 matrix over R :

$$A = \begin{pmatrix} x^{56} & x^{12} & x^{52} & x^{13} & x^{39} & x^{11} & x^{61} \\ x^{61} & x^{56} & x^{12} & x^{52} & x^{13} & x^{39} & x^{11} \\ x^{11} & x^{61} & x^{56} & x^{12} & x^{52} & x^{13} & x^{39} \\ x^{39} & x^{11} & x^{61} & x^{56} & x^{12} & x^{52} & x^{13} \\ x^{13} & x^{39} & x^{11} & x^{61} & x^{56} & x^{12} & x^{52} \\ x^{52} & x^{13} & x^{39} & x^{11} & x^{61} & x^{56} & x^{12} \\ x^{12} & x^{52} & x^{13} & x^{39} & x^{11} & x^{61} & x^{56} \end{pmatrix}$$

$$b(x) = 1 + x + x^6 \in R$$

The code is then defined via the block parity-check matrices

$$H_X = [A, \quad b(x)I_m], \quad H_Z = [b(x)^T I_n, \quad A^*],$$

where A^* denotes the transpose of A with each entry polynomial-reversed modulo $x^\ell - 1$, and I_m, I_n are identity matrices over R [22].

(13) 3D and 4D surface codes: The instances of the 3D surface code are $[[12,1,2]]$ for $L = 2$ and $[[51,1,3]]$ for $L = 3$, while the 4D surface code instances are $[[33,1,4]]$ for $L = 2$ and $[[241,1,9]]$ for $L = 3$. These codes are constructed via the hypergraph product of $(L - 1) \times L$ repetition code chain complexes. The $[[33,1,4]]$ 4D surface code and the construction method for these codes are described in Ref. [28].

(14) Generalized toric codes on twisted tori: We consider the code instances $[[292,18,8]]$, $[[392,8,\leq 28]]$, and $[[396,8,\leq 36]]$, as summarized in Tables III and IV of Ref. [29].

(15) Homological product codes: We consider the code instances $[[117,9,4]]$ and $[[225,9,6]]$, as reported in Table III of Ref. [30].

(16) Double Homological Product codes: We consider the code instances $[[241,1,9]]$ and $[[486,6,9]]$, as reported in Table I of Ref. [31].

B. Shor compilation

The number of shuttling operations to carry out the fault-tolerant Shor-syndrome-extracting X and Z circuits for these codes under various levels of compilation is displayed in Table I. The first number in the parenthesis corresponds to the value for the X check syndrome circuit, and the second value corresponds to the Z circuit. Levels of compilation are defined as follows:

(1) *Uncompiled*. The number of shuttles to run the original circuit provided by QuantumClifford.jl, reading the circuit from left to right, shuttling whenever is needed to run the next gate.

(2) *GateShuffled*. The number of shuttles required after δ sorting the circuit as discussed in Sec. II A.

(3) *AHR*. The number of shuttles after reindexing the ancilla by the simple sorting heuristics discussed in Sec. III.

(4) *SSSC (Shor syndrome specialized compilation)*. The number of shuttles after reindexing the ancilla by our specialized Shor-syndrome circuit algorithm, discussed in Sec. IV.

(5) *NumChains and Blanks*. The “NumChains” value corresponds to the minimum number of shuttles possible, achievable by inserting “blanks” into the ancilla array. The “Blanks” value is then a naive upper bound on the number of blanks required to achieve this. It was calculated by counting the number of gaps in the chains. Chains were introduced in Sec. IV, and the idea of adding blanks to achieve the minimum number of possible shuttles is discussed further in Sec. VI.

In Table I, we can see that our techniques find quite good shuttling schedules, trivially so on codes with regular column weight (and so SSSC does not add anything). For codes with nonregular column weight, with only some exception, SSSC does further reduce the number of shuttles compared to our heuristic sorting methods. Furthermore, notice that we need not use the same reindexing technique for both the X and Z circuits. For example, on the $[[117,9,4]]$ homological product code tested, we could use AHR to compile the X circuit to run in 12 shuttles, while using SSSC to compile the Z circuit to run in 7 shuttles. This is also exemplified in the $[[117,9,4]]$ La-Cross code, which can run with five shuttles on the X

circuit after compiling via AHR, and then five shuttles on the Z circuit after compiling it with SSSC.

Another thing to notice is that since these methods are taking as input a circuit algorithmically generated via its stabilizer tableau, codes like the 2D toric code compile to different shuttle numbers for its X and Z circuits. This is simply due to the stabilizer tableaux for its X and Z checks being literally different due to how they are generated in QuantumClifford.jl [14]. The manifestation of the 2D toric code's stabilizer generators forms a column-irregular matrix, at least as defined in QuantumClifford.jl, and so optimal compilation is not promised. An interesting avenue for future work might be on techniques that use Gaussian elimination on the stabilizer tableau to work in tandem with our methods that merely permute it.

C. Naive compilation

The number of shuttling operations to carry out the naive syndrome-extracting circuits for these codes is displayed in Table II. The definitions of the types of compilation are the same as for Shor-syndrome extraction. Naturally, there are fewer, as we developed some specialized methods for Shor-syndrome extraction. For most of the codes tested, AHR does not reduce the number of shuttles for their naive syndrome circuit. This is due to naive compilation being a much harder problem than Shor-syndrome compilation, which we show is NP hard in Appendix A2. The AHR methods are after all heuristics and need not be used if the original ordering of ancilla is better. Development of algorithms to improve the performance of more general or other special classes of syndrome-extraction circuits (flag qubit circuits are perhaps more interesting than these naive circuits), we leave to future research.

VI. OPTIMAL COMPILATION OF REGULAR LDPC CODES

So far, we have discussed techniques for general compilation, the special case of Shor-syndrome extraction, as well as simulation of how these techniques can lower pseudothresholds. In this section, we illuminate a family of LDPC codes that can always be reordered in polynomial time such that the number of resulting shuttles in its Shor-syndrome-extraction circuit is perfectly minimized.

First, we notice how a stabilizer tableau or parity check matrix maps to the input S of our ancillary qubit reindexing problem. For example, take the Steane code (treating the X and Z checks together to create a more pedagogical example):

$$\begin{aligned} & \begin{bmatrix} + & & & & X & X & X & X \\ + & & X & X & & & X & X \\ + & X & & X & & X & & X \\ + & & & & Z & Z & Z & Z \\ + & & Z & Z & & & Z & Z \\ + & Z & & Z & & Z & & Z \end{bmatrix} \\ & \rightarrow \begin{bmatrix} & & & 3 & 2 & 1 & 0 \\ & 5 & 4 & & & 1 & 0 \\ 6 & & 4 & & 2 & & 0 \\ & & & 3 & 2 & 1 & 0 \\ & 5 & 4 & & & 1 & 0 \\ 6 & & 4 & & 2 & & 0 \end{bmatrix}. \end{aligned} \quad (3)$$

If we were to build a syndrome-extracting circuit of this code, we would take each row of the Steane code, attributing the first column to the first data qubit (the top of the circuit diagram), the second column to the second qubit on our diagram, and so on, then attach gates to ancilla wherever we find a nonzero entry. For example, the first row would cause four gates to be connected to the fourth, fifth, sixth, and seventh data qubits, respectively. When we apply our ancillary qubit reindexing techniques to this, the ancillary qubits connected to the bottom (seventh) data qubit would have a value of 0 in their primitive set, the ancilla connected to the sixth data qubit would have a 1 in their primitive set, and so on. Thus, in the case of Shor-style syndrome extraction, our set S will be comprised of the entries in the right matrix in Eq. (3) (while in naive syndrome circuits, the primitive sets will correspond to the rows of the right side of that same matrix).

We now notice that if our error correction code has regular column weight w_c (each column has the same number of nonzero entries), then this would mean that all elements that occur in S , occur the same number of times. Because each element occurs w_c times, we can then make w_c identical chains of the form $n-1, n-2, \dots, 0$, each of which has no gaps, and will pack perfectly. Therefore, all LDPC codes with regular column weight can be compiled to run on $2 \times n$ architectures in exactly w_c shuttles. It is worth noting that this would then include all “regular” LDPC codes, as they have constant column and row weight [32]. Furthermore, the $[[144, 12, 12]]$ gross code [25] also has regular column weight and would be an excellent choice of error-correcting code for this architecture, as well as trivariate tricycle codes [26], generalized toric codes on twisted tori [29], lifted product codes and its subfamilies including generalized bicycle and two block group algebra codes [22]. In addition, we observed some irregular codes that also compiled perfectly such as $[[882, 48, 16]]$, and $[[882, 12, ?]]$ quasicyclic generalized hypergraph product codes [22].

Optimal compilation on irregular LDPC codes

Until now, we have viewed the problem of reindexing the ancillary qubits under Shor-style syndrome extraction as a problem that implicitly involves maintaining the tightest possible spacing of the qubits. If we relax this assumption and allow for spaces in the array, then we can trivially achieve optimal compilation in regard to the number of required shuttles by deriving the set of minimal chains, and placing those consecutively on an “extended” staircase to accommodate the gaps in some of the chains. The number of additional quantum dots required to do this will, in the most naive placement scheme, be equal to the sum of the sizes of all gaps in the set of minimal chains. Before, we were maintaining optimal spacing, while trying to minimize shuttles; however, one could instead aim to fix the number of shuttles to optimal (equal to the maximum column weight of the code), while minimizing the spacing. We prove the former to be NP hard in Appendix A2, while we conjecture the latter to also be NP hard.

VII. DISCUSSION AND CONCLUSION

We have demonstrated methods for scheduling the shuttles of qubits in $2 \times n$ one-dimensional architectures, highlighting

a large family of LDPC codes that under Shor-syndrome extraction only ever requires a number of shuttles exactly equal to the code's column weight, which notably includes bivariate bicycle quantum LDPC codes, which were shown in recent work to have a threshold on par with the surface code [25]. Other codes that would make great candidates for this architecture would include: trivariate tricycle codes [26], generalized toric codes on twisted tori [29], lifted product codes and its subfamilies including generalized bicycle and two block group algebra codes [22], as already mentioned in Sec. VI.

The most natural next problem to address would be the state preparation of “interwoven” cat states needed to perform Shor-syndrome extraction. “Interwoven” corresponds to after ancilla reindexing, we might require the ancilla array contain distinct multiple cat states that need to be arranged in the ancilla array such that the physical qubits constituting one of the cat states are not all neighbors to one another (see Fig. 13). Throughout this work, we assumed that we had access to fault-tolerantly prepared cat states to load into the ancilla array; however, the details of how to do and optimize this, especially the process of interweaving, we leave to future work. In Appendix A 3, we propose a sketch for how one could start with an array of consecutive cat states and interweave them, using an additional empty quantum dot array.

For other future work, we plan to investigate the design of more elegant algorithms for compilation, perhaps exploring the trade-off between minimizing the size of the ancilla array and minimizing the number of shuttles. We also hope to consider other paradigms and pieces of the quantum error correction pipeline possible in the specific hardware implementation; for instance, incorporating flag qubits, architectures with more than two rails, how to optimize the performance of logical operations, and the discovery of other family codes that can have their syndrome extracted in ways that require only a few shuttling operations.

Nevertheless, what we have presented here is already a substantial achievement toward embarking on what we believe to be a rich area of research, both from the perspectives of theoretical computer science as well as practical quantum error correction. That is, the rather algorithmic questions raised by optimizing error-correcting codes to run on practical hardware are challenging enough to be interesting, yet seemingly constrained enough to allow for progress to be made. The culmination of which is the practical result of having more efficient quantum error correction schemes specifically tailored to a specific and promising type of quantum hardware. We hope to see this work expanded upon, as well as other works that explore these sorts of optimizations for other qubit platforms as we continue to approach a future of fault-tolerant quantum computation.

ACKNOWLEDGMENTS

We would like to thank our colleagues Filip Rozepedek and Guus Avis for their fruitful discussions regarding gate commutativity. We also thank Nithin Raveendran for his recommendation to see how bivariate bicycle codes compile on this architecture. We would like to acknowledge colleagues Ferdinand Kuemmeth and Fabrizio Berritta as well. This work




FIG. 9. Application of heuristics on the Steane code. Here, we have the primitive sets of the naive syndrome-extraction circuit of the Steane code expressed in block form. We can see that either of the first two heuristics discussed in Sec. III will allow use to find an optimal solution for this case of ancillary qubit reindexing. Note: As this example is pedagogical in nature, this corresponds to reindexing the X and Z syndrome circuits together, something that in practice will be avoided.

was supported by funding from the Inge Lehmann Programme of the Independent Research Fund Denmark, the US Army Research Office (ARO) under Award No. W911NF-24-2-0043, and the NSF under Grants No. 1941583, No. 2346089, and No. 2402861.

DATA AVAILABILITY

The data that support the findings of this article are openly available [13]; embargo periods may apply. The underlying simulator is available as a registered open-source package in the Julia ecosystem: CircuitCompilation2xn.jl [13]. The goal of that package is to allow users to provide any stabilizer code's tableau, or any syndrome-extraction circuit, and have the package compile it for better shuttling performance on the $2 \times n$ architecture.

APPENDIX

In this Appendix, we provide examples of the block stacking visualization on valid error-correcting codes, an illustration of how to translate between the “ladder” and “staircase” notations, a proof for the NP hardness of the ancillary qubit reindexing problem, a discussion of how to prepare the necessary interwoven cat states needed for compiled Shor-syndrome circuits, and a presentation of plotting logical versus physical error rates for different compilation schemes, following a circuit-level noise model.

1. More block visualization examples

Refer to Figs. 9 and 10 to see how the simple heuristics from Sec. III can be applied to the naive syndrome-extracting circuits of Steane's $[[7,1,3]]$ code [18] and Shor's $[[9,1,3]]$ code [17]. In Sec. IV, we introduced a condensed version of the staircase notation that is specifically for Shor-syndrome circuits. Figure 11 illustrates how these two notations are equivalent.




FIG. 10. Application of heuristics on the Shor code. Here, we have the primitive sets of the naive syndrome-extraction circuit of the Shor code expressed in block form. We can see that either of the first two heuristics discussed in Sec. III will allow use to find an optimal solution for this case of ancillary qubit reindexing. Note: As this example is pedagogical in nature, this corresponds to reindexing the X and Z syndrome circuits together, something that in practice will be avoided.

2. Ancillary qubit reindexing is NP hard

We prove here that the assignment problem we devised heuristics for in Secs. III and IV is in fact NP hard. We do so by first noting that if we could solve the general problem of naive-syndrome reindexing defined in Sec. III, then we could also solve the special case of Shor-syndrome reindexing, as Shor-syndrome reindexing is simply the case where all primitive sets are of size 1. Therefore, the complexity of the general problem is at least as hard as the Shor-syndrome reindexing problem. With this in mind, we prove that Shor reindexing is NP hard. First, let us recall the problem definition from Sec. IV:

Shor reindexing problem formulation. Given a multiset of s integers $\mathcal{S} = \{p_1, p_2, \dots, p_s\}$, find a bijection $\pi : [s] \rightarrow [s]$, such that $|\mathcal{G}|_0$ is minimized, which can be written as follows:

$$\min_{\pi} |\mathcal{G}|_0 = \min_{\pi} \left| \bigcup_{i=1}^s \{p_i + \pi(i)\} \right|. \quad (\text{A1})$$

We would like to prove that given an $\mathcal{S}$, with its most common element occurring t times, determining whether or not there exists a bijection π , such that $|\bigcup_i \{p_i + \pi(i)\}| = t + 1$, to be NP hard.

We prove this by reduction from 3-partition. In the 3-partition problem, we are given $3m$ positive integers as a multiset $A = \{a_1, \dots, a_{3m}\}$. The problem is to see if we can partition these into m triples (multisets of size 3), such that each set sums to the same value T , and we may assume that $T/4 < a_i < T/2$ [33]. This is an NP-complete problem so if we can show that any algorithm that could solve our problem generally would also be able to solve arbitrary instances of 3-partition under a polynomial time mapping to our problem, then we know that our problem must be at least as hard as 3-partition-, and therefore NP hard.

For this reduction, we take an arbitrary 3-partition problem and map it to our Shor reindexing problem by mapping the elements of A to a multiset $\mathcal{S}$, where $\mathcal{S}$ is the input to our problem as defined above. This $\mathcal{S}$ can then be decomposed into no fewer than $3m + 1$ sets, each corresponding to a chain

c_i , where $i \in \{0, 1, \dots, 3m\}$. Rather than specifying each element in $\mathcal{S}$, we instead define each c_i below, noting that $\mathcal{S} = \sum_{i=0}^{3m} c_i$ (addition of multisets here is defined as combining them, e.g., $\{1, 2\} + \{2, 3\} = \{1, 2, 2, 3\}$.) Later, we prove that with how our $\mathcal{S}$ is defined, no other set of chains would be feasible in obtaining a solution to the reindexing problem, thus excluding the possibility of a solution to Shor reindexing existing that does not also map to a valid solution of the given 3-partition problem. We define $c_0, c_1, \dots, c_{3m}$ as follows:

(1) First, let

$$a^* = \max_A - \min_A < \frac{T}{2} - \frac{T}{4} \leq \frac{T}{4}. \quad (\text{A2})$$

We use this to define $c_1, \dots, c_{3m}$, noting that all of these chains each correspond to an original element in A :

$$c_i = \{0, 1, \dots, a^* + a_i - 1\}, \forall i \in \{1, \dots, 3m\}. \quad (\text{A3})$$

Each chain is length $a + a^*$, corresponding to each $a \in A$, and each chain has at least the first a^* elements in common.

(2) c_0 is defined as follows:

$$F(i, j) = i(T + 3a^*) + j + a^*(i - 1) + (|S| + T), \quad (\text{A4})$$

$$c_0 = \{(|S| + T), F(i, j) | i \in \{1, \dots, m\}, j \in \{1, \dots, a^*\}\}. \quad (\text{A5})$$

Later, we prove, as they are invoked, that this definition of c_i , where $i \in \{0, 1, \dots, 3m\}$, guarantees the following lemmas:

- (1) c_0 contains m gaps, each of size $T + 3a^*$.
- (2) Between each gap in c_0 , there exists a run of a^* elements.
- (3) It is not possible for any $e \in c_0$ to participate in another chain.
- (4) If we were to construct an alternate collection of $3m$ sets (chains) from elements in the multiset $\sum_{i=1}^{3m} c_i$, then that alternate collection would require at least one chain with a gap in it.

Finally, we notice that 0 occurs $3m$ times, so this construction needs at least $3m$ shuttles. Furthermore, due to Eq. (3), our construction will require a minimum of $3m + 1$ shuttles, and our problem is to determine whether or not a bijection $\pi : [s] \rightarrow [s]$, recall $s = |S| = \sum_{i=0}^{3m} |c_i|$, exists such that

$$\left| \bigcup_{p \in \mathcal{S}} \{p + \pi(i)\} \right| = 3m + 1. \quad (\text{A6})$$

To structure our proof, we first prove that if a solution exists in the 3-partition instance, then a solution of exactly $3m + 1$ shuttles must exist for our problem. We prove afterward that if a solution to our problem exists, then a solution to the 3-partition problem must also exist.

Direction A1. 3-partition $\rightarrow$ Shor reindexing. If a solution to the 3-partition instance exists, then we know that A can be partitioned into m triples such that each triple sums to T . In the Shor reindexing problem, we can place c_0 in the only allowable position (placement refers to the “ladder” picture as in Fig. 7). Invoking Lemma A1, we know that there will be m gaps of size $T + 3a^*$ to place the remaining elements after placing c_0 . For each of the m triples given by the 3-partition solution, we consider the corresponding triple of




FIG. 11. Translation between “staircase” and “ladder” notation. (a) The “staircase” packing problem for some arbitrary circuit. If all primitive sets are of size one, we can condense each chain of blocks to a single number denoting where the nonair block is, and similarly we can represent each stair in the staircase as a single number. (b) The condensed notation for the same problem. We sometimes refer to this notation as a “ladder” in this work. Numbers summing to the same value in the ladder notation correspond to blocks occupying the same column in the staircase notation. For problems where the ladder notation is applicable (primitive sets of size one), both the ladder and staircase notation are equivalent; however, the ladder notation is considerably more concise.

chains; that is, let us consider an arbitrary triple (c_i, c_j, c_k) such that we know $a_i + a_j + a_k = T$. Chain c_i has a length equal to $a_i + a^*$, therefore placing these three chains sequentially would occupy the following amount of space:

$$(a_i + a^*) + (a_j + a^*) + (a_k + a^*) = a_i + a_j + a_k + 3a^* = T + 3a^*.$$

Therefore, each triple of chains would be packable into the gaps that occur within c_0 . Therefore, all chains can be packed into s positions exactly, implying the existence of a $3m + 1$ shuttle solution.

Lemma A1. c_0 contains exactly m gaps, each of size $T + 3a^*$.

Proof. In the construction of c_0 , we notice that the difference of the largest and smallest element is

$$F(m, a^*) - (|S| + T) = m(T + 3a^*) + a^* + a^*(m - 1) = m(T + 4a^*).$$

Therefore, $m(T + 4a^*) - 1$ is the number of spaces between the first and last elements, where the -1 is due to the number of spaces between, for example, 2 and 0 being 1 in this context. As provided by Lemma A2, in those spaces, we have $(m - 1)$ ridges of size a^* , one ridge of size $a^* - 1$. So, there are a total of $(m - 1)a^* + (a^* - 1)$ elements in those spaces, meaning that the total size of all gaps in c_0 is given equal to

$$\begin{aligned} F(m, a^*) - F(1, 1) - 1 - ((m - 1)a^* + (a^* - 1)) \\ = m(T + 4a^*) - 1 - ((m - 1)a^* + (a^* - 1)) \\ = m(T + 4a^*) - 1 - (ma^* - a^* + a^* - 1) \end{aligned}$$

$$\begin{aligned} &= m(T + 4a^*) - 1 - (ma^* - 1) \\ &= m(T + 4a^*) - ma^* \\ &= m(T + 3a^*). \end{aligned}$$

Thus, there is a total of exactly $m(T + 3a^*)$ gaps in c_0 as defined. Now, to prove this corresponds to exactly m gaps, where each is size $T + 3a^*$. We show this by first noticing that aside from the first element, the gaps in c_0 occur on between elements $F(n + 1, 1)$ and $F(n, a^*)$, for some $n \in \{1, \dots, m - 1\}$. The first gap is size:

$$\begin{aligned} F(1, 1) - |S| + T - 1 \\ = (T + 3a^*) + 1 + (|S| + T) - (|S| + T) - 1 \\ = (T + 3a^*). \end{aligned}$$

Then, for all other gaps:

$$\begin{aligned} F(n + 1, 1) - F(n, a^*) - 1 \\ = (n + 1)(T + 3a^*) + 1 + a^*n + (|S| + T) \\ - (n(T + 3a^*) + a^* + a^*(n - 1) + (|S| + T)) - 1 \\ = T + 3a^* + a^*n - (a^* + a^*(n - 1)) \\ = T + 3a^*. \end{aligned}$$

Therefore, the lemma holds, and we have shown that c_0 contains exactly m gaps of each size $T + 3a^*$. ■

Lemma A2. Between each gap in c_0 , there exists a run of a^* elements

Proof. We first notice that if written in ascending order, the elements in c_0 are $|S| + T, F(1, 1), F(1, 2), \dots, F(1, a^*), F(2, 1), \dots, F(m, a^*)$.

Gaps occur between the first and second element and also between $F(n+1, 1)$ and $F(n, a^*)$, as discussed in Lemma A1. Therefore, between $F(n, 1)$ and $F(n, a^*)$ are $a^* - 2$ elements, and a total of a^* elements following the gap defined by the difference of $F(n, 1)$ and $F(n-1, a^*)$. This is true for all $n \in \{2, 3, \dots, m\}$, as well as between the first and second elements, that is, between $F(1, 1)$ and $|S| + T$.

Therefore, each gap in c_0 is followed by a^* elements. ■

Direction A2. Shor reindexing $\rightarrow$ 3-partition. Now, we prove that any $3m+1$ shuttle solution to Shor reindexing would also be a solution to the 3-partition problem. We prove this by proving that our construction could only ever provide a $3m+1$ shuttle solution if our definitions of $c_0, c_1, \dots, c_{3m}$ are followed exactly, which would then mean that each chain, aside from c_0 , represents an element from A encoded in unary, and as discussed in Direction A1, the packing of those chains corresponds to triples from A each summing to T . In other words, we designed our reduction so that technically speaking, the idea of packing chains is not enforced; however, by design, the only possible $3m+1$ shuttle solution will also be the solution to packing our chains as defined, and no other collection of chains would be feasible. To show that no other collection if feasible, we first prove the remaining two lemmas.

Lemma A3. It is not possible for any $e \in c_0$ to participate in another chain.

Proof. We observe that $a^* < \frac{T}{4}$, and that $a < \frac{T}{2}$ for all $a \in A$. Therefore, $\max_{e \in \sum_{i=1}^{3m} c_i} e < \frac{T}{4} + \frac{T}{2}$. The smallest value in c_0 is $|S| + T$, so if the largest element in $\sum_{i=1}^{3m} c_i$ and the smallest element in c_0 were in the same chain, a gap of size at least $|S| + T/4 - 1$ would be created. Because we are looking for a bijection π that maps from $1, \dots, |S|$ to $1, \dots, |S|$, the longest possible chain (including its gaps) would be limited to length $|S|$, and the largest gap possible would be $|S| - 2$. Therefore, because $|S| + T/4 - 1 > |S| - 2$, as $T/4 > 0$, it would not be possible for any element of c_0 to participate in one of the other chains. ■

Lemma A4. If we were to construct an alternate collection of $3m$ sets (chains) from elements in the multiset $\sum_{i=1}^{3m} c_i$, then that alternate collection would require at least one chain with a gap in it. (A chain with a gap in it refers to a nonconsecutive set of integers.)

Proof. First, we notice that all chains $c_1, \dots, c_{3m}$ have the first a^* elements in common by definition. Because chains are defined as sets and not multisets, an alternate set of $3m$ chains would also need to start with same first a^* elements in common. Furthermore, we notice that because all chains have the first element in common, for an alternate set $c'_1, \dots, c'_{3m}$ to be distinct from $c_1, \dots, c_{3m}$, it would need to contain some chains with gaps in them. This is because chains without gaps correspond to values in A plus a constant a^* , encoded in unary. An alternate set that also did not have any gaps would correspond to the same multiset of numbers in A encoded in unary, and therefore $\{c_1, \dots, c_{3m}\}$ and $\{c'_1, \dots, c'_{3m}\}$ would be equivalent.

To illustrate this further, consider our input multiset $\sum_{i=1}^{3m} c_i$, and consider creating $3m$ empty sets that will be our chains. If we take all the 0's from our input and distribute them among these chains, and then 1's and do the same, and

so on, eventually we will arrive at a number n_1 with less than $3m$ copies. The chains that do not receive a copy of n_1 will no longer need to be considered if we are interested in generating sets of chains without gaps (as the lack of n_1 means we cannot add further elements without inducing a gap). If we repeat this process until all values in our input have been assigned to chains, then we naturally will be left with a collection of chains that have lengths equal to the values in A plus the constant a^* . It should be straightforward to see there is only one collection of gap-free chains possible in this problem, where a collection has no ordering to it.

Therefore, we can say that any set of $3m$ chains partitionable from $\sum_{i=1}^{3m} c_i$ yet distinct from $\{c_1, \dots, c_{3m}\}$ would need to include some gaps in them. ■

We use the above lemmas to prove that no other collection of chains would result in a $3m+1$ solution. For the sake of contradiction, let's assume that such an alternate set exists. By Lemma A3, we know that c_0 must remain unchanged, as c_0 constitutes 1 shuttle, and the other $3m$ shuttles must come from the other elements, as 0 occurs $3m$ times. Therefore, this alternate set of chains would be $\{c_0, c'_1, \dots, c'_{3m}\}$, and a gap must occur at least once in a chain in $\{c'_1, \dots, c'_{3m}\}$ by Lemma A4. By Lemma A1, we know there are m gaps of size $T + 3a^*$, and thus to have a $3m+1$ solution, we must be able to pack $\{c'_1, \dots, c'_{3m}\}$ into those m gaps of size $T + 3a^*$. Therefore, the gaps introduced by using an alternate set of chains must be filled by other chains. As all gaps are by definition of size less than a^* , the ridges (Lemma A2) within c_0 cannot be used to fill them, which then leaves the only possibility to be that within each of the m gaps in c_0 , other chains would fill gaps. However, the structure of any possible packing of chains from $\{c'_1, \dots, c'_{3m}\}$ into a triple (c_i, c_j, c_k) would be a run of at least a^* elements, followed potentially by a region of gaps with possible some elements in between (all of those gaps are size less than a^*), followed again by a run of at least a^* elements, another region of gaps, and so on. This corresponds to each chain beginning with a consecutive sequence of at least a^* elements. Because of this, it is not possible for the beginning of say chain c_j to fill the gaps present in chain c_i . Therefore, this alternate set of chains would not give way to a valid $3m+1$ solution, as it would be impossible to pack, and therefore there will not exist a bijection that satisfies Eq. (A6). Therefore by contradiction, the only possible $3m+1$ solution is to use the initially defined sets $\{c_0, c_1, \dots, c_{3m}\}$.

Therefore, if we have a $3m+1$ solution to Shor reindexing, then it follows that we have successfully packed the elements of A plus some constant a^* , encoded in unary, into m gaps, each of size $T + 3a^*$. From this point, we can recover a solution to the arbitrary 3-partition problem instance.

Therefore, Shor reindexing is at least as hard as 3-partition. Therefore Shor reindexing is NP hard.

3. Preparation of cat states

Here, we present a possible simple algorithm for reordering a linear chain of qubits, provided that we have access to another blank row of qubits, which can then be applied to generating the interwoven cat states that go along with compiled




FIG. 12. Illustration of the full hardware running a Shor-syndrome circuit. (a) After performing syndrome extraction by shuttling our qubits into the proper configurations and applying the appropriate two-qubit gates, the ancilla must then be shuttled to an area where quantum measurement can be carried out. Here, it is represented as a continuation of the 1D architecture; however, this is mostly illustrative. (b) After generating the required ancilla in their proper interwoven cat states, the ancilla are then shuttled into the data qubit area where we perform syndrome extraction, shuttling to different values of δ , as described in the main text. (c) In this area, we propose using two rows of quantum dots to generate the interwoven cat states required by how we compiled the syndrome-extraction circuit.

Shor-syndrome circuits, as we assume that nearest-neighbor cat states can be generated.

Referring to Fig. 13(a), we start with an array of adjacent cat states. After applying the bijection π found by compilation to the ancillary qubit indices, we will often need to reorder this initial cat state such that each stabilizer check is still being interacted with the correct state. Figure 13(b) is an example arbitrary ordering we might need to obtain. To achieve this, we generate the cat states in an adjacent array of quantum dots, as shown in Fig. 12(c). We now can read the labels of the qubits from left to right, shuttling to the main array the largest consecutive sequence that starts at 1 [the first step of this is depicted in Fig. 12(c)]. Next, we can shuttle these qubits to the left and out of the way of the remaining qubits in the top row, so that we can repeat this process until all ancillary qubits on present on the bottom row of qubits, in the proper order. Each time we shuttle down qubits, we start the sequence at one more than the previous sequence's last element. For example, again in Fig. 12(c), the second step would shuttle down qubits 3 and 4, the third step would shuttle down 5, 6, 7, and the final would bring down qubit 8.




FIG. 13. Illustration of the cat state preparation problem. (a) An initial collection of cat states, generated utilizing nearest neighbor connectivity. (b) The order of cat states necessitated by the compilation of a Shor-syndrome circuits. As this is the desired order, we enumerates these consecutively. “Un-weaving” this gives us the numbers present in panel (a), and now all we need to do is use empty quantum dots to sort our initial sequence of cat states.

4. Simulation results

Through simulation, we show that our methods of circuit compilation can be used to reduce the threshold of various error-correcting codes. We model Pauli noise with highly efficient Pauli-frame simulation techniques [34,35] implemented in Julia [14]. We simulate at least 20 000 samples for each set of noise parameter values (each data point in our plots) in order to calculate the logical error rate for a single round of error correction. The simulated circuits contain the following steps:

(1) *Encoding*. For the encoding step, one can use either an encoding circuit or measure the stabilizers to project onto a stabilized space, tracking initial corrections as needed (in Pauli-frame simulation, these corrections can be ignored). As the latter is more resilient to error, that is how we carry out




FIG. 14. Performance of toric 3×3 code's Shor-syndrome circuit on $2 \times n$ architecture. The y axis is $\log_{10}$ of the logical error rate, and the x axis is $\log_{10}(p_{\text{mem}})$. The different colors correspond to different levels of compilation, noting that “Naive compilation” corresponds to running canonical syndrome-extraction circuits without any specialization to this hardware. “All to all” corresponds to setting p_{wait} and p_{shuttle} to 0. This is only the plot of the logical X error rate, although the plot for the logical Z error rate looks quite similar. See Fig. 15 for the logical Z error rate, as well as the plots of this code's performance on a naive syndrome-extraction circuit.




FIG. 15. Compilation performance of the 3×3 toric code. The toric code's naive syndrome-measuring circuit is an example of where the heuristics break down a little. They still perform better than doing nothing, but as shown here, and in Table II, gate shuffling alone finds a decent solution of 12 shuttles regardless of the lattice size, whereas the ancillary qubit reindexing heuristics find a number of shuttles that grows with lattice size.

encoding in our simulations, although the same result can be achieved by simulating an error-free encoding circuit.

(2) *Memory noise.* After encoding, we apply a round of Pauli noise on each data qubit with independent probability p_{mem} . This is the typical place to insert errors when evaluating a code or decoder in nonfault-tolerant simulation.

(3) *Syndrome measurement.* Syndrome measurement will consist of a few pieces. First is the syndrome circuit itself, which is the output of our compilation algorithm (our algorithm transforms a syndrome circuit into an equivalent syndrome circuit). Then, we initialize a variable $\delta_{\text{current}} = 0$

to track the current state of the $2 \times n$ hardware. If the next gate $\mathcal{G}_{\text{next}}$ to be run in the syndrome circuit has $\delta(\mathcal{G}_{\text{next}}) \neq \delta_{\text{current}}$, then we simulate the noise from a shuttle by applying Pauli Z errors independently to each of the idling data qubits with probability p_{wait} , and likewise we apply Pauli noise to each ancilla qubit with probability p_{shuttle} . Afterword, we set $\delta_{\text{current}} = \delta(\mathcal{G}_{\text{next}})$. Furthermore, we apply Pauli noise on both qubits after each two-qubit gate with probability p_{gate} .

(4) *Logical operator measurement.* In order to have something to compare our after-decoding results to, we measure the logical XorZ operators at the end of each syndrome circuit, as

if it were an additional parity check. More specifically, we cannot measure both operators in the same circuit, so we need to run two simulations in order to obtain both the logical X and Z error rates.

(5) *Decoding.* We now run our decoder of choice on the syndromes, and see if the predicted error would cause the same logical error reported in the logical X/Z syndrome. If we predicted the logical error correctly, then that counts as successfully decoded. The logical error is then defined as the ratio of successfully decoded samples to total samples. A lookup table decoder is used for small ($d = 3$) codes such as Steane's $[[7,1,3]]$ code [18], and minimum-weight perfect matching (MWPM) [36] is used for toric codes.

In our plots, we vary p_{mem} along the horizontal axis and use constant values for other types of error. Specifically, we use a value of $p_{\text{shuttle}} = 0.0001$ from Ref. [4], a potentially near-term achievable value of $p_{\text{gate}} = 0.9995$, and $p_{\text{wait}} = 1 - e^{-\frac{(2\mu\text{s}+12.5\mu\text{s})}{28\text{ms}}}$. The shuttle time was said to be $2\mu\text{s}$ ramp and

the fastest time between shuttles was stated to be $12.5\mu\text{s}$ in Ref. [4]. Twenty-eight microseconds was the T_2 time reported in Ref. [37] (recently, great effort has gone into improving to these times [3]). These near-term parameters are represented by the triangles in Fig. 14. Notice that these curves have error floors due to constant error terms that do not depend on p_{mem} . However, the curves plotted by circles correspond to scaling p_{wait} , p_{shuttle} , and p_{gate} each by $10 \times p_{\text{mem}}$, creating curves with pseudothresholds, as one may be used to.

Figure 15 presents the compilation performance of both naive and Shor-syndrome circuits for the 3×3 toric code. These plots are similar to the style of Fig. 14, except we now have both the logical X and Z error rates corresponding to the top and bottom rows of the figure, respectively. The left column corresponds to the performance on naive syndrome circuits, while the right column corresponds to performance on fault-tolerant Shor-syndrome circuits.

-
- [1] A. Zwerver, Scaling spin qubits in quantum dots more—distant—industrial, Dissertation (TU Delft), Delft University of Technology, 2022, <https://doi.org/10.4233/uuid:d09dedd5-de2b-4da9-b17f-5a61bc439805>.
 - [2] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, and J. R. Petta, Semiconductor spin qubits, *Rev. Mod. Phys.* **95**, 025003 (2023).
 - [3] C.-A. Wang, V. John, H. Tidjani, C. X. Yu, A. S. Iyev, C. Déprez, F. van Riggelen-Doelman, B. D. Woods, N. W. Hendrickx, W. I. L. Lawrie, L. E. A. Stehouwer, S. D. Oosterhout, A. Sammak, M. Friesen, G. Scappucci, S. L. de Snoo, M. Rimbach-Russ, F. Borsoi, and M. Veldhorst, Operating semiconductor quantum processors with hopping spins, *Science* **385**, 447 (2024).
 - [4] A. M. J. Zwerver, S. V. Amitonov, S. L. de Snoo, M. T. Mądzik, M. Rimbach-Russ, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, Shuttling an electron spin through a silicon quantum dot array, *PRX Quantum* **4**, 030303 (2023).
 - [5] F. Ansaloni, A. Chatterjee, H. Bohuslavskiy, B. Bertrand, L. Hutin, M. Vinet, and F. Kuemmeth, Single-electron operations in a foundry-fabricated array of quantum dots, *Nat. Commun.* **11**, 6399 (2020).
 - [6] F. Ansaloni, H. Bohuslavskiy, F. Fedele, T. Rasmussen, B. Brovang, F. Berritta, A. Heskes, J. Li, L. Hutin, B. Venitucci, B. Bertrand, M. Vinet, Y.-M. Niquet, A. Chatterjee, and F. Kuemmeth, Gate reflectometry in dense quantum dot arrays, *New J. Phys.* **25**, 033023 (2023).
 - [7] A. Siegel, A. Strikis, and M. Fogarty, Towards early fault tolerance on a $2 \times n$ array of qubits equipped with shuttling, *PRX Quantum* **5**, 040328 (2024).
 - [8] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, *PRX Quantum* **2**, 040101 (2021).
 - [9] D. Gottesman, Fault-tolerant quantum computation with constant overhead, *Quantum Inf. Comput.* **14**, 1338 (2024).
 - [10] A. Siegel, Z. Cai, H. Jnane, B. Koczor, S. Pexton, A. Strikis, and S. Benjamin, Quantum snakes on a plane: Mobile, low dimensional logical qubits on a 2D surface, *PRX Quantum* (2025).
 - [11] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of 37th Conference on Foundations of Computer Science* (IEEE, Burlington, VT, USA, 1996), pp. 56–65.
 - [12] <https://github.com/amicciche/CircuitCompilation2xn>.
 - [13] A. Micciche and S. Krastanov, CircuitCompilation2xn, Zenodo (2024), <https://doi.org/10.5281/zenodo.14320748>.
 - [14] QuantumClifford Contributors, QuantumSavory/QuantumClifford.jl: v0.10.1 (v0.9.12), Zenodo (2024), <https://doi.org/10.5281/zenodo.13950271>.
 - [15] A. Y. Kitaev, Quantum error correction with imperfect gates, in *Quantum Communication, Computing, and Measurement*, edited by O. Hirota, A. S. Holevo, and C. M. Caves (Springer US, Boston, MA, 1997), pp. 181–188.
 - [16] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, *arXiv:quant-ph/9811052*.
 - [17] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493 (1995).
 - [18] A. Steane, Multiple particle interference and quantum error correction, *Proc. R. Soc. London A* **452**, 2551 (1996).
 - [19] A. J. Landahl, J. T. Anderson, and P. R. Rice, Fault-tolerant quantum computing with color codes, *arXiv:1108.5738*.
 - [20] J. Roffe, L. Z. Cohen, A. O. Quintavalle, D. Chandra, and E. T. Campbell, Bias-tailored quantum LDPC codes, *Quantum* **7**, 1005 (2023).
 - [21] N. Raveendran, N. Rengaswamy, F. Rozpędek, A. Raina, L. Jiang, and B. Vasić, Finite rate qLDPC-GKP coding scheme that surpasses the CSS hamming bound, *Quantum* **6**, 767 (2022).
 - [22] P. Panteleev and G. Kalachev, Degenerate quantum LDPC codes with good finite length performance, *Quantum* **5**, 585 (2021).
 - [23] H.-K. Lin and L. P. Pryadko, Quantum two-block group algebra codes, *Phys. Rev. A* **109**, 022407 (2024).
 - [24] L. Voss, S. J. Xian, T. Haug, and K. Bharti, Multivariate bicycle codes, *Phys. Rev. A* **111**, L060401 (2025).
 - [25] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead

- fault-tolerant quantum memory, *Nature (London)* **627**, 778 (2024).
- [26] A. Jacob, C. McLauchlan, and D. E. Browne, Single-shot decoding and fault-tolerant gates with trivariate tricycle codes, [arXiv:2508.08191](https://arxiv.org/abs/2508.08191).
- [27] L. Pecorari, S. Jandura, G. K. Brennen, and G. Pupillo, High-rate quantum LDPC codes for long-range-connected neutral atom registers, *Nat. Commun.* **16**, 1111 (2025).
- [28] N. Berthussen, J. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, D. Gresh, N. Hewitt, M. Mills, S. A. Moses, B. Neyenhuis, P. Siegfried, and D. Hayes, Experiments with the four-dimensional surface code on a quantum charge-coupled device quantum computer, *Phys. Rev. A* **110**, 062413 (2024).
- [29] Z. Liang, K. Liu, H. Song, and Y.-A. Chen, Generalized toric codes on twisted tori for quantum error correction, *PRX Quantum* **6**, 020357 (2025).
- [30] Q. Xu, H. Zhou, G. Zheng, D. Bluvstein, J. P. B. Ataiades, M. D. Lukin, and L. Jiang, Fast and parallelizable logical computation with homological product codes, *Phys. Rev. X* **15**, 021065 (2025).
- [31] E. T. Campbell, A theory of single-shot error correction for adversarial noise, *Quantum Sci. Technol.* **4**, 025006 (2019).
- [32] Regular LDPC code, in *The Error Correction Zoo*, edited by V. V. Albert and P. Faist (2023), https://errorcorrectionzoo.org/c/regular_ldpc.
- [33] E. D. Demaine and M. L. Demaine, Jigsaw puzzles, edge matching, and polyomino packing: Connections and complexity, *Graphs Combin.* **23**, 195 (2007).
- [34] C. Gidney, Stim: A fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [35] E. Knill, Quantum computing with realistically noisy devices, *Nature (London)* **434**, 39 (2005).
- [36] O. Higgott and C. Gidney, Sparse blossom: Correcting a million errors per core second with minimum-weight matching, *Quantum* **9**, 1600 (2025).
- [37] M. Veldhorst, J. C. C. Hwang, C. H. Yang, A. W. Leenstra, B. de Ronde, J. P. Dehollain, J. T. Muhonen, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, An addressable quantum dot qubit with fault-tolerant control-fidelity, *Nat. Nanotechnol.* **9**, 981 (2014).