Subsystem Codes with High Thresholds by Gauge Fixing and Reduced Qubit Overhead

Oscar Higgott[*] and Nikolas P. Breuckmann[†]

Department of Physics & Astronomy, University College London, WC1E 6BT London, United Kingdom

We introduce a technique that uses gauge fixing to significantly improve the quantum-error-correcting performance of subsystem codes. By changing the order in which check operators are measured, valuable additional information can be gained, and we introduce a new method for decoding which uses this information to improve performance. Applied to the subsystem toric code with three-qubit check operators, we increase the threshold under circuit-level depolarizing noise from 0.67% to 0.81%. The threshold increases further under a circuit-level noise model with small finite bias, up to 2.22% for infinite bias. Furthermore, we construct families of finite-rate subsystem low-density parity-check codes with three-qubit check operators and optimal-depth parity-check measurement schedules. To the best of our knowledge, these finite-rate subsystem codes outperform all known codes at circuit-level depolarizing error rates as high as 0.2%, where they have a qubit overhead that is 4.3× lower than the most efficient version of the surface code and 5.1× lower than the subsystem toric code. Their threshold and pseudo-threshold exceeds 0.42% for circuit-level depolarizing noise, increasing to 2.4% under infinite bias using gauge fixing.

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

The realization of scalable quantum computing depends on our ability to correct errors which arise due to inevitable interactions between the device and the environment. These errors can be corrected by introducing redundancy in the form of quantum-error-correcting codes. The most widely studied quantum-error-correcting code, both theoretically and experimentally, is the surface code [1], which has a high tolerance for realistic circuit-level noise and uses four-qubit measurements that are geometrically local in two dimensions.

Despite these advantages, the surface code has several shortcomings. First, it is estimated that thousands of physical qubits will be required to encode each logical qubit for fault-tolerant quantum computing in a noise regime of practical interest [2]. While there has been significant progress in the construction of families of codes called quantum low-density parity-check (LDPC) codes that have improved theoretical parameters relative to the surface code [3–8], none of these codes have been shown to have a lower qubit overhead than the surface code once

circuit-level noise is taken into account. Second, even weight-4 check operator measurements can be too large for some architectures. In some superconducting qubit architectures, for example, the degree of the surface-code interaction graph can lead to frequency collisions [9]. Finally, physical error rates observed experimentally in devices are still above the threshold [10], and the standard implementation of the surface code is not well suited to handle biased noise models that can arise in some physical systems [11]. As a result, improving the tolerance of the surface code to biased noise models is an active area of research [12–15].

In this work, we tackle all three of these problems by introducing new decoding techniques and constructions for subsystem codes. Most quantum codes considered in the literature are stabilizer codes, which are defined in terms of a set of Pauli operators [16]. Subsystem codes are a slight generalization of stabilizer codes where only a subset of the available encoded degrees of freedom are used [17]. They can simplify the measurements which are part of the error-correction procedure by reducing the number of physical qubits involved [18], or by enabling bare-ancilla fault tolerance even when the check weights are large [19]. Furthermore, subsystem codes allow for a procedure called *gauge fixing* which is useful to manipulate the encoded quantum information. Gauge fixing effectively allows us to change the code midcomputation, and in Ref. [20] the authors exploit this to switch between codes which have complementary sets of logical operations.

These advantages have motivated experimentalists to pursue subsystem codes for implementing fault-tolerant

[*]oscar.higgott.18@ucl.ac.uk
[†]n.breuckmann@ucl.ac.uk

quantum computation. This includes IBM, who plan to implement the heavy-hexagon subsystem code [9] to reduce frequency collisions in their superconducting quantum processors [21]. Notably, the Bacon-Shor subsystem code [22] has recently been implemented experimentally in a trapped-ion architecture [23], where the fidelity of the encoded logical operations exceeded that of the entangling physical operations used to implement them.

However, subsystem codes have usually had lower thresholds, an issue which can be attributed to their higher-weight stabilizers. They have also typically had a larger qubit overhead, since some logical qubits are not used to encode information. We introduce constructions and decoding techniques that instead demonstrate that subsystem codes can be used to *increase* the threshold and *reduce* the qubit overhead. The decoding technique we introduce, called schedule-induced gauge fixing, improves the error-correcting performance of a wide class of subsystem codes, especially under biased noise models. By changing the order in which check operators are measured, valuable additional information can be gained, and we introduce a new method for decoding which uses this information to improve performance. In previous work, the Bacon-Shor code has been used as a template to construct elongated compass codes, which can be tailored to biased noise models [24]. However, this method requires changing interactions at the hardware level, as well as measuring high-weight stabilizers directly, since elongated compass codes are not subsystem codes themselves. In contrast, our technique can be implemented entirely in software, and requires only measuring the low-weight gauge operators of the code. In essence, schedule-induced gauge fixing allows us to switch repeatedly between different codes such that more information can be inferred about potential errors.

We also reduce the qubit overhead for quantum-error correction by introducing a construction for subsystem codes that encode a number of logical qubits k proportional to the number of physical qubits n, while using only three-qubit check operators. These codes are derived from hyperbolic tessellations, and we use the symmetry group of the tessellation to derive quantum circuits for measuring the check operators that use only four time steps, which is optimal. From simulating their performance in circuit-level depolarizing noise, we find that these finite-rate subsystem codes have a qubit overhead that is $4.3\times$ lower than the most efficient version of the surface code for error rates as high as 0.2%, which is a noise regime often considered for practical surface-code quantum computing [25]. The potential advantages of quantum LDPC codes have previously been shown only under a simplistic phenomenological noise model [26], in some cases at very low error rates [27]. Once circuit-level noise is taken into account, the potential reduction in qubit overhead can be lost [28]. Therefore, to the best of our knowledge, the results for our finite-rate subsystem codes are the first demonstration of a quantum

code outperforming the surface code in a practical regime of circuit-level depolarizing noise.

In Sec. II, we review the stabilizer formalism, subsystem codes, and gauge fixing before reviewing the subsystem surface code in Sec. III in the context of our more general subsystem code construction. We introduce our construction for finite-rate subsystem LDPC codes in Sec. IV, where we analyze their properties and show how to construct efficient stabilizer measurement circuits for them. In Sec. V, we introduce schedule-induced gauge fixing, our technique for improving the quantum-error-correcting performance of subsystem codes. We present our numerical results in Sec. VI, which includes the application of schedule-induced gauge fixing to the subsystem toric code, as well as a performance analysis of our finite-rate subsystem LDPC codes. We discuss broader applications of schedule-induced gauge fixing and our constructions in Sec. VII, before concluding in Sec. VIII.

II. PRELIMINARIES

A quantum stabilizer code is defined by an Abelian subgroup $\mathcal{S}$ of the Pauli group operating on n physical qubits. The code space is the common $+1$ eigenspace of all elements of the stabilizer group. If there exists a generating set of $\mathcal{S}$ such that each generator acts nontrivially on the physical qubits as either Pauli X or Pauli Z, only then the code is called a *Calderbank-Shor-Steane* (CSS) code [29,30].

A subsystem code is a stabilizer code in which a subset of logical operators is chosen not to store information [17]. In a subsystem code, the overall Hilbert space $\mathcal{H}$ can be decomposed as

$$\mathcal{H} = (\mathcal{H}_{\mathcal{L}} \otimes \mathcal{H}_{\mathcal{G}}) \oplus \mathcal{C}^{\perp}, \tag{1}$$

where only $\mathcal{H}_{\mathcal{L}}$ stores information and any operations applied only on $\mathcal{H}_{\mathcal{G}}$ are ignored. The Pauli operators that act trivially on $\mathcal{H}_{\mathcal{L}}$ form the *gauge group* $\mathcal{G}$ of the code. The stabilizer group $\mathcal{S}$ is the center of $\mathcal{G}$ up to phase factors, $\langle iI, \mathcal{S} \rangle = Z(\mathcal{G}) := C(\mathcal{G}) \cap \mathcal{G}$. Hence, up to phase factors, operators from $\mathcal{G}$ are either stabilizers (acting trivially on $\mathcal{H}_{\mathcal{L}} \otimes \mathcal{H}_{\mathcal{G}}$), or act nontrivially on $\mathcal{H}_{\mathcal{G}}$ only. Logical operators that act nontrivially only on $\mathcal{H}_{\mathcal{L}}$ are called *bare* logical operators $\mathcal{L}_{\text{bare}}$ and are given by $C(\mathcal{G}) \backslash \mathcal{G}$. The *dressed* logical operators $\mathcal{L}_{\text{dressed}} = C(\mathcal{S}) \backslash \mathcal{G}$ act nontrivially on both $\mathcal{H}_{\mathcal{L}}$ and $\mathcal{H}_{\mathcal{G}}$. A dressed logical operator is a bare logical operator multiplied by a gauge operator in $\mathcal{G} \backslash \mathcal{S}$. The distance d of a subsystem code is the weight of the minimum-weight dressed logical operator, $d = \min_{P \in C(\mathcal{S}) \backslash \mathcal{G}} |P|$. The number of physical qubits n, logical qubits k, independent stabilizer checks r, and gauge qubits g are related as

$$n - k = r + g. \tag{2}$$

One advantage of introducing gauge qubits is that they can enable simpler stabilizer measurements if the

generators of the gauge group $\mathcal{G}$ (the *gauge generators*) have a lower weight than the generators of the stabilizer group $\mathcal{S}$. Since $\mathcal{S} \subseteq \mathcal{G}$, the outcomes of the gauge generator measurements can be used to infer the eigenvalues of the stabilizers, provided the gauge generators are measured in the appropriate order (since $\mathcal{G}$ is generally not Abelian) [29,30]. We refer to standard stabilizer codes, where all logical qubits are used to store quantum information, as *subspace* codes, to distinguish them from subsystem codes.

The technique called *gauge fixing* applied to a subsystem code consists of adding an element $g \in \mathcal{G}$ into the stabilizer group $\mathcal{S}$, as well as removing every element $h \in \mathcal{G}$ that anticommutes with g. Gauge fixing was introduced by Paetznick and Reichardt and can be useful for performing logical operations [20,31,32], including for code deformation and lattice surgery [33]. Gauge fixing can also be used for constructing codes: Both the surface code and the heavy-hexagon code [9] are gauge fixings of the Bacon-Shor subsystem code [22], all belonging to the larger family of 2D compass codes [24]. These constructions are static, as the fixed gauge stays the same over time. In this work, we show how gauge fixing can be used to improve the quantum-error-correcting performance of subsystem codes. Here, we consider a dynamical approach to gauge fixing; i.e., in contrast to the code constructions mentioned earlier, we change which gauge degrees of freedom are being fixed over time. This allows us to improve the error-correction capabilities of subsystem codes.

III. THE SUBSYSTEM SURFACE CODE

We now describe a method for constructing subsystem codes from hexagonal lattices, which we see is equivalent to the subsystem toric code of [34]. In Sec. IV, we generalize this construction to other tessellations to obtain subsystem *hyperbolic* codes.

Take a torus which is subdivided into hexagons. The quantum stabilizer code associated with this lattice is the hexagonal toric code constructed by placing a data qubit on each edge of the lattice and associating each face and vertex with a Z and X stabilizer, respectively. Now consider a face f of the lattice and two of its vertices v and u that are not direct neighbors and do not have any neighbors in common either. We identify the vertices u and v and call the new vertex w. This deforms the face f into a shape like a bow tie with w in the center (see Fig. 1). Any edge which was incident to either v or u before is incident to w after this identification.

There is a canonical subdivision of a bow-tie-shaped face: We can simply consider either half. Similarly, the neighborhood of a merged vertex in the middle of the bow tie can be subdivided into two disjoint sets. In terms of the associated quantum code, there is a canonical way to *break* the X and Z checks associated with the vertices and faces (see Fig. 1, right). The four operators obtained by the breaking procedure are gauge operators and do not
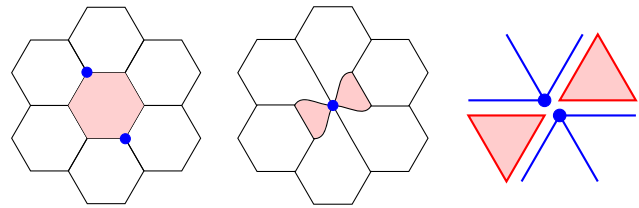

FIG. 1. Left and middle: merging inside a hexagonal lattice. After merging, the resulting vertex has degree 6. Note that the surrounding faces are unaffected (besides being deformed). Right: after merging we can break the X check (blue) and Z check (red) into two pairs of operators. These operators all have weight 3. Operators of different types pairwise anticommute, but they commute with all remaining stabilizers in the lattice.

commute. Importantly, the gauge operators of a bow tie, while not commuting among themselves, commute with all other check operators. By merging the upper left and lower right vertices of each hexagon, as shown in Fig. 1, we obtain the subsystem toric code of Ref. [34] shown as a tiling by bow ties in Fig. 2 (left).

An alternative representation of the subsystem toric code can be found by placing a qubit on the middle of each edge and on each vertex of a square tiling.

Each gauge operator is now represented by a triangle, with a qubit associated with each of its vertices. This lattice can be obtained from the bow-tie lattice by substituting edges for vertices. Borrowing terminology from Ref. [34], each gauge generator is referred to as a *triangle operator*, and consists of a Pauli operator acting nontrivially on its three qubits. There are four types of triangle operators in each face of the square lattice: two Z-type triangle operators defined in the northwest and southeast corners, and two X-type operators defined in the northeast and southwest corners. These four types of triangle operators are highlighted in Fig. 3 for the $L = 2$ subsystem toric code. Within each face, the product of each pair of Z-type triangle operators forms a six-qubit Z stabilizer, and the product of each pair of X-type triangle operators forms a six-qubit X-type stabilizer. The subsystem toric code has $3L^2$ data
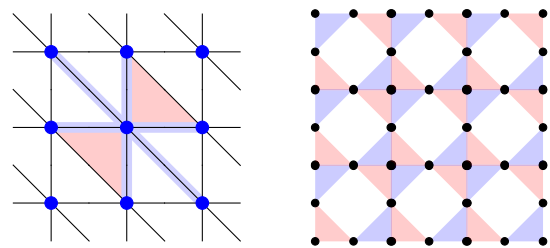

FIG. 2. Left: merging top left and bottom right vertices of all faces of a hexagonal tessellation leads to a tiling of bow ties. The X and Z stabilizers belonging to the merged vertex in the center are highlighted in blue and red. Both are weight-6 operators. Right: we can redraw the lattice by exchanging edges with vertices representing the broken stabilizers as triangles.
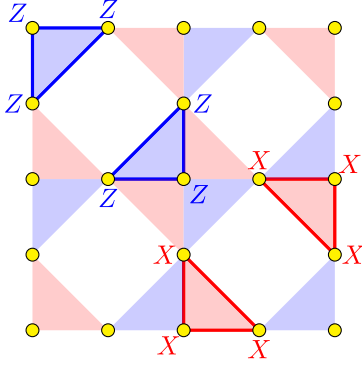
FIG. 3. The subsystem toric code of Ref. [34]. Data qubits (yellow-filled circles) are placed in the middle of each edge and on each vertex of a square lattice of the toric code. Opposite sides are identified. The gauge group is generated by three-qubit *triangle operators*. The two Z triangle operators in the top left face are outlined with a blue border, and their product forms a six-qubit Z stabilizer. Similarly, in the bottom right face, two X triangle operators are outlined with a red border, and their product is a six-qubit X stabilizer.

qubits (there are L^2 vertices and $2L^2$ edges of the square lattice) and $2(L^2 - 1)$ independent stabilizer generators forming a stabilizer code with $L^2 + 2$ logical qubits, L^2 of which are gauge qubits, with the remaining two logical qubits encoding quantum information. It can be verified that all triangle operators commute with the stabilizers and are therefore logical operators for the stabilizer code (since they are not stabilizers). The logical $\bar{Z}$ and $\bar{X}$ operators for each gauge qubit are chosen to be the northwest and northeast triangle operators of each face, respectively. The remaining two pairs of logical operators are the same as for the toric code, each acting nontrivially only on data qubits lying on a (horizontal or vertical) homologically nontrivial loop of the torus. In Ref. [34], it was shown that the minimum distance of the subsystem toric code is L, and therefore, the code has parameters $[[3L^2, 2, L]]$.

In Ref. [34], a planar subsystem surface code was also introduced (with two-qubit stabilizers on the boundary), which has code parameters $[[3L^2 - 2L, 1, L]]$, and in Ref. [35], a planar rotated subsystem code was introduced (with three-qubit stabilizers on the boundary) which has parameters $[[\frac{3}{2}L^2 - L + \frac{1}{2}, 1, L]]$. These parameters compare to the parameters $[[2L^2, 2, L]]$, $[[L^2 + (L-1)^2, 1, L]]$, and $[[L^2, 1, L]]$ for the toric, planar, and rotated surface codes, respectively [1,36].

By mapping the threshold to the phase transition in the random-bond Ising model on the honeycomb lattice, the subsystem toric code has been found to have a threshold of around 7% for maximum likelihood decoding, the independent Z/X noise model, and perfect syndrome measurements [34]. Under the same noise model, the threshold using a minimum-weight perfect-matching decoder is 6.5% [37]. Syndrome extraction can be done by measuring only

the three-qubit triangle operators, and it has a threshold under a circuit-level depolarizing noise model of around 0.6% [34], which is below that of the standard surface code, which has a threshold approaching 1% for a similar circuit-level depolarizing noise [38–40].

IV. FINITE-RATE LDPC SUBSYSTEM CODES

While the subsystem toric code can be derived from the hexagonal tessellation of a Euclidean surface, we now show how we can also obtain subsystem codes derived from more general tessellations, including of hyperbolic surfaces. A regular tessellation of a surface can be denoted by its Schläfli symbol $\{r, s\}$, which indicates that each face in the tessellation is an r-gon and s faces meet at each vertex. Regular tessellations of hyperbolic surfaces satisfy $1/r + 1/s < 1/2$. Hyperbolic codes, which are subspace codes derived from hyperbolic tessellations, have a finite encoding rate k/n and distance scaling as $O(\log n)$ [5,6], and it has been shown that they can require a smaller qubit overhead compared to the toric code and surface code for a target logical error rate under a phenomenological error model [26]. However, the stabilizer weight of hyperbolic codes is larger than for the toric code, making syndrome extraction more challenging, and a key benefit of the subsystem hyperbolic code construction we now give is that syndrome extraction can be done with only weight-3 check operators. Hyperbolic codes are a promising candidate for experimental realization in systems that allow variable length connectivity between qubits [41], such as modular architectures [42,43], and reduced check weight simplifies stabilizer readout, as well as reduces cross talk [9].

A subsystem hyperbolic code can be obtained by merging multiple vertices of each face of a hyperbolic tessellation. For example, a subsystem hyperbolic code can be constructed from a $\{12, 3\}$ hyperbolic tessellation by merging four vertices of each 12-gon face, as shown for a single 12-gon in Fig. 4. In general, breaking an m-clover-shaped face introduces m local loop operators that do not mutually commute. They will be interpreted as logical operators of the gauge qubits. Note that for a clover with m leaves, we introduce $m - 1$ linearly independent local loop operators as the product of all m local loops is the original face, which is a stabilizer.

As for the subsystem toric and subsystem surface codes, there exists an alternative representation with qubits placed on vertices and where a triangle operator will be placed in each corner of each face of the hyperbolic tessellation. For example, we can construct an $\{8, 4\}$ subsystem hyperbolic code by placing triangle operators in the corner of each face of an $\{8, 4\}$ tessellation, as shown in Fig. 5. Each Z stabilizer is the product of all Z triangle operators within a face of the $\{8, 4\}$ tessellation, and similar for X stabilizers and X triangle operators. Note that this code obtained from placing triangle operators in an $\{8, 4\}$ lattice can
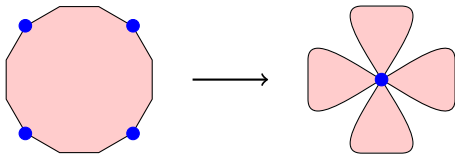
FIG. 4. Merging vertices of a 12-gon. As this procedure effectively removes three independent X checks, we introduce three gauge qubits.
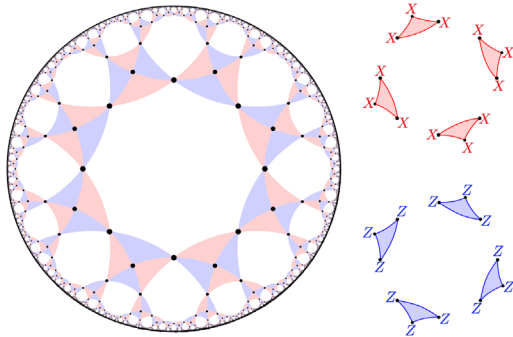

FIG. 5. The $\{8, 4\}$ subsystem hyperbolic code. A qubit (each represented by a black-filled circle) is placed in the center of each edge and on each vertex of an $\{8, 4\}$ tessellation of a closed hyperbolic surface. A three-qubit triangle operator is placed in each corner of each face. Each X stabilizer is the product of the four X triangle operators within a face (top right). Similarly, each Z stabilizer is the product of the four Z triangle operators within a face (bottom right).

equivalently be constructed by merging faces in a $\{12, 3\}$ tessellation.

We adopt the former approach (qubits on vertices) and construct the subsystem hyperbolic codes directly by requiring that, as for the subsystem toric code, any pair of triangle operators that belong to the same face of the tessellation and overlap on a single qubit must be of opposite Pauli types. Similarly, any two triangle operators belonging to the same vertex of the original tessellation and overlapping on two qubits must be of the opposite Pauli type. In other words, adjacent triangle operators related by a single rotation about a face or a vertex must be of opposite Pauli types, and we say that a tessellation that allows such an assignment of triangle operators is *colorable*. For a tessellation to be colorable, each face must have an even number of sides, and an even number of faces must meet at each vertex (so for regular $\{r, s\}$ tessellations, both r and s must be even). Furthermore, to ensure that our stabilizers commute, we further require that four faces meet at each vertex of the tessellation. In the Appendix C, we show that a regular tessellation of a *closed* surface is colorable if a particular function f (which we define) extends to a homomorphism from the symmetry group of the tessellation to the cyclic group $\mathbb{Z}_2$.

A. Properties of subsystem hyperbolic codes

We now consider some more properties of subsystem hyperbolic codes, each derived from a $\{2c, 4\}$ tessellation with edges E, vertices V, and faces F. Since we place a qubit on each vertex and in the center of each edge of this tessellation, our subsystem hyperbolic code will have $|E| + |V|$ data qubits. Each vertex in the tessellation has degree 4, and so $2|V| = |E|$. Furthermore, we also place n_a ancilla qubits within each triangle operator. While we can always use $n_a = 1$ ancillas per triangle operator by using schedules with some idle qubit locations (if necessary), we parallelize many of our schedules, which in some cases requires $n_a = 2$. Each vertex is adjacent to four triangle operators and each triangle operator is adjacent to a single vertex. Therefore, in total there are $n = \frac{3}{2}|E|$ data qubits and $2n_a|E|$ ancilla qubits in our subsystem hyperbolic codes. For the subsystem toric code, where $|E| = 2L^2$, there are $3L^2$ data qubits and $4n_aL^2$ ancilla qubits.

The number of faces in the $\{r, s\}$ tessellation satisfies $r|F| = 2|E|$. Since the product of all X-type (or Z-type) stabilizers is the identity, and since there are no other relations that the stabilizers satisfy, the number of independent stabilizers is $4|E|/r-2$. Therefore, the total number of logical qubits (including gauge qubits) is $(3/2 - 4/r)|E| + 2$.

Aside from the triangle operators introduced within each face, the number of remaining bare logical operators (those in $C(\mathcal{G})\backslash\mathcal{G}$ is determined from the topology of the tessellation from which it is derived. Therefore, excluding gauge qubits, the number of logical qubits k that a subsystem hyperbolic code derived from an $\{r, 4\}$ tessellation encodes is given by [6]

$$k = \frac{|E|}{2} - \frac{2|E|}{r} + 2. \tag{3}$$

This leaves $(1 - 2/r)|E|$ gauge qubits, or $r/2 - 1$ gauge qubits per face. The triangle operators act nontrivially on these gauge qubits. The encoding rate of the subsystem hyperbolic code is therefore

$$\frac{k}{n} = \frac{1}{3} - \frac{4}{3r} + \frac{2}{n}. \tag{4}$$

There are $4n_a/3$ ancilla qubits per data qubit, leading to $(4n_a/3 + 1)n$ qubits in total. Note that this expression does not depend on r: The number of ancilla qubits is proportional to the number of data qubits, and the constant of proportionality is the same regardless of which $\{2c, 4\}$ tessellation we use.

In Appendix F, we show that the distance d of a subsystem hyperbolic or semihyperbolic code is bounded by $d_X/2 \le d \le d_X$, where d_X is the X distance of the *subspace* hyperbolic or semihyperbolic code derived from the same tessellation. The X distance of the subspace code is always less than or equal to its Z distance for the codes

we consider, and so the distance of the subsystem code is at least half, and at most the same as the distance of the subspace code. We analyze the distances of the codes we construct in Appendix F and find codes with distances that span this full range.

B. Condition for consistent scheduling

In order to determine the syndrome used for decoding, we require a stabilizer measurement schedule, which is the sequence of gates applied to data and ancilla qubits in order to measure the eigenvalues of the stabilizers. We now show that any valid stabilizer measurement schedule defined within a single face of the subsystem toric code and chosen to be periodic in space (i.e., identical for every vertex or face) can be generalized for a subset of $\{4c, 4\}$ subsystem hyperbolic codes, for $c \in \mathbb{Z}^+$. The measurement schedule used by Bravyi et $al.$ [34] is an example of such a periodic schedule.

We first assign an element of the cyclic group $\mathbb{Z}/4\mathbb{Z}$ to each of the four types of triangle operators within a face and call such an assignment a *labeling*. We choose to label the northwest, northeast, southeast, and southwest triangle operators with the elements 0, 1, 2, and 3 of $\mathbb{Z}/4\mathbb{Z}$, respectively [see Fig. 6(a)]. Note that, for a translationally invariant schedule, each triangle operator with a given *label* in the subsystem toric code is assigned an identical schedule. Triangle operators with different labels have different measurement schedules. In order to apply this measurement schedule to the subsystem hyperbolic code, we label every triangle operator as one of these four types in such a way that the schedule always looks locally the same as for the subsystem toric code to ensure that it remains correct. More precisely, for each triangle operator with a given label in the subsystem hyperbolic code, its neighborhood of triangle operators it shares qubits with (and their labels) must be the same as for a triangle operator with the same label in the subsystem toric code. We call a labelling
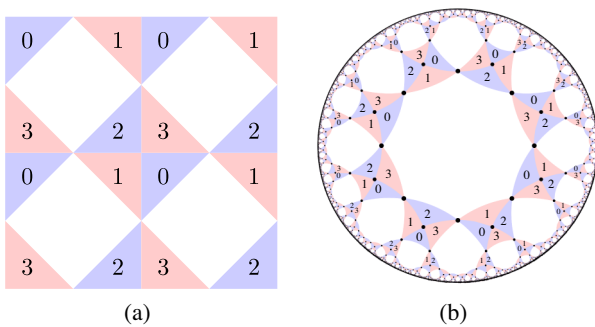

FIG. 6. (a) An $L = 2$ subsystem surface code. The four types of triangle operators are labeled as 0, 1, 2, and 3. (b) Labeling of the four types of triangle operators on an $\{8, 4\}$ tessellation of the hyperbolic plane. The neighborhood of each triangle operator (the types and relative locations of triangle operators it overlaps with) is the same as in the toric code.

that has this property a *valid labeling*, and a *schedulable* code is one that admits a valid labeling. In Appendix D, we show that a regular tessellation of a closed hyperbolic surface admits a valid labeling if a particular function h (which we define) extends to a homomorphism from the proper symmetry group of the tessellation to the cyclic group $\mathbb{Z}/4\mathbb{Z}$. We show that a subset of $\{4c, 4\}$ regular tessellations of closed hyperbolic surfaces satisfies this property. An example of a valid scheduling of the $\{8, 4\}$ tessellation of the hyperbolic plane is shown in Fig. 6(b).

C. Subsystem semihyperbolic codes

The $\{8, 4\}$ subsystem hyperbolic code has stabilizers of weight 12, which is double that of the subsystem toric code. Despite the check operators still being weight 3, we find that the large stabilizer weight results in a lower threshold of 0.31(1)% compared to 0.666(1)% for the subsystem toric code. The intuition behind this is the following: If a stabilizer has higher weight, it provides less information about the location of an error and requires more gates to be used when measured, making it harder to measure precisely.

To address this issue, we can construct subsystem codes derived from *semihyperbolic* tilings introduced in Ref. [26]. The idea is to fine-grain the tessellation leading to lower-weight stabilizers. A semihyperbolic tiling is derived from a $\{4, q\}$ regular tessellation of a closed hyperbolic manifold for $q > 4, q \in \mathbb{Z}^+$. Each (square) face of the $\{4, q\}$ tessellation is tiled with an $l \times l$ square lattice. By doing so, the curvature of the surface is weakened. The subspace quantum code derived from the semihyperbolic tessellation (a semihyperbolic code) has larger distance and reduced check weight compared to a code derived from the original $\{4, q\}$ tessellation. This comes at the cost of requiring l^2 times more qubits and, since the number of logical operators is unchanged, the encoding rate is reduced by a factor of l^2. An important advantage of semihyperbolic codes is that, by increasing l, we obtain a family of codes with distance scaling like $\sqrt{n}$ (as for the toric code), while expecting to retain a reduced qubit overhead relative to the toric code [26]. The same advantages apply for the subsystem semihyperbolic codes we construct in this work.

Recall that the tessellations that we derive subsystem hyperbolic codes from must have vertices of degree 4, and each face must have $4c$ sides (where $c \in \mathbb{Z}^+$). On the other hand, a $\{4, q\}$ semihyperbolic tiling instead has faces with four sides, while vertices have degree 4 or q. We can therefore derive a subsystem code from the *dual lattice* of $\{4, 4c\}$ semihyperbolic tessellation. In Appendix D, we show that if an $\{8, 4\}$ tessellation is schedulable, then so is the semihyperbolic tessellation derived from it. Therefore, each schedulable closed $\{8, 4\}$ tessellation defines a family of subsystem semihyperbolic codes (each code in the family having a different lattice parameter l), and where each code in the family is schedulable.

We say that an l, $\{4c, 4\}$ subsystem semihyperbolic code is the code derived by placing a triangle operator in each corner of each face of the dual lattice of a semihyperbolic lattice, where that semihyperbolic lattice was constructed by tessellating each face of the $\{4, 4c\}$ tessellation with an $l \times l$ square lattice. The subsystem semihyperbolic codes we construct and analyze in this work are $l = 2$, $\{8, 4\}$ subsystem semihyperbolic codes. The irregular tessellations these codes are derived from therefore contain both square and octagonal faces, with four faces meeting at each vertex.

V. IMPROVED ERROR CORRECTION BY GAUGE FIXING

We now introduce some general techniques that improve the quantum-error-correcting performance of a wide class of subsystem codes. We alter the stabilizer measurement procedure *in software*, in such a way that the individual gauge operator measurements themselves yield useful information. This use of individual gauge operator measurements is in contrast to existing methods for decoding subsystem codes in the literature, where individual gauge operator measurements themselves are never treated as syndrome bits, and only their products (the stabilizers) are used for decoding.

While we analyze these techniques numerically using the subsystem code constructions given in Secs. III and IV, the key ideas can be applied to the vast majority of subsystem codes considered in the literature, for which stabilizer eigenvalues can be inferred by measuring gauge operators. In fact, these techniques address one of the main drawbacks of subsystem codes, which is that they typically have lower thresholds. Low thresholds arise partly because stabilizer eigenvalues are determined by combining the outcomes of many gauge operator measurements, each of which may be faulty, making their measurement less reliable. Additionally, these high-weight stabilizers provide less information about which qubit has suffered an error, further reducing the threshold. The most dramatic example of this effect is the Bacon-Shor code which, although it has weight-2 check operators, has no threshold, as the stabilizer operators grow with system size. The techniques we introduce can also be used when applying logical operations with *subspace* codes, as we explain in Sec. VII B, since lattice surgery and code deformation for surface codes can be interpreted as gauge fixing of a larger subsystem code [33].

We call the general method schedule-induced gauge fixing, since we alter the schedule of the stabilizer measurement circuits in such a way that gauge fixing can be used to significantly improve the error-correcting performance when decoding. We refer to it simply as gauge fixing when the meaning is clear from context.

Schedule-induced gauge fixing can be applied to a large class of subsystem codes, for which there are stabilizers s that are the product of gauge operators

$s = g_0, g_1, \ldots, g_{m-1}, g_i \in \mathcal{G} \backslash \mathcal{S}$. We call these gauge operators *gauge factors* $\mathcal{G}^s$ of s,

$$\mathcal{G}^s := \{g_0, \ldots, g_{m-1} | g_i \in \mathcal{G} \backslash \mathcal{S}, s = g_0, g_1, \ldots, g_{m-1}\}, \quad (5)$$

and stabilizers which admit such a decomposition are referred to as *composite stabilizers*. In general, there can be more than one such decomposition for a given stabilizer, though we are typically most interested in the minimum-weight decomposition, where the average weight of gauge factors $g_i \in \mathcal{G} \backslash \mathcal{S}$ is minimized. For the codes we construct in this work, there is a unique minimum-weight decomposition for each stabilizer, though in general there can be more than one [44]. For CSS subsystem stabilizer codes, the gauge factors of each stabilizer mutually commute and can be measured in any relative order. For more general subsystem codes, the order of measurements of gauge factors $g_0, g_1, \ldots, g_{m-1}$ of each stabilizer $s \in \mathcal{S}$ must be chosen such that each gauge factor measurement g_i commutes with the product $g_0 g_1 \ldots g_{i-1}$ of gauge factor measurements before it. In Ref. [29], this condition was shown to be both necessary and sufficient to guarantee that the stabilizer can indeed be recovered from the product of individual measurements. Schedule-induced gauge fixing will typically be most useful for subsystem codes which have at least one composite stabilizer, and for which the weight of each composite stabilizer is greater than the weight of each of its gauge factors. In the case of the subsystem codes studied in this work, the gauge factors of each Z stabilizer associated with a face are the Z triangle operators belonging to that face (and similar for X stabilizers and X triangle operators).

When decoding subsystem codes with existing methods, the syndrome used consists of eigenvalues of stabilizers. In other words, where a stabilizer is composite, measured by taking the product of the measurements of its m gauge factors $g_i \in \mathcal{G}^s$, it is the product that is used, not the result of each gauge factor measurement individually. Therefore, for each stabilizer, we measure m bits of information and using only a single bit (their parity) for decoding. For the most simple stabilizer measurement schedules typically used, the parity is indeed all the useful information that can be used for decoding. This is because $\mathcal{G}$ is not Abelian and, by definition, each gauge factor $g_i \in \mathcal{G}^s$ must anticommute with at least one other gauge operator $h \in \mathcal{G}$. Once h is measured, either h or $-h$ becomes a stabilizer, and a subsequent measurement of g_i will result in either 1 or -1 at random with $P(1) = P(-1) = 0.5$. Consider a schedule W of measurements of check operators $K_0, K_1, \ldots, K_{N-1}$, chronological order from left to right, where each check operator K_i is either a gauge factor or a stabilizer that is not composite, and where each K_i is measured once. If this measurement schedule W is simply repeated periodically, then every consecutive pair of measurements of any check operator K_i will be separated by one measurement of every

other check operator. As a result, if the check operators in W generate $\mathcal{S}$ as required, every measurement of a gauge factor will give a random outcome and will not be useful for decoding, since its eigenvalue will not have been preserved between consecutive measurements. In fact, the eigenvalue of any product of check operators that is not in $\mathcal{S}$ will also not be preserved between consecutive measurements, following similar reasoning.

However, we can instead choose a measurement schedule W, again repeated periodically, where some gauge factors g_i are measured multiple times within W, with no anticommuting check operators measured between consecutive measurements of g_i within W. In this case, only the first measurement of g_i in W will have a uniform random outcome, whereas the remaining measurements of g_i within W will have fixed outcomes (if no error occurs), as the quantum state will (temporarily) be an eigenstate of g_i—we can think of g_i (or $-g_i$) as a temporary stabilizer.

A. Gauge fixing for CSS codes

We now restrict our attention to CSS subsystem codes, for which the gauge group $\mathcal{G}$ can be decomposed into a set of operators each in $\{I, X\}^n$, which we denote $\mathcal{G}_X$, and a set of operators each in $\{I, Z\}^n$, which we denote $\mathcal{G}_Z$, with $\mathcal{G} = \mathcal{G}_X \cup \mathcal{G}_Z$. The stabilizer group can similarly be decomposed into either X-type or Z-type Pauli operators. For CSS subsystem codes, the most common measurement schedule consists of alternating between measuring all X-type- and all Z-type-check operator measurements in a repeating sequence. In other words, the sequence of measurements for measuring the X or Z stabilizers is of the form $(ZX)^r$, where $2r$ is the number of rounds of stabilizer measurements, and the chronological order is from left to right. We call such a sequence of measurements a *homogeneous* schedule, since all stabilizers of the same Pauli type are given identical measurement schedules. Equivalently, for the subsystem codes we construct, a homogeneous schedule assigns the same schedule to each face of the lattice from which it is derived. We sometimes denote a schedule just by its longest repeating subsequence if the number of repetitions is not relevant [i.e., denoting the above schedule by ZX rather than $(ZX)^r$].

For the ZX schedule, each X gauge operator measurement comes directly after the measurement of a Z gauge operator that it anticommutes with (and vice versa), and so the outcome of each individual gauge operator measurement is random. However, by repeating X- or Z-check operator measurements, we can temporarily *fix* some gauge operators as stabilizers. As an example, consider a homogeneous schedule of the form $(Z^2X^2)^r$. The first in each pair of X gauge operator measurements will give a random outcome, whereas the second is simply a repetition of the first and, provided no error has occurred, will give the same outcome as the first measurement. The same is true for the first and second Z gauge operator measurement outcomes.

B. Gauge-fixing matching graph: Vertex splitting and merging

We now show how this additional gauge operator information can be used when decoding a CSS subsystem code using a method based on minimum-weight perfect matching, which introduces the additional requirement that the code must have no more than two stabilizers of a given Pauli type acting nontrivially on each qubit. Subsystem codes which satisfy these properties include the subsystem surface code [34], the Bacon-Shor code [18,22], and some 2D compass codes [24], including heavy-hexagon codes [9].

As an example, let us first consider the 2D matching graphs of the subsystem toric code, assuming perfect stabilizer measurements. Each vertex in the X-type (or Z-type) matching graph corresponds to an X (or Z) stabilizer, and each edge corresponds to a qubit (and therefore, a possible error). For the stabilizer group of the subsystem toric code with no gauge operators fixed, both the X-type and Z-type matching graphs are triangular lattices, as shown in Fig. 7 (left) for the X-type matching graph. This triangular lattice-matching graph has a minimum-weight perfect-matching (MWPM) threshold of 6.5% with perfect measurements [37]. However, once we measure all the X-type gauge operators, they become gauge fixed as stabilizers (up to signs that can be corrected in software), and the stabilizer group we obtain is that of the hexagonal toric code [37]. The new associated X-type matching graph instead has an improved MWPM threshold with perfect measurements of 15.6% [37], exceeding that of the toric code on a square lattice of 10.3% [1]. If we measure all the Z-type gauge operators, we instead obtain the dual of the hexagonal toric code, and now the Z-type matching graph is a hexagonal lattice.

When using the standard ZX schedule for the subsystem toric code, the stabilizer group is indeed constantly switching (up to signs) between the hexagonal toric code and its dual, both Abelian subgroups of the gauge group $\mathcal{G}$. However, each gauge operator is only ever fixed immediately after it is measured, and is randomized by the time the same gauge operator is next measured, since an anticommuting gauge operator of the opposite Pauli type is measured in between these consecutive measurements of the same gauge operator. However, by making more than
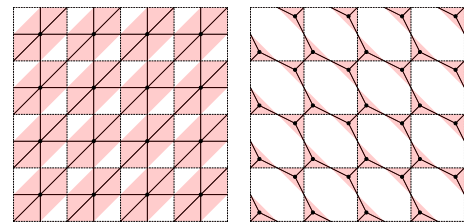

FIG. 7. Matching graphs (X type) for the subsystem toric code with no triangle operators fixed as stabilizers (left) and all triangle operators fixed as stabilizers (right).

one consecutive measurement of gauge operators of a given Pauli type, we now show that we can gauge fix into the hexagonal toric code (and its dual) for longer durations, thereby making more valuable use of the individual gauge operator outcomes themselves.

Since measurements themselves can be faulty, we must instead use a 3D matching graph when decoding the subsystem toric code. Each vertex in the matching graph corresponds to a stabilizer measurement, and each edge (u, v) corresponds to a single fault that can occur, creating a defect (-1 syndrome) at vertices u and v. In order to handle measurement errors, each stabilizer measurement is repeated $T \geq L$ times [1], and a syndrome for a stabilizer at time step t takes the value -1 if its value differs from its measurement in time step $t-1$. Measurement errors correspond to timelike edges, and memory (data qubit) errors correspond to spacelike edges. There are also single circuit faults that can induce *diagonal* edges, which have vertices that differ in both space and time. We can label each vertex in the matching graph with a coordinate (s, t), where t is the time step and $s = g_0, \ldots, g_{m-1}$ denotes the stabilizer using its gauge factors $g_i \in \mathcal{G}^s$. We depict the 3D matching graph for the subsystem toric code in Fig. 8.

For the ZX schedule used in the previous literature, gauge operators are never fixed and stabilizers are always the product of gauge operators, whereas for many of the schedules we use, we can fix a subset of the gauge operator measurements and obtain (temporarily) stabilizers consisting of single gauge operators. In our matching graph, we can fix a measurement of a gauge operator g as a stabilizer if no gauge operator h which anticommutes with g is measured since the last measurement of g. The matching graphs for the schedules $(ZX)^6$ and $(Z^3X^3)^2$ are shown in Fig. 9. For the $(ZX)^6$ schedule, gauge operators can never be fixed as stabilizers in the matching graph, whereas for the $(Z^3X^3)^2$ schedule, two-thirds of the gauge operator measurements can be fixed as stabilizer measurements.
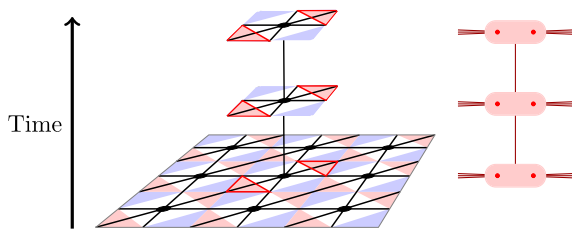

FIG. 8.   The 3D X-type matching graph for the subsystem toric code. Left: we show three time steps of the 3D matching graph for a single X stabilizer (highlighted in red), with black lines denoting edges. Right: we also use this more simple, abstract notation to depict the same 3D matching graph in our work, restricted to a single face of the lattice. Here, each pale red rounded rectangle corresponds to an X stabilizer in one of three consecutive time steps in the matching graph. Red dots denote X triangle operator measurements (two of which within a face form a stabilizer), and red lines denote edges in the 3D matching graph.
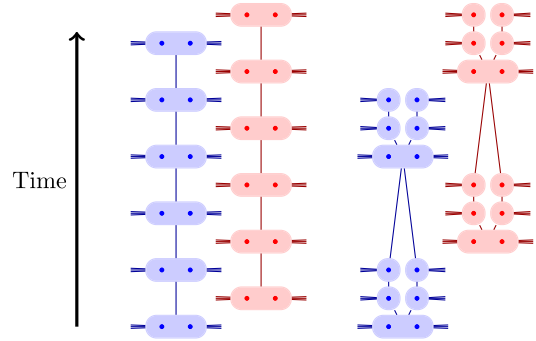

FIG. 9.   Matching graph for a single face of the subsystem toric code using a homogeneous $(ZX)^6$ schedule (left) and a homogeneous $(Z^3X^3)^2$ schedule (right). The vertical axis corresponds to time, with the direction of time being from bottom to top. Small blue- and red-filled circles correspond to Z and X gauge operator measurements, respectively, with each vertical column of small filled circles corresponding to a single gauge operator. Large light-blue- and light-red-filled rounded rectangles (or rounded squares) correspond to stabilizers being the product of the gauge operators they enclose. Diagonal edges (between stabilizers that differ in space and time) are omitted for clarity. Blue and red lines correspond to edges in the Z and X matching graphs, respectively.

Since each gauge operator has weight 3, by fixing some gauge operators as stabilizers, we can reduce the weight of some stabilizer measurements from 6 down to 3.

Since the stabilizers can change between consecutive time steps when using gauge fixing, we must generalize our definitions of the difference syndrome and vertical edges in the matching graph. For our generalized difference syndrome, we set the syndrome of stabilizer s to be -1 in time step t if its eigenvalue differs from that of the same product of gauge operators in time step $t-1$. We draw a vertical edge in the matching graph between a stabilizer measurement s_t in time step t and measurement s_{t-1} in time step $t-1$ if s_t and s_{t-1} have at least one gauge factor in common.

As an example, we now consider the case where a stabilizer has two gauge factors, as is the case for the subsystem toric code. Suppose a stabilizer is the product $g_0 g_1$ of gauge factors g_0 and g_1 in time step $t-1$, but both g_0 and g_1 are fixed as stabilizers in time step t. We say that the stabilizer vertex is *split* into two vertices in time step t, with the matching graph locally looking like (with time propagating upward)
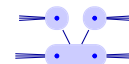
and a measurement error in time step $t-1$ on gauge factor measurement g_0, e.g., at the vertex $(g_0 g_1, t-1)$, will cause a -1 difference syndrome at vertex $(g_0 g_1, t-1)$ as well as vertex (g_0, t). Therefore, this measurement error corresponds to flipping the vertical edge $[(g_0 g_1, t-1), (g_0, t)]$.

The same argument holds for a measurement error on g_1 in time step $t-1$ corresponding to flipping the other vertical edge $[(g_0g_1, t-1),(g_1,t)]$. Similarly, we can fix g_0 and g_1 as stabilizers in time step $t-1$ but instead have the stabilizer g_0g_1 in time step t (the vertices are *merged* in time step t). This merging occurs if gauge operators are measured in between time steps $t-1$ and t that anticommute with g_0 and g_1. The matching graph locally looks like
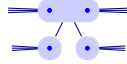
and we find that a measurement error that occurs at the vertex $(g_0, t-1)$ results in a -1 syndrome at both $(g_0, t-1)$ and (g_0g_1, t), corresponding to flipping the edge $[(g_0, t-1),(g_0g_1,t)]$. Similarly, a measurement error at vertex $(g_1, t-1)$ corresponds to flipping the edge $[(g_1, t-1),(g_0g_1,t)]$. While, in this example, we consider stabilizers which have only two gauge factors (which is the case for subsystem toric codes), the definition of the difference syndrome can be applied to stabilizers with any number m of gauge factors. For example, we have $m=4$ for the $\{8,4\}$ subsystem hyperbolic codes considered in this work, since these have four triangle operators (gauge factors) in each face.

In a stabilizer round in which all gauge operators are fixed (matching graph vertices are split), there are two distinct advantages which gauge fixing can offer. First, vertical timelike edges have a lower error probability, since the syndrome corresponding to a vertex is obtained from only a single check operator measurement, rather than taking the product of multiple measurements. Second, the degree of vertices in the matching graph is reduced. The advantage that this can offer becomes clear when we again consider the (spacelike) matching graph of the subsystem surface code when all gauge operators are fixed, compared to the matching graph when they are not fixed. We find that the hexagonal lattice matching graph when gauge operators are fixed (Fig. 7, right) has a threshold of around 4.1% under a phenomenological noise model. On the other hand, for the triangular lattice-matching graph when no gauge operators are fixed (Fig. 7, left), we find a threshold of 2.0% with a phenomenological noise model (see Fig. 30). Furthermore, the outcomes of the weight-3 checks are more reliable, since their measurement circuits are shorter. However, a potential disadvantage of gauge fixing is that by repeating X checks, more errors accumulate for the next measurement of Z checks, for which Z gauge operators cannot be fixed. We show in Sec. VI A 1 how this trade-off leads to an optimal homogeneous schedule for the threshold under a circuit-level depolarizing noise model.

C. Homogeneous stabilizer measurement circuits

In order to measure the triangle operators (and therefore stabilizers), we require a circuit to measure each triangle

operator using an ancilla qubit. We now show how these circuits can be constructed for *homogeneous schedules*, where the same schedule is applied to each face in the lattice. As we discuss in Sec. IV B, all triangle operators with the same label share the same schedule in the subsystem toric, hyperbolic, and semihyperbolic codes we use, so we need specify only four parity-check circuits, one for each label. Each triangle operator consists of three data qubits and at least one ancilla, and can be measured using three controlled-NOT (CNOT) gates, along with state preparation and measurement of an ancilla. A time step is defined as the time taken for a CNOT gate, and we assume that state preparation and measurement combined take a single time step. This assumption is similar to the assumption of non-demolition measurements in Refs. [39,45], except we assume both state preparation and measurement errors, rather than just the latter. In Ref. [34], the authors instead assume that state preparation and measurement each take a time step, and use an additional ancilla to parallelize state preparation and measurement into a single time step. The parity-check measurement circuit therefore takes four time steps.

The measurement schedules we use are shown in Fig. 10. The schedule shown on the left of Fig. 10 is for alternating measurement of the Pauli-Z and Pauli-X operators (ZX schedule), and is the same as that used in Ref. [34]. The right-hand diagram in Fig. 10 shows the schedule for measuring ZZ (blue labels) as well as the schedule for measuring XX (red labels). All three of these schedules have period 4, and so the time steps which each gate is labeled with are given modulo 4. Note that the first half of the ZZ schedule matches the Z component of the ZX schedule, and the first half of the XX schedule matches the X component of the ZX schedule. Therefore, the schedule for *any* homogeneous sequence can be implemented by concatenating these three schedules (or subsets of them). For the standard ZX schedule, we need only a single ancilla qubit for each triangle operator. For schedules which contain ZZ, we use two ancillas per Z triangle operator
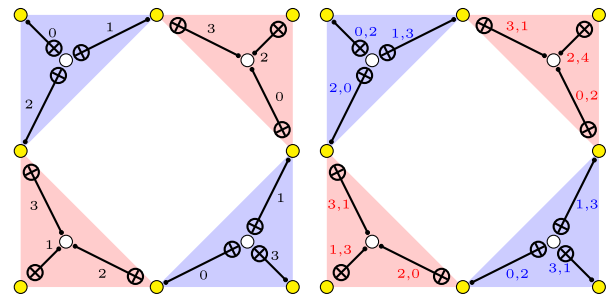

FIG. 10. Parity-check measurement schedule for the subsystem surface code using a homogeneous $(ZX)^r$ sequence (left) [34], a homogeneous Z^r sequence (right, blue text), and a homogeneous X^r sequence (right, red text). CNOT gates are labeled with the time step(s) they are applied in, which are given modulo 4, since all schedules have period 4.

to parallelize consecutive triangle operator measurements, and similarly, we use two ancillas per X triangle operator for parallelized schedules containing XX.

For the subsystem hyperbolic and semihyperbolic codes, we generalize the schedule in Fig. 10 by using the same schedule for triangles with the same label, as we explain in Secs. IV B and Appendix D. Each individual fault in the measurement circuit results in at most a single data qubit error, a property that is made possible by the weight-3 gauge operators. As a result of this bare-ancilla fault tolerance of the measurement circuits, we can correct up to the full code distance for all the codes we construct.

D. Edge weights

In order to decode the subsystem surface codes using minimum-weight perfect matching, we construct a matching graph, where each individual fault that can occur flips an edge in the matching graph [1,34]. We assign each edge a weight $w = \log[(1 - p)/p]$, where p is the total probability that any individual fault will result in the edge being flipped [1,39,46].

We first consider the matching graph obtained by measuring only X-type-check operators and fixing *all* X-type gauge operators as stabilizers. We see later that all other matching graphs for arbitrary homogeneous schedules can be obtained by merging edges and/or vertices in this matching graph. There are two types of X-type gauge operators in the subsystem surface code, as shown in Fig. 6, labeled by 1 and 3, which we refer to as T_1 and T_3, respectively. Every spacelike or diagonal edge is from a T_1 to a T_3 (or vice versa), and the neighborhood of every triangle operator with the same label is identical. All seven types of edges in the matching graph for X-type checks are shown in Fig. 11. All edges are undirected but are denoted by directed arrows in the diagram to remove any ambiguity in the definition of the diagonal edges. The purely spacelike edges are labeled 0, 1, and 2, purely timelike errors are
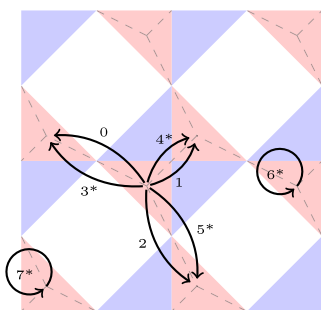
labeled 6 and 7, and diagonal edges are labeled 3, 4, and 5. Diagonal and timelike errors are drawn from time step t to time step $t + 1$, whereas spacelike edges connect vertices within a single time step. Therefore, each vertex in this matching graph has degree 8 (since each vertex is both the source and target of a timelike edge).

If an X-type-check operator is measured directly after a Z-type-check operator that anticommutes with it, then this X-type-check operator cannot be fixed, and the matching graph shown in Fig. 11 is not quite valid. However, we can use the vertex-merging procedure detailed in Sec. V B to give the matching graph the correct structure. When the X-type-check operators within a face of the lattice cannot be fixed, then the corresponding X-type matching graph vertices from that face (each vertex v_{g_i} corresponding to a gauge factor g_i) are *merged* into a single vertex v_s. The edges incident to v_s each correspond to an edge incident to a gauge factor vertex v_{g_i}. This process can result in more than one edge (a multiedge) between the same pair of vertices [such as for timelike edges in homogeneous $(ZX)^r$ schedules]. When this happens, we replace the multiedge with a single edge and assign it a flip probability equal to the probability that an odd number of edges in the multiedge would flip.

In order to calculate the probability p that each edge flips (both for edge weights and for simulations), we count the number of single faults (of each type) that can lead to each type of edge flipping. These counts are given in Table I for the X matching graph (for X-type-check operators). The operators G_1^X and G_1^Z are Pauli errors from CNOT gates in the X or Z measurement schedule, respectively,

FIG. 11. The different types of edges in the 3D matching graph of the subsystem surface code for X-type checks only, when all X-type gauge operators are fixed. Each unique edge type is labeled with a number. If an asterisk is present in the label, the edge is from time step t to $t + 1$, otherwise, the edge is purely spacelike. The whole X matching graph for a single time step is drawn with gray dashed lines.

TABLE I. Number of single faults that can cause each type of edge to flip in the 3D matching graph for X-type-check operators. Each G_1^X or G_1^Z fault is a single Pauli error arising from a CNOT gate in the measurement circuit for an X-type or Z-type gauge operator, respectively. Each G_2^Z fault is a *pair* of Pauli errors arising from a single CNOT gate in the measurement circuit for a Z-type gauge operator. P_X and M_X are state preparation and measurement errors in the X-type-check operator measurement schedule, respectively. r_Z is the number of rounds of Z-type-check operator measurements that have occurred since the last X-type-check operator measurement. For example, $r_Z = 1$ always for $(ZX)^r$ schedules, and $r_Z = 2$ always for $(ZZX)^r$ schedules. The edge types are shown in Fig. 11. Faults for the Z matching graph can be found by exchanging Z and X in the table.

Edge type	G_1^X	G_1^Z	G_2^Z	P_X	M_X
0	2	$2r_Z$	0	0	0
1	2	$2r_Z$	$2r_Z$	0	0
2	2	$2r_Z$	$2r_Z$	0	0
3	2	0	0	0	0
4	2	0	0	0	0
5	2	0	0	0	0
6	3	0	0	1	1
7	3	0	0	1	1

TABLE II. The probability of a fault occurring for each type of circuit element under the two error models considered in this work, as well as for the depolarizing error model used in Ref. [34] for reference.

Error type	G_1^X	G_2^X	G_1^Z	G_2^Z	P_X	M_X	P_Z	M_Z
Depolarizing	$\frac{4}{15}p$	$\frac{8}{15}p$	$\frac{4}{15}p$	$\frac{8}{15}p$	$\frac{2}{3}p$	$\frac{2}{3}p$	$\frac{2}{3}p$	$\frac{2}{3}p$
Independent	$\frac{1}{3}p_X$	$\frac{2}{3}p_X$	$\frac{1}{3}p_Z$	$\frac{2}{3}p_Z$	p_X	p_X	p_Z	p_Z
Ref. [34]	$\frac{1}{4}p$	$\frac{1}{2}p$	$\frac{1}{4}p$	$\frac{1}{2}p$	p	p	p	p

corresponding to either an XI, IX, or XX error acting after the gate. In the standard depolarizing model, G_1^X or G_1^Z errors occur with probability $4p/15$. See Table II for the gate error probabilities under the independent noise model we use. G_2^Z errors correspond to a pair of G_1^Z errors from the same CNOT gate in the Z measurement circuit that both cause the same edge to flip. For example, both XI and XX errors on a CNOT gate may cause the same edge to flip, and since these errors are mutually exclusive on the same gate, the chance of either of these errors occurring is exactly twice the probability that one of them occurs. The number of G_1^Z or G_2^Z errors that can cause an edge to flip depends on r_Z, the number of Z-check operator measurements that have occurred since the most recent prior X-check operator measurement. We can recover the matching graph for the standard $(ZX)^r$ schedule used in Ref. [34] by setting $r_Z = 1$ and merging all vertices within each face (up to small differences in the error model shown in Table II).

E. Noise models

We consider two different types of noise models: a circuit-level depolarizing noise model and a circuit-level independent noise model. The depolarizing noise model is widely used in the literature and is useful for comparing to previous work. Later, we consider biased noise, for which we use the independent noise model.

The circuit-level depolarizing noise model is the same as that used in Refs. [9,46] and is parametrized by a single variable p. Ancilla state preparation and measurement errors each occur with probability $2p/3$. With probability p, each CNOT gate is followed by a two-qubit Pauli error drawn uniformly from $\{I, X, Y, Z\}^{\otimes 2} \backslash I \otimes I$. A single-qubit Pauli error drawn uniformly from $\{X, Y, Z\}$ occurs with probability p after each idle single-qubit gate location. Note that many of our syndrome extraction circuits are fully parallelized and do not contain single-qubit gates or idle locations.

In our circuit-level independent noise model, Z-type errors and X-type errors are independent. For a given error probability parametrized by p_0, we choose a high-rate-error probability for Z-type errors $p_Z = p_0 \eta / (\eta + 1)$ and the low-rate-error probability $p_X = p_0 / (\eta + 1)$ for X-type errors. The bias $\eta = p_Z / p_X$ parametrizes the relative strengths of Z-type and X-type errors. The total probability of any error is

$$p_{\text{tot}} = 1 - (1 - p_X)(1 - p_Z)$$
$$= p_0 - \frac{p_0^2 \eta}{(\eta + 1)^2}. \tag{6}$$

Each with probability p_Z, a CNOT gate is followed by an error in $\{IZ, ZI, ZZ\}$, chosen uniformly at random, an X-type ancilla is prepared or measured in an orthogonal state, and a single qubit idle for one time step undergoes a Z error. Similarly, each with probability p_X, a CNOT gate is followed by an error randomly chosen from $\{IX, XI, XX\}$, a Z-type ancilla is prepared or measured in an orthogonal state, and a single qubit idle for one time step undergoes an X error. Biased noise models are common in many physical realizations of quantum computers, and bias-preserving CNOT gates can be realized using stabilized cat qubits [11]. We note that our techniques significantly improve performance even for small finite bias ($\eta \leq 10$), which may be achievable even with CNOT gates that do not fully preserve bias, as is the case in many architectures [47,48].

The probability of each different type of circuit element undergoing a fault for our two error models (as well as the error model in Ref. [34] for comparison) is given in Table II.

VI. NUMERICAL ANALYSIS

For all of the numerical results in this section, we used a local variant of the minimum-weight perfect matching (MWPM) decoder described in Appendix A and available at Ref. [49], along with the BLOSSOM V implementation of the Blossom algorithm [50,51].

A. Subsystem toric codes

1. Gauge fixing for depolarizing noise

We now show how gauge fixing can be used to improve the quantum-error-correcting performance of the subsystem toric code under a depolarizing noise model. For this unbiased noise, we use *balanced* schedules, which we define to be of the form $Z^a X^a$ for some $a \in \mathbb{Z}^+$. We find that schedules that allow gauge fixing increase the threshold from 0.666(1)% for the standard $(ZX)^r$ schedule used in Ref. [34] to 0.811(2)% for the $(Z^4 X^4)^r$ schedule, where gauge operators are fixed for three in every four rounds of measurements. In Fig. 12, we show the thresholds for the ZX, $Z^2 X^2$, and $Z^3 X^3$ schedules. We see that both the $Z^2 X^2$ and $Z^3 X^3$ schedules are higher than the standard ZX schedule, but the crossing is at a higher logical error rate. For these balanced schedules $(Z^a X^a)^r$ under depolarizing noise, we find that $a = 4$ is optimal (see Table V). Therefore, schedule-induced gauge fixing makes the threshold of the subsystem toric code under depolarizing noise much more
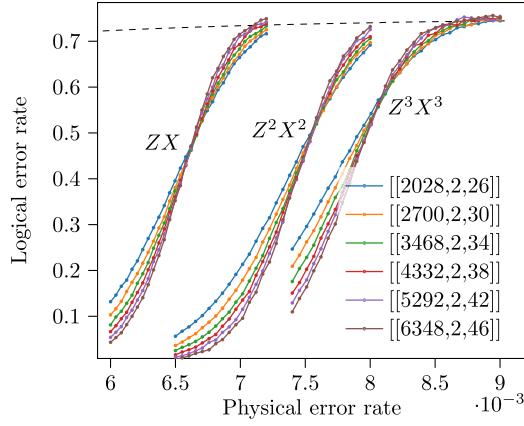
FIG. 12. Threshold plots for subsystem toric codes using a $(ZX)^{92}$ schedule (left), $(Z^2X^2)^{46}$ schedule (middle), and $(Z^3X^3)^{31}$ schedule (right) using a depolarizing noise model.

competitive with the rotated surface code, which we find has a threshold of around 0.97% under the same noise model and assumptions (state preparation and measurement each take half the time of a CNOT gate, and the logical error rate per time step is used). However, in Sec. VI A 2 we show that schedule-induced gauge fixing with the subsystem toric code can be used to outperform the rotated surface code for small finite bias $\eta > 2.3$.

By using gauge fixing (setting $a > 1$), we reduce the average stabilizer weight in the 3D matching graph, since the stabilizers introduced from gauge fixing have weight 3. The mean stabilizer weight in the 3D (X check) matching graph for a $\{2c, 4\}$ subsystem surface or hyperbolic code using a $(Z^qX^a)^r$ schedule (for any $q \geq 1$ or $r \geq 1$) is given by $3ca/[c(a-1)+1]$. So for the subsystem toric code ($c = 2$), the mean stabilizer weights for the $(ZX)^r$, $(Z^2X^2)^r$, and $(Z^3X^3)^r$ schedules are 6, 4, and 3.6, respectively. We also reduce the average degree of vertices in the matching graph. For $a = 1$, the mean vertex degree is 14, whereas for $a > 1$, the mean vertex degree is $8ca/[c(a-1)+1]$, and so the $(ZX)^r$, $(Z^2X^2)^r$, and $(Z^3X^3)^r$ schedules have mean vertex degrees of 14, 32/3, and 9.6, respectively, for the subsystem toric code. More properties of matching graphs for some homogeneous schedules with the subsystem toric code are given in Table III.

TABLE III. The mean $|\bar{s}|$, maximum $|s|_{\max}$, and minimum $|s|_{\min}$ stabilizer weight and mean $\bar{d}$, maximum Δ, and minimum δ degree of the X-check 3D matching graphs for various homogeneous schedules with the subsystem toric code.

Schedule	$|\bar{s}|$	$|s|_{\max}$	$|s|_{\min}$	$\bar{d}$	Δ	δ
Z^qX	6	6	6	14	14	14
Z^qX^2	4	6	3	10.67	16	8
Z^qX^3	3.6	6	3	9.6	16	8
Z^qX^5	3.33	6	3	8.89	16	8
Z^qX^{10}	3.16	6	3	8.42	16	8

While we expect that reducing the average stabilizer weight and vertex degree in the matching graph should improve the threshold, increasing a in balanced Z^aX^a schedules also alters the edge-fault probabilities. In time steps where gauge operators are fixed, $r_Z = 0$ in Table I, reducing the edge weights for some edges of type 0, 1, and 2. However, in the time steps where gauge operators are not fixed, $r_Z = a$, and so increasing a also increases the edge-fault probability for these edges of type 0, 1, and 2. Therefore, increasing a increases the proportion of time steps where a spacelike slice of the matching graph is a degree-3 hexagonal lattice with small edge-fault probabilities, but also *increases* the edge-fault probabilities for the remaining time steps where the matching graph is *not* fixed, and is instead a degree-6 triangular lattice. There is therefore a trade-off between increasing the edge weights and decreasing the stabilizer weights and vertex degrees, and the $a = 4$ schedule is the optimal compromise for schedules of the form $(Z^aX^a)^r$ for a circuit-level depolarizing noise model.

Since changing the schedule alters both the matching graph via gauge fixing, as well as the edge-fault probabilities, we can better understand how these two factors contribute to performance by studying them separately. In Fig. 13, we plot the threshold as a function of a for balanced schedules Z^aX^a both with and without using gauge fixing. The thresholds that do not use gauge fixing are decoded by always merging gauge factors of a stabilizer into a single vertex in the matching graph, even in time steps where they could be split (gauge factors fixed) using the techniques we introduce. We see that for schedules that do not use gauge fixing, there is almost no improvement for $a > 1$, with performance degrading for $a > 4$. This demonstrates that almost all the improvement in threshold for depolarizing noise is due to the use of gauge fixing, rather
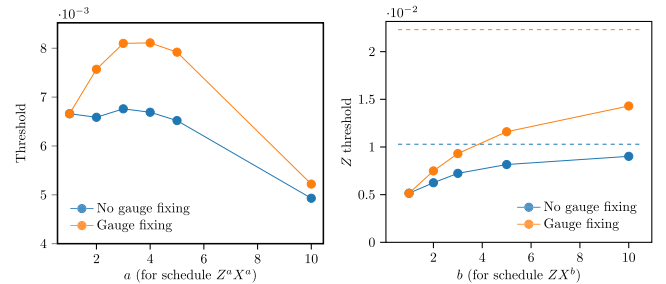

FIG. 13. Left: circuit-level depolarising threshold as a function a for schedules of the form Z^aX^a, with and without gauge fixing. Right: Z thresholds as a function of b for schedules of the form ZX^b, both with (orange) and without (blue) gauge fixing, using a circuit-level independent noise model. The orange and blue dashed lines are the threshold achievable under infinite bias (using an X schedule) with and without gauge fixing, respectively. Error bars are smaller than the marker size and are omitted for clarity.

than the change in the noise model induced by the different schedule alone.

2. Tailoring the 3D matching graph to biased noise using gauge fixing

By using *unbalanced* schedules, where X-check operators are measured more frequently than Z-check operators (or vice versa), we can use gauge fixing to improve performance under biased noise models. Since we correct X errors and Z errors independently, we can define the Z threshold p_Z^{th} and X threshold p_X^{th} as the threshold for only Z-type or only X-type errors, respectively. In Fig. 14, we plot the Z threshold for the unbalanced ZX, ZX^2, ZX^{10}, and X schedules, under the independent circuit-level noise model. Increasing the ratio of X checks to Z checks significantly increases the Z threshold from 0.52% for the ZX schedule up to 2.22% for the X schedule, which sets an upper bound.

By measuring X checks more frequently, we also reduce the noise on data qubits caused by the CNOT gates used to measure Z checks. To determine how much of the improvement in threshold comes from this reduced noise in the measurement schedule compared to the use of gauge fixing in the matching graph, we determine the thresholds both with and without using gauge fixing in Fig. 13. We see that even without using gauge fixing, increasing the ratio of X checks to Z checks increases the Z threshold, as expected. However, gauge fixing significantly boosts the Z threshold further, and even a ZX^5 schedule using gauge fixing outperforms the best achievable Z threshold without gauge fixing (using the X schedule).

However, by increasing the ratio of X to Z checks, we also reduce the X threshold of the code, which we must take into account when determining the total threshold under
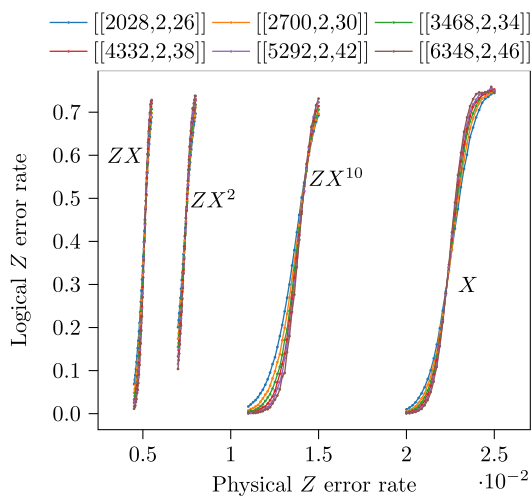
biased noise models. We now ask what the threshold is under the biased independent circuit-level noise model described in Sec. V E, with bias parameter η. Specifically, for a given η, we wish to find the total physical error rate $p_{\text{total}}^{\text{th}}$ below which the total logical error probability $p_{\text{total}}^{\log}$ of both logical $\bar{X}$ or $\bar{Z}$ errors vanishes as the distance L of the code increases to infinity. A sufficient and necessary condition for a total error probability p_{total}' to be below the accuracy threshold for a decoder that decodes Z and X errors independently is that the probability of a Z-type error p_Z' be below p_Z^{th} and the probability of an X-type error p_X' be below p_X^{th}.

The total error probability $p_{\text{total}}^{Z_{\text{th}}}$ when $p_Z = p_Z^{\text{th}}$ is

$$p_{\text{total}}^{Z_{\text{th}}} = p_Z^{\text{th}} + p_Z^{\text{th}}(1 - p_Z^{\text{th}})\frac{1}{\eta}, \tag{7}$$

and the total error probability $p_{\text{total}}^{X_{\text{th}}}$ when $p_X = p_X^{\text{th}}$ is

$$p_{\text{total}}^{X_{\text{th}}} = p_X^{\text{th}} + p_X^{\text{th}}(1 - p_X^{\text{th}})\eta. \tag{8}$$

The total threshold $p_{\text{total}}^{\text{th}}$ is therefore given by

$$p_{\text{total}}^{\text{th}} = \min(p_{\text{total}}^{Z_{\text{th}}}, p_{\text{total}}^{X_{\text{th}}}). \tag{9}$$

In Fig. 15, we plot $p_{\text{total}}^{\text{th}}$ as a function of the bias parameter η for the subsystem toric code and for a few different choices of homogeneous schedule. For the ZX schedule used in Ref. [34] and the Z^3X^3 schedule with gauge fixing, the optimal bias is $\eta = p_Z/p_X = 1$. This is as expected, since the X threshold is identical to the Z threshold for these symmetric schedules. From Eqs. (7)–(9) we see that at $\eta = 0$ and $\eta = \infty$ the total threshold is simply the X threshold and Z threshold, respectively.

FIG. 14. Z thresholds for unbalanced schedules of the form ZX^b, as well as an X schedule, which gives an upper bound on the Z threshold achievable using unbalanced schedules.
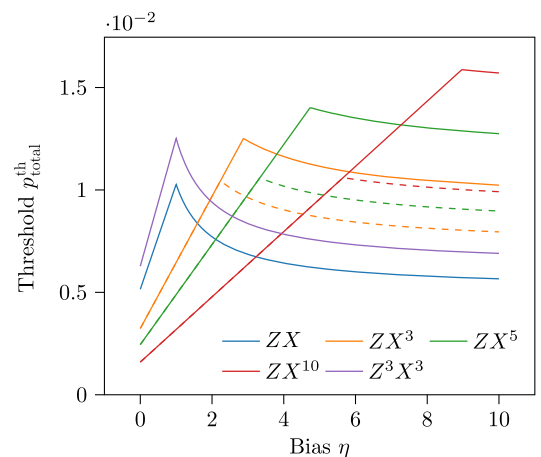

FIG. 15. Threshold $p_{\text{total}}^{\text{th}}$ [see Eq. (9)] as a function of the bias for different homogeneous schedules and under a circuit-level independent noise model. Dashed lines use the same schedule as the corresponding solid line of the same color, except gauge fixing is not used, for the purpose of comparison.
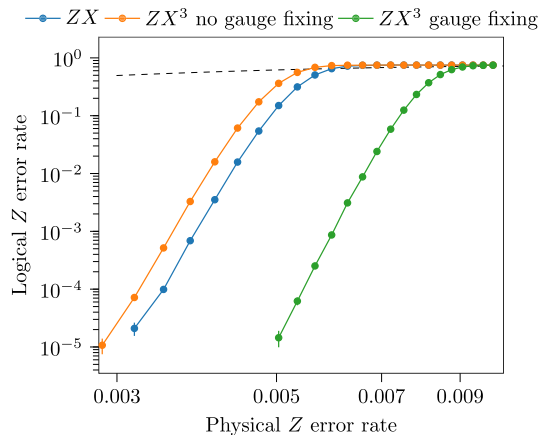
FIG. 16. Logical $\bar{Z}$ error rate of the $[[2028, 2, 26]]$ subsystem toric code using a $(ZX)^{36}$ schedule, as well as a $(ZX^3)^{12}$ schedule (using only a single ancilla by introducing idle time steps) with and without using gauge fixing in the matching graph. All schedules use 144 time steps, and the independent circuit-level noise model is used. The dashed black line is the probability that either of two physical qubits will suffer a Z error during 144 time steps without using error correction.

For each of the schedules for which $p_{\text{total}}^{\text{th}}$ is plotted in Fig. 15, there are two regimes: to the left and to the right of the peak. To the left of the peak, the threshold is limited by the X threshold, and is therefore given by Eq. (8), which is linear in η. To the right of the peak, the threshold is limited by the Z threshold, and is therefore given by Eq. (7), which is linear in $1/\eta$. The optimal η for a given schedule can be found by setting $p_{\text{total}}^{Z_{\text{th}}} = p_{\text{total}}^{X_{\text{th}}}$.

Even for small finite bias, using unbalanced schedules and gauge fixing significantly improves the total threshold compared to the traditional ZX schedule, with a 2.8× increase in threshold at $\eta = 9$. With infinite bias, the threshold rises to 2.22% which is 4.3× higher than the threshold of 0.52% using the standard ZX schedule. Each dashed line in Fig. 15 uses the same schedule as the corresponding solid line of the same color but without using gauge fixing to decode. For high bias, we see that approximately half of the improvement over the ZX schedule can be attributed to the effect the new schedule has on the noise model, with the remainder attributed to the extra information used by gauge fixing when decoding.

For the rotated surface code using the same schedule as in Ref. [35], we find a threshold under circuit-level independent noise of 0.741(2)%. Therefore, the subsystem toric code (with a ZX^3 schedule and using gauge fixing) outperforms the rotated surface code for biases $\eta > 2.3$.

Note that, for all the thresholds we report so far, we use fully parallelized schedules. Whereas the ZX schedule is fully parallelized with only $n_a = 1$ ancilla qubits per triangle operator, the unbalanced ZX^b schedules require two ancilla qubits per X-check operator ($n_a = 1.5$), and the balanced $Z^a X^a$ schedules require two ancilla qubits per

X-check operator *and* per Z-check operator ($n_a = 2$). Since there are $4n_a/3$ ancilla qubits per data qubit, this leads to a larger qubit overhead when using gauge fixing with parallelized schedules. We can choose not to parallelize the schedules, and instead simply omit gates in the ZX schedule to construct our other schedules (e.g., an unparallelized ZX^2 schedule can be constructed by omitting every other Z measurement in the ZX schedule). These schedules incur no qubit overhead but instead introduce idle errors. The threshold with infinite bias using an unparallelized X schedule is 1.25%, compared to 2.22% using a parallelized X schedule, both an improvement over the 0.52% threshold using the ZX schedule. Near the threshold, using additional ancillas is clearly worthwhile, whereas far below threshold it may be beneficial to use an unparallelized schedule, using the additional qubits to instead construct a code with a larger distance.

To analyze the performance below threshold, we compare a ZX schedule to an unparallelized ZX^3 schedule ($n_a = 1$) using the $L = 26$ subsystem toric code, both with and without using gauge fixing to decode. When using gauge fixing, the logical Z error rate is reduced by around 4 orders of magnitude compared to the ZX schedule (see Fig. 16). Without using gauge fixing, the logical error rate with the unparallelized ZX^3 schedule is slightly worse than with the ZX schedule, since idle qubit errors are worse than qubit errors in the standard depolarizing noise model [40].

B. Performance of the finite-rate LDPC subsystem codes

We simulate the performance of $l = 2$ $\{8, 4\}$ subsystem semihyperbolic codes under the circuit-level depolarizing noise model. We are interested in finding the threshold value below which the logical error rate per logical qubit tends to zero as the code distance tends to infinity. Since the number of logical qubits k increases with distance for this family of finite-rate codes, we fix the number of logical qubits by using multiple independent copies of the smaller codes. In Fig. 17, we plot the probability that at least one of 338 logical qubits suffers a Z failure as a function of the depolarizing error rate p. The $[[8064,338,10]]$ code has the lowest logical error rate per logical qubit for physical error rates below 0.42%, from which we conclude that the threshold is at least 0.42%. We are not able to obtain an upper bound on the threshold, since all codes have an error rate (per 338 logical qubits) of 1 for physical error rates above 0.42%, within the precision provided by our numerical experiments.

We now analyze the performance of the $[[8064,338,10]]$ $l = 2$ $\{8, 4\}$ subsystem semihyperbolic code, which has the best ratio $n/(kd^2) = 0.24$ of the codes we construct. In Fig. 18, we compare its performance with that of the $L = 4$, 6, 8, 9, and 10 subsystem toric codes. We use 169 independent copies of the subsystem toric codes in order to keep the number of logical qubits (338) constant, and the
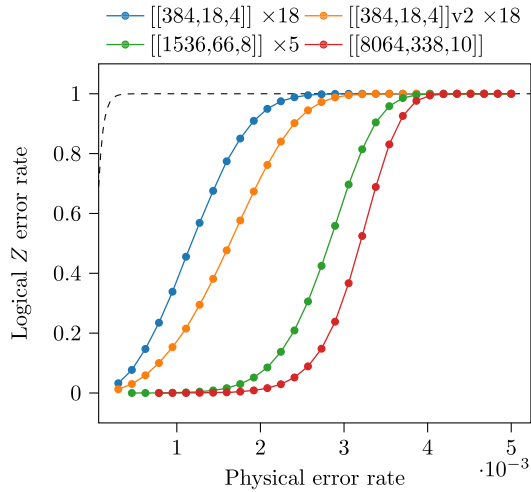
FIG. 17. Performance of the extremal $l = 2$ $\{8, 4\}$ subsystem semihyperbolic codes under a circuit-level depolarizing noise model. Here, we fix the number of logical qubits to at least 338 for all codes by using multiple copies of the smaller codes. A homogeneous $(ZX)^{20}$ schedule is used for all codes, and the y axis is the probability that at least one logical Z error occurs. The dashed black line is the probability of a Z error occurring on at least one of 338 physical qubits without error correction under the same error model for the same duration (80 time steps). For each code that encodes $k < 338$ logical qubits, we use $m = \lfloor k/338 \rfloor$ copies and plot the failure rate as $p^*_{\log} = 1 - (1 - p_{\log})^m$.

total number of physical qubits used (including ancillas) is given in the legend. We find that the [[8064,338,10]] subsystem semihyperbolic code (which uses 18 816 physical qubits) outperforms the $L = 4$ subsystem toric code (which uses 18 928 physical qubits to encode 338 logical qubits) by around 3 orders of magnitude at $p = 0.15\%$. At a physical error rate of 0.2%, the performance of the [[8064,338,10]] subsystem semihyperbolic code is similar to the $L = 9$ subsystem toric code, which uses 95 823 physical qubits to achieve the same logical error rate. This demonstrates that the [[8064,338,10]] subsystem semihyperbolic code requires $5.1\times$ fewer resources to achieve the same level of protection that the subsystem toric code would provide at a physical error rate of 0.2%.

We also compare the [[8064,338,10]] subsystem semihyperbolic code with the rotated surface code, which is the leading candidate for realizing fault-tolerant quantum computation, and has the optimal ratio $n/d^2 = 1$ for surface codes [36]. This comparison is shown in Fig. 19, where we again keep the number of logical qubits fixed by using 338 independent copies of the rotated surface codes. At a circuit-level depolarizing error rate of 0.15%, the subsystem semihyperbolic code using 18 816 physical qubits has a similar performance to $L = 11$ rotated surface codes using 81 458 physical qubits, a $4.3\times$ reduction in qubit overhead. We also compare the performance of the [[8064,338,10]] subsystem semihyperbolic code with a distance-6 rotated surface code, which has a slightly lower
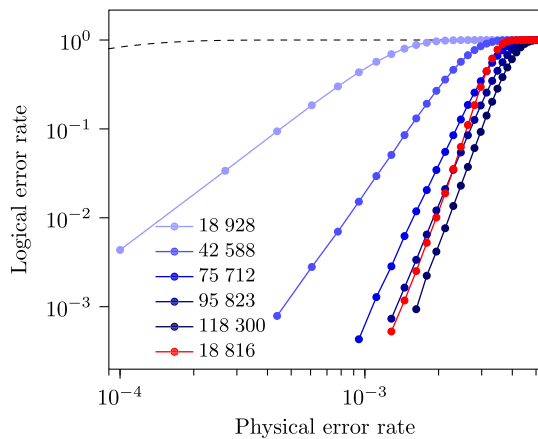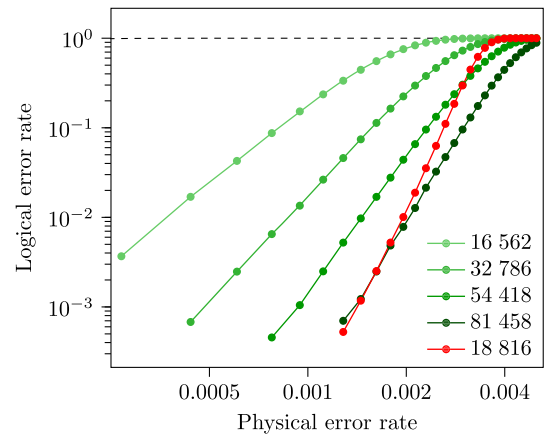

FIG. 18. Comparison of the [[8064,338,10]] $l = 2$ $\{8, 4\}$ subsystem semihyperbolic code (red), which has 8064 data qubits and 10 752 ancillas, with $L = 4, 6, 8, 9,$ and 10 subsystem toric codes (shades of blue), using a $(ZX)^{20}$ schedule (no gauge fixing) and a circuit-level depolarizing noise model. We fix the number of logical qubits by plotting the probability that at least one of 169 independent copies of the subsystem toric codes suffers a logical Z failure [i.e., we plot $1 - (1 - p_{\log})^{169}$ for the subsystem toric codes where $p_{\log}$ is the probability that a single copy of the code suffers a logical Z error]. The total number of physical qubits (including ancillas) is given in the legend. The black dashed line is the probability that at least one of 338 physical qubits would suffer a Z failure without error correction over the same duration.

FIG. 19. Comparison of the [[8064,338,10]] $l = 2$ $\{8, 4\}$ subsystem semihyperbolic code (red), with $L = 5, 7, 9,$ and 11 rotated surface codes (shades of green) using a $(ZX)^{20}$ schedule (no gauge fixing) for the subsystem semihyperbolic code and a $(ZX)^{16}$ schedule for the rotated surface codes (both schedules require 80 time steps). We use a circuit-level depolarizing noise model. We fix the number of logical qubits by plotting the probability that at least one of 338 independent copies of the rotated surface code suffers a logical Z failure. The legend gives the total number of qubits (ancilla and data qubits) used. The black dashed line is the probability that at least one of 338 physical qubits would suffer a Z failure without error correction over the same duration.
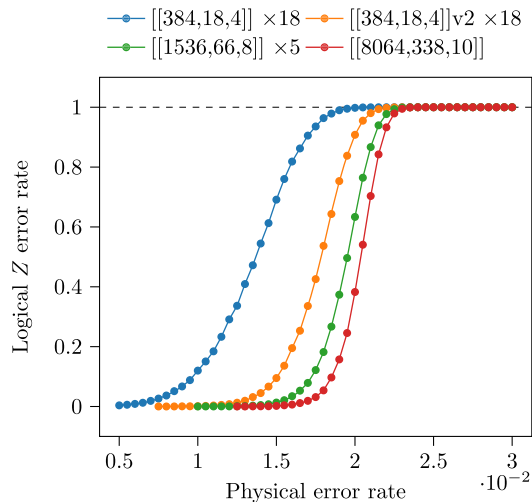
FIG. 20. Performance of the extremal subsystem $\{8,4\}$ $l = 2$ semihyperbolic codes under a circuit-level independent noise model and using an X schedule.

encoding rate (including ancillas) and find that the subsystem semihyperbolic code has a lower logical error rate below 0.43%.

To the best of our knowledge, the rotated surface code is the best-performing code in the literature in terms of qubit overhead in the regime of around 0.15% to 0.2% circuit-level depolarizing noise, which is roughly the same physical error rate assumed for practical implementations of fault-tolerant quantum computing [25,52]. Since our subsystem semihyperbolic codes have a qubit overhead that is 4.3× smaller than the rotated surface code at $p = 0.15\%$, as demonstrated in Fig. 19, we therefore believe that they outperform all known quantum-error-correcting codes in terms of qubit overhead in this regime.

Furthermore, we can use schedule-induced gauge fixing for the subsystem hyperbolic and semihyperbolic codes just as we did for the subsystem toric code. In Fig. 20, we plot the threshold of the $l = 2$, $\{8,4\}$ subsystem semihyperbolic codes under the independent circuit-level noise model using an X schedule, and find a threshold of at least 2.4%, exceeding that of the subsystem toric code (2.22%). This threshold sets an upper bound on the thresholds that can be achieved using gauge fixing under biased noise models, and we expect that large gains can still be found even for small finite bias, as we find for the subsystem toric codes.

VII. BROADER APPLICATIONS OF OUR TECHNIQUES

A. Inhomogeneous schedules

We so far consider only homogeneous schedules; however, sometimes it may be advantageous to use schedules that are *inhomogeneous*, where check operators in different faces of the lattice are given different schedules.

As an example, consider two different unparallelized ZX^4 schedules, which we call L_0 and L_1, obtained by omitting three quarters of the Z-check operator measurements in the ZX schedule, and such that L_1 is identical to L_0 other than a lag of four check operator measurements. A section of eight rounds of X-check operator measurements for these schedules looks like

$(ZX)^8$	Z	X	Z	X	Z	X	Z	X	Z	X	Z	X	Z	X	Z	X
L_0		X		X		X	Z	X		X		X		X	Z	X
L_1		X	Z	X		X		X		X	Z	X		X		X

where each column corresponds to a measurement round of either X-type- or Z-type-check operators. We can assign either the L_0 or L_1 schedule to each face of the planar subsystem surface code independently, since each schedule is a subset of the ZX schedule, for which we have a consistent measurement circuit for every face. Let G_0^X be the set of X triangle operators in faces assigned the L_0 schedule, and let G_1^X be the set of X triangle operators in faces assigned the L_1 schedule. Note that in each round of X-check operator measurements, either G_0^X, G_1^X, or $G_0^X \cup G_1^X$ may be fixed.

Can an inhomogeneous schedule be used to increase the Z distance of a subsystem code? For the planar subsystem surface code, the only Z logical is a Pauli Z operator applied to each qubit in a column of the lattice, corresponding to a path in the matching graph joining the north and south boundaries. Consider the inhomogeneous schedule where we alternate between using the L_0 and L_1 schedule in each row of the lattice: We assign the schedule $L_{(i \bmod 2)}$ to faces in the ith row of the lattice. For a planar subsystem surface code with an odd distance, in each round of X-check operator measurements at least half of the gauge operators can be fixed: We can fix gauge operators in all rows, then in even rows, then all rows again, then odd rows, and so on in a cycle. In Fig. 21, we plot spacelike slices (single time steps) of the 3D matching graph for when all rows, odd rows, and even rows of gauge operators are fixed. Within each of these slices of the 3D matching graph, the shortest path between the north and south boundary is *larger* than the Z distance of the subsystem surface code itself. We expect that the shortest path between the north and south boundaries of the overall 3D lattice is also larger,
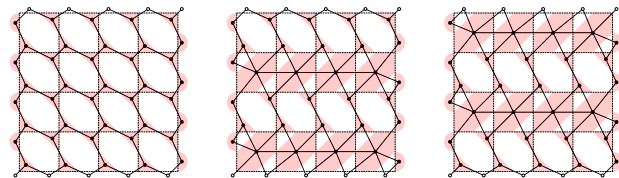

FIG. 21. Matching graphs (X type) for the $L = 5$ subsystem surface code with triangle operators fixed in all rows (left), odd rows (middle), and even rows (right). Filled and hollow circles correspond to stabilizers and boundary nodes, respectively.

leading to an increased Z distance of $d_Z = \lfloor 3(L-1)/2 \rfloor + 1$ but do not prove this here. The X distance cannot increase in this schedule, since none of the Z gauge operators can be fixed.

Note that homogeneous schedules cannot increase the Z or X distance of the code, since there are always time steps where all X gauge operators are measured simultaneously, as well as time steps where all Z gauge operators are measured simultaneously. Measuring all X gauge operators removes all Z gauge operators from the stabilizer group, leaving time steps where none of the Z gauge operators can be fixed (and therefore, not increasing the X distance), and similarly, there are also time steps where no X gauge operators can be fixed.

B. Lattice surgery and code deformation

It was shown in Ref. [33] that the techniques of lattice surgery [53] and code deformation [54] can be interpreted as switching between different gauge fixes of a subsystem code. We can use this perspective to apply some of the techniques in this work to lattice surgery and code deformation. As an example, consider performing lattice surgery on two rotated surface-code patches. During the merging step of lattice surgery, the weight-$2X$ stabilizers on the opposing boundaries of the two patches are merged into weight-4 square stabilizers. These weight-4 stabilizers can be interpreted as stabilizers of a subsystem code, with the weight-2 checks that they are merged from being gauge operators of the subsystem code. This procedure is shown for distance-3 codes in the left side of Fig. 22, for which a
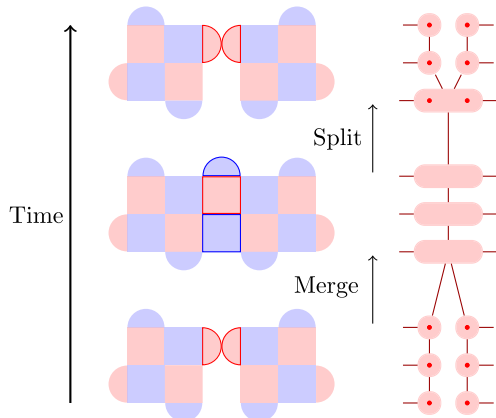

FIG. 22.   A slice of the matching graph for lattice surgery, which can be interpreted as switching between different gauge fixes of a subsystem code. Left: the three stages of lattice surgery are shown for a distance-3 rotated surface code. Red (blue) squares and semicircles denote X (Z) stabilizers, with data qubits at their corners. Right: a slice of the matching graph for the X stabilizers at the boundaries of the two codes where the merge takes place (denoted with red borders in the left diagram). Stabilizer measurements are repeated three times for each stage of lattice surgery, with the generalized difference syndrome used to connect the stabilizer with its gauge factors.

single pair of weight-$2X$ checks (with red borders) is merged into a single square stabilizer. Since each pair of these weight-2 boundary X checks is a pair of gauge factors of the corresponding weight-4 stabilizer, we can use the merging and splitting technique given in Sec. V B to construct the matching graph and decode them. This is shown on the right side of Fig. 22, where three repetitions are used for each of the three stages of lattice surgery. With this technique, each of the consecutive stages of lattice surgery can be connected using the generalized difference syndrome, leading to a single matching graph that can be used for error correction with the overlapping recovery method of Ref. [1], and with information from the weight-2 boundary X checks used directly where possible. The same ideas can also be readily applied to code deformation, which can also be viewed as gauge fixing of a subsystem code [33] and involves merging surface-code patches in a similar manner [54].

C. Subspace codes from gauge fixing

Another use of gauge fixing is to derive families of subspace codes from subsystem surface, toric, and hyperbolic codes by choosing different Abelian subgroups of the gauge group $\mathcal{G}$ to be the stabilizer group $\mathcal{S}$, permanently fixing some gauge operators as stabilizers. For example, by fixing all the X-type triangle operators in the subsystem toric code as stabilizers, we obtain the hexagonal toric code, and by fixing X-type triangle operators in the $\{8, 4\}$ subsystem hyperbolic code as stabilizers, we obtain the $\{12, 3\}$ hyperbolic code.

By fixing different subsets of the triangle operators in the subsystem toric code, we can interpolate between the hexagonal toric code and its dual. To achieve this, we define hexagonal surface-density codes inspired by the surface-density and Shor-density codes of Ref. [24]. To construct a (subspace) hexagonal surface-density code with parameter q_f from a subsystem toric code, we fix the X-type gauge operators in each face with probability q_f, else we fix the Z-type gauge operators. When $q_f = 1$, we obtain the hexagonal surface code, and at $q_f = 0$ we construct its dual, but setting $0 < q_f < 1$ allows us to interpolate between these two extremes. With $q_f = 0.5$, there are both weight-6 and weight-$3X$-type and Z-type stabilizers, and both X-type and Z-type stabilizers have average weight 4. The same idea can be directly applied to subsystem hyperbolic codes: Applied to the $\{8, 4\}$ subsystem hyperbolic code, we can interpolate between the $\{12, 3\}$ hyperbolic code and its dual, for example.

For the subsystem hyperbolic codes, we can choose to fix only a subset of the triangle operators within each face. Consider the code obtained by fixing a single Z triangle operator (chosen at random) within each face of the $\{6, 4\}$ subsystem hyperbolic code, as well as the single X triangle operator that commutes with it (an example of this for a
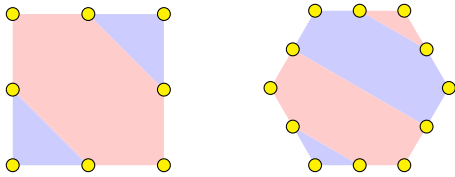
FIG. 23. Gauge fixings of a square (left) and hexagonal (right) face of a subsystem toric and $\{6, 4\}$ subsystem hyperbolic code, respectively. Yellow-filled circles are data qubits, and X and Z stabilizers are denoted by red- and blue-filled polygons, respectively.

single face is shown in Fig. 23). For both the X and Z stabilizers, half have weight 6, and the other half have weight 3. This hyperbolic code derived from an *irregular* lattice has average stabilizer weight 4.5 for both X and Z stabilizers, an improvement on the weight-5 stabilizers in the $\{5, 5\}$ hyperbolic code, which has the smallest stabilizer weight of hyperbolic codes derived from self-dual regular lattices.

We can also use our choice of Abelian subgroup of the gauge group to tailor codes to spatially inhomogeneous noise models, where the noise is biased toward Z-type errors in some regions of the lattice and biased toward X-type errors in other regions. We can fix X-type gauge operators in regions where there is a Z bias, locally reducing the vertex degree and stabilizer weight in the X-type matching graph, and likewise, we can fix Z-type gauge operators where there is X bias. This method of tailoring a code to spatially inhomogeneous noise models has been demonstrated in Ref. [24] using gauge fixes of the Bacon-Shor code, and the same ideas can be readily applied here to gauge fixes of subsystem surface, toric, and hyperbolic codes.

VIII. CONCLUSION

In this work, we introduce new techniques and constructions for quantum-error correction that improve upon the widely studied surface code in several ways. While the surface code requires four-qubit measurements and encodes a single logical qubit, we introduce families of quantum-error-correcting codes that use only three-qubit measurements and encode a number of logical qubits k proportional to the number of physical qubits n. Furthermore, we introduce a technique, which we call schedule-induced gauge fixing, that improves the performance of a wide class of codes, especially under biased noise models.

Schedule-induced gauge fixing changes the order in which check operators are measured in subsystem codes. While the check operators of subsystem codes do not all mutually commute, we find that grouping blocks of mutually commuting check operators together allows us to obtain more useful information without increasing the total number of measurements. By making consecutive

measurements of the same gauge operators they can be treated temporarily as stabilizers, and we introduce a method for decoding based on MWPM that takes advantage of this additional information. When applied to the subsystem surface code with three-qubit check operators, we can switch repeatedly between the hexagonal surface code and its dual, both of which are Abelian subgroups of the gauge group of the code. We find that the threshold under circuit-level depolarizing noise can be increased from 0.67% to 0.81% by making four consecutive measurements of each gauge operator in the measurement schedule. The improvement is even more significant under biased noise models. With an independent Z-biased circuit-level noise model, X-check operators can be repeated (and fixed) more frequently, leading to an even higher threshold under small finite bias, up to 2.22% under infinite bias. Below threshold, gauge fixing reduces the logical error rate by several orders of magnitude for biased noise models.

Schedule-induced gauge fixing can be applied *in software*, with no changes to the underlying hardware interactions necessary. This allows both the code and the decoder to be tailored to the noise model even if it cannot be fully characterized prior to device fabrication. Furthermore, the decoding method changes only the structure of the matching graph, with no additional overhead in runtime, and other decoders such as UNION-FIND [46,55], which has almost-linear runtime, can be directly substituted for MWPM in our procedure.

The same techniques can also be directly applied to a broad class of subsystem codes beyond the subsystem surface code, including the Bacon-Shor code [22], the heavy-hexagon code [9], and some compass codes [24], and future work could investigate the performance improvements achievable using schedule-induced gauge fixing with these codes. It would also be interesting to generalize the decoding method to other subsystem codes where syndrome defects do not come in pairs, such as the gauge color code [31], among others [29,56].

A drawback of subsystem codes is that they typically have a smaller encoding rate k/n compared to their subspace counterparts. To address this issue, we generalize the subsystem surface code to surfaces with negative curvature, constructing families of quantum LDPC subsystem codes with a finite encoding rate and only three-qubit check operators. We call these codes subsystem hyperbolic and subsystem semihyperbolic codes and show how the symmetry group of the tessellation can be used to construct check operator measurement circuits which require only four time steps to implement. Thanks to the weight-3 check operators, these measurement circuits allow us to correct up to the full code distance fault tolerantly.

By simulating the performance of subsystem semihyperbolic codes under circuit-level depolarizing noise, we find that they can require 4.3× fewer physical qubits than the rotated surface code and 5.1× fewer physical qubits

than the subsystem toric code to achieve the same physical error rate at around 0.15% to 0.2%. To the best of our knowledge, they therefore outperform all known quantum-error-correcting codes in terms of qubit overhead in this practical regime of circuit-level depolarizing noise. Furthermore, these subsystem semihyperbolic codes belong to a family of codes that achieve distance scaling as $\sqrt{n}$, and that we expect to maintain a reduced qubit overhead relative to the surface code even at higher distances. These codes are *locally* Euclidean, which is encouraging for the prospect of physical implementations in modular architectures [41–43].

We also find a threshold of 0.42% for the subsystem semihyperbolic codes under a circuit-level depolarizing noise. All of the techniques for schedule-induced gauge fixing that applied to the subsystem toric code can also be applied to subsystem semihyperbolic codes, and we find a threshold of 2.4% under infinite bias, exceeding that of the subsystem toric code.

Our work focuses on reducing the qubit overhead of quantum-error correction; however, reducing the time overhead of implementing logical gates is also an important problem. In Ref. [26], it was shown how lattice surgery and Dehn twists can be used to implement logical gates in hyperbolic codes. While these techniques should generalize straightforwardly to the subsystem hyperbolic codes we introduce, in the future it would be interesting to compare the time overhead of these methods with those used for surface codes, as well as to investigate alternative methods for implementing fault-tolerant logical operations.

A key advantage of all the subsystem codes we construct and use in this work is that they all use check operators of weight 3, compared to the weight-4 stabilizers of the surface code. Besides enabling bare-ancilla fault tolerance and efficient measurement schedules, weight-3 gauge operators can be helpful for handling leakage errors [35], and direct three-qubit parity-check measurements have been proposed in Ref. [57]. Since the average degree of the interaction graph is lower than the surface code, we also expect these codes to suffer from fewer frequency collisions and less cross talk than the surface code in superconducting qubit architectures [9]. On the other hand, if high-weight stabilizer measurements are available in hardware, then it may be possible to reduce the qubit overhead of our subsystem codes even further (likely at the cost of a lower threshold) by using a single ancilla qubit per stabilizer rather than per gauge operator and measuring along the gauge operators to retain bare-ancilla fault tolerance [19].

While there has been significant progress in the development of quantum LDPC codes with improved parameters $[[n, k, d]]$ relative to the surface code, our work provides the first evidence that these improvements can be retained even under circuit-level depolarizing noise. We demonstrate the advantages that can be had from codesigning an error-correcting code along with its parity-check measurement

schedule. In particular, we show that subsystem codes offer great promise in reducing the weight of check operators in quantum LDPC codes, as well as enabling improved performance under biased noise models through the use of schedule-induced gauge fixing. Furthermore, our results show that symmetries present in the construction of quantum LDPC codes can also be crucial for optimizing parity-check measurement schedules. We hope that our work will inspire the construction of new families of quantum LDPC codes designed using similar principles, further reducing the overhead of fault-tolerant quantum computation.

APPENDIX A: LOCAL MATCHING DECODER

In order to reduce the complexity of the MWPM decoder, we use an approximate version which we call *local matching*. In our local matching decoder, we check only if each defect can be matched to one of the k closest defects in the matching graph (rather than considering all other defects in the matching graph).

Given a matching graph G (containing a vertex for each stabilizer measurement or boundary and a weighted edge for each single error) and a syndrome vector $\mathbf{z}$ (where $\mathbf{z}[i] = 1$ if stabilizer i is measured to be -1, otherwise $\mathbf{z}[i] = 0$), the first step of a standard implementation of MWPM is to construct the *defect graph* V, which contains a vertex for each defect v (where by definition $\mathbf{z}[v] = 1$) in G and an edge for each possible pair of defects weighted with the distance between them in G. Edmond's Blossom algorithm is then used to find a minimum-weight perfect matching in V [50], and the product of Pauli operators, each corresponding to the shortest path between matched pairs of defects, is returned as a correction.

In our local matching algorithm, we include fewer edges in the defect graph V than used in standard MWPM. For each defect v in G, rather than finding the distance to every other defect using Dijkstra's algorithm, we instead find the distance to the m nearest defects using our *local Dijkstra* algorithm and include m edges to these defects in V, each weighted by their distance from v in G. The pseudocode for the local Dijkstra algorithm for finding the m nearest
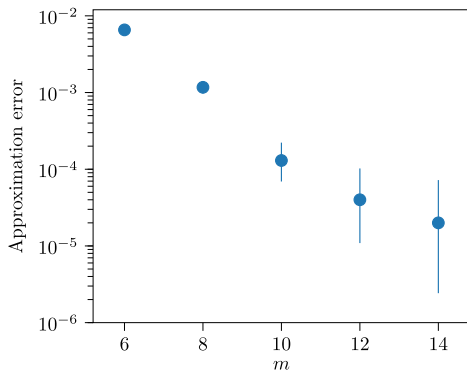
FIG. 24.   The approximation error of the local matching decoder defined as the fraction of trials for which the weight of the minimum-weight perfect matching found by our local matching decoder differs from an exact minimum-weight perfect matching. We use 10^5 trials for each m, using an independent noise model with $p = 6\%$ (and noiseless syndromes) for an $L = 30$ toric code. Error bars are 95% Clopper-Pearson binomial proportion confidence intervals. For $m = 16$, 18, and 20, we also run 10^5 trials and find that the local matching is equivalent to exact matching for all trials.

neighbors of a defect $s \in G$ is given in Algorithm 1, which outputs a list l of the m nearest defects, their distances d to s, and a predecessor list p. We use Kolmogorov's BLOSSOM V implementation of MWPM in C++ to find the MWPM in V [51].

Note that our local matching algorithm is similar to the strategy used in Ref. [45], where defects are initially matched only with other defects that are within a given radius r (determined by their coordinates in the 3D surface-code matching graph). While this strategy is effective for codes derived from tessellations of Euclidean surfaces, it does not generalize well to the codes we derive from

Algorithm   1.   Local Dijkstra algorithm.

Function *LocalDijkstra* (*G*, **z**, *m*, *s*)
> For each $u \in G$, $d[u] = \infty$, $p[u] = u$;
> $d[s] = 0$;
> Initialize priority queue Q
> $Q.insert(s)$
> Initialize empty list of found defects l
> **while** Q *is not empty and* $length(l) < m$ **do**
>> $u = Q.ExtractMin()$;
>> **if** **z**$[u] = 1$ **then**
>>> $l.Insert(u)$;
>>
>> **for** *each vertex* v *adjacent to* u *in* G **do**
>>> **if** $weight(u, v) + d[u] < d[v]$ **then**
>>>> $d[v] = weight(u, v) + d[u]$;
>>>> $p[v] = u$
>>>> **if** $d[v]$ *previously equal to* ∞ **then**
>>>>> $Q.Insert(v)$
>>>>
>>>> **else**
>>>>> $Q.DecreaseKey(v)$

---

tessellations of hyperbolic surfaces, where using coordinates as a proxy for finding the closest defects is less straightforward. Our approach uses only the structure of the matching graph itself, not the coordinates of defects, and therefore readily generalizes to the codes we derive from hyperbolic tessellations. Our adaptation of MWPM is also similar to the decoder used in Ref. [14], where rather than restricting the number of neighbors of each vertex in V as we do here, they instead impose a threshold on the maximum distance between vertices in V.

We analyze the accuracy of our local matching decoder at approximating exact MWPM. In Fig. 24 we show that, empirically, the approximation error decreases exponentially with m. For $m > 14$, we do not observe any differences in the weight of the matchings found by local matching and exact MWPM in any of the 10^5 trials run on the $L = 30$ toric code. Note that where our local matching differs from exact MWPM the solution given is still very good (low weight), so differences in the logical error rate between local matching and exact MWPM are likely far more rare than the differences in the exact minimum-weight matching solution measured here. We use $m = 20$ for all simulations in this work. Our implementation of the local matching decoder is available as a PYTHON package at Ref. [49].

APPENDIX B: TESSELLATIONS OF CLOSED SURFACES

We now give some additional background on tessellations of closed Euclidean and hyperbolic surfaces, since these tessellations are used to construct the subsystem hyperbolic and semihyperbolic codes in this work. An $\{r, s\}$ tessellation subdivides a surface into disjoint faces, where each face is an r-gon, and s faces meet at each vertex. Using Wythoff's kaleidoscopic construction, an $\{r, s\}$ tessellation can be related to a symmetry group $G_{r,s}$ of distance-preserving maps (isometries). $G_{r,s}$ is generated by reflections on the edges of one of the $2r$ right triangles induced by the symmetry axes of a face (r-gon) of the tessellation. Each triangle has internal angles $\pi/2$, π/r, and π/s, and are from now on referred to as a *fundamental triangle*. In Figs. 25(a) and 26(a), we draw a fundamental triangle of the $\{4, 4\}$ and $\{8, 4\}$ tessellations, respectfully, with sides labeled by the reflections a, b, and c which act on them, and which generate $G_{r,s}$. Note that the isometries $a^2, b^2, c^2, (ac)^2, (ab)^r$, and $(ca)^s$ are equivalent to doing nothing and, since these are the only relations satisfied by $G_{r,s}$, the group has presentation

$$G_{r,s} = \langle a, b, c | a^2 = b^2 = c^2 = (ac)^2$$
$$= (ab)^r = (bc)^s = e \rangle, \qquad \text{(B1)}$$

where e is the identity element. By fixing one fundamental triangle as a fundamental domain of $G_{r,s}$, every other
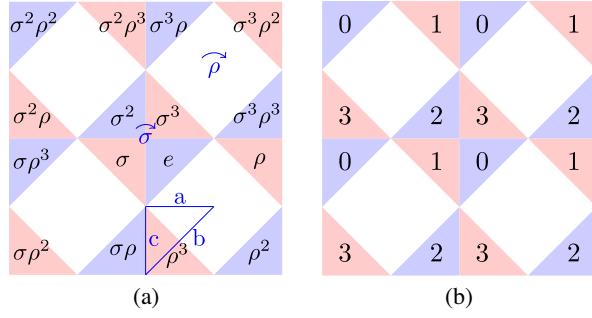
FIG. 25. An $L = 2$ subsystem surface code. (a) After associating a triangle operator with the identity element e, every triangle operator is in one-to-one correspondence with an element of the proper symmetry group $G_{4,4}^{H+}$ of the tessellation. In blue we label a fundamental triangle with sides a, b, and c, as well as the rotations $\rho = ab$ and $\sigma = bc$. (b) Each triangle operator can be labeled with an element of the cyclic group $\mathbb{Z}_4$ using the homomorphism $h(\rho) = h(\sigma) = 1$ from $G_{4,4}^{H+}$ to $\mathbb{Z}_4$.
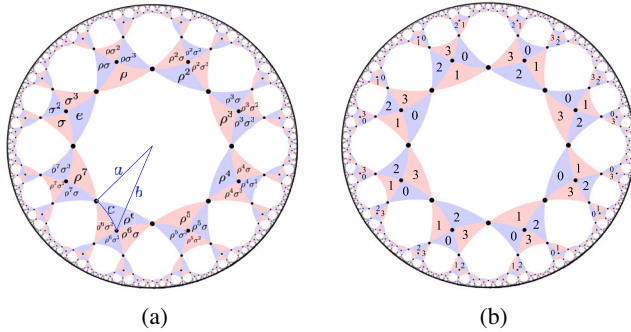

FIG. 26. The $\{8, 4\}$ subsystem hyperbolic code. (a) Each triangle operator can be uniquely identified with an element of the proper symmetry group $G_{8,4}^{H+}$ of the lattice (after identifying a triangle operator with the identity element e). We label a fundamental triangle in blue. (b) Each triangle operator can be labeled with an element of $\mathbb{Z}_4$ using the homomorphism $h(\rho) = h(\sigma) = 1$ from $G_{8,4}^{H+}$ to $\mathbb{Z}_4$. The neighborhood of each triangle operator (the labels and relative locations of triangle operators it overlaps with) is the same as in the toric code.

fundamental triangle can be labeled uniquely by an element of $G_{r,s}$.

We construct codes derived from $\{r, s\}$ tessellations of *closed* Euclidean and hyperbolic surfaces. The process of defining a closed surface is called *compactification*. A regular tessellation of a closed surface can be defined by a quotient group $G_{r,s}^H := G_{r,s}/H$, where H is a finite-index normal subgroup of $G_{r,s}$ with no fixed points (see Ref. [6] for more details). Note that the generators of H become relations in the presentation of $G_{r,s}/H$, so compactification can be interpreted as adding additional relations into the presentation of the symmetry group of the tessellation of the hyperbolic plane. An important subgroup of $G_{r,s}$ is the

proper symmetry group $G_{r,s}^+$ generated by double reflections, or *rotations*, $\rho = ab$ and $\sigma = bc$. This group has presentation

$$G_{r,s}^+ = \langle \rho, \sigma | (\rho\sigma)^2 = \rho^r = \sigma^s = e \rangle, \quad \text{(B2)}$$

where e is again the identity element. Regular tessellations of *orientable* closed surfaces can be constructed from a quotient group $G_{r,s}^{H+} := G_{r,s}^+/H$, where H is a normal subgroup of $G_{r,s}^+$.

APPENDIX C: SYMMETRY GROUPS THAT ADMIT SUBSYSTEM HYPERBOLIC CODES

In Sec. IV of the main text, we introduce subsystem hyperbolic codes, which are derived from $\{2c, 4\}$ tessellations of hyperbolic surfaces, where $c \in \mathbb{Z}^+$ and $c > 2$. In this section, we show how a subsystem hyperbolic code can be described in terms of the symmetry group of the tessellation from which it is derived. By doing so, we show what conditions must be satisfied by the compactification procedure for a $\{2c, 4\}$ tessellation of a closed hyperbolic surface to be used for constructing a subsystem hyperbolic code.

Let us first consider some properties of the subsystem toric code in group theoretic terms. These properties are later used as requirements for the subsystem hyperbolic codes we define. First, note that each triangle operator (gauge generator) of the subsystem toric code can be identified by a *pair* of fundamental triangles related by a b reflection in $G_{4,4}$. In other words, each triangle operator is identified by a *left coset* of the subgroup $\langle b \rangle$ given by $g\langle b \rangle := \{g, gb\}$ for some $g \in G_{4,4}^H$, and thus, each element $g \in G_{4,4}^H$ identifies a unique triangle operator (but not vice versa). For now, we consider only the Pauli type of each triangle operator, which can be either Z type (blue) or X type (red). We call an assignment of a Pauli type to each triangle operator a *coloring*. For the subsystem toric code, note that blue triangle operators are always mapped to red triangle operators by either an a or c reflection, and vice versa. We make this property a requirement of our subsystem hyperbolic codes and call a coloring that satisfies this property a *valid* coloring.

Since each triangle operator can be identified by the coset $g\langle b \rangle$ of an element $g \in G_{r,s}^H$, and after identifying each *color* of triangle operator with a different element of the cyclic group $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$, a coloring of the triangle operators can be achieved by defining an appropriate function $f : G_{r,s}^H \to \mathbb{Z}_2$. The constraint that either a or c reflections map a triangle operator to another of a different type, with b reflections leaving it invariant, defines the image of the generators and identity element e of $G_{r,s}^H$ by f to be

$$f(a) = f(c) = 1,$$
$$f(b) = f(e) = 0. \tag{C1}$$

Since we require that, by definition of the code, the action of a reflection a, b, or c should have the same effect on the color of a triangle operator no matter which triangle operator we apply it to, this implies that $f(g_i g_j) = f(g_i) + f(g_j) \ \forall \ g_i, g_j \in G_{r,s}^H$. This condition implies (from the definition of a homomorphism) that f must extend to a *group homomorphism* from $G_{r,s}^H$ to $\mathbb{Z}_2$. For each triangle operator to be assigned a unique color, we must also have that $f(r_i) = 0$ for each relation r_i in the presentation of $G_{r,s}^H$. This latter condition is in fact also a necessary and sufficient condition for the function f to extend to a homomorphism from $G_{r,s}^H$ to $\mathbb{Z}_2$ [58]. This constraint $f(r_i) = 0$ holds not just for the $\{4, 4\}$ tiling but also $\{r, s\}$ tilings for which r and s are even, since $(ab)^r = e$ and $(bc)^s = e$ are relations. The constraints do not hold if either r or s are odd. However, we also have the constraint $f(g_i) = 0$ on the generators g_i of the normal subgroup H defining the compactification (since these generators are relations in $G_{r,s}^H$) and, therefore, only a subset of the possible compactifications of these regular tessellations admit valid colorings.

We must also ensure that each triangle operator in a colored tessellation commutes with every stabilizer, and that all stabilizers mutually commute (since by definition $\mathcal{S}$ is Abelian and the center of $\mathcal{G}$). We will now show that this condition further restricts us to tessellations where $s = 4$ faces meet at each vertex. For regular tessellations of closed Euclidean or hyperbolic surfaces, we are already restricted to $s \geq 3$, and we already require that s be even to ensure a valid coloring. For all $s \in \{6, 8, 10, \ldots\}$ we see that each triangle operator anticommutes with the stabilizer (of the opposite Pauli type) belonging to the face related to it by a $(bc)^3$ rotation, since it overlaps with this stabilizer on only a single qubit. On the other hand, for $s = 4$, it can be directly verified that each triangle operator commutes with all stabilizers, since each triangle operator overlaps on either zero or two qubits with stabilizers of the opposite Pauli type. Since stabilizers are products of nonoverlapping triangle operators, all stabilizers must also mutually commute. We are therefore restricted to tessellations with $s = 4$ faces meeting at each face and with $r = 2c$ sides to each face, and for which $f(g_i) = 0$ for each generator g_i of the normal subgroup H defining the compactification.

APPENDIX D: GROUP THEORETIC CONDITION FOR CONSISTENT SCHEDULING

In Sec. IV B of the main text, we show that any translationally invariant schedule for the subsystem toric code assigns the same schedule to each triangle operator with the same *label*, where a label is an assignment of an element of the cyclic group $\mathbb{Z}_4$ to each triangle operator as shown in Fig. 25(b). We now describe this labeling of the

triangle operators of the subsystem toric code in terms of the proper symmetry group $G_{r,s}^{H+}$ of orientation-preserving symmetries of the lattice generated by the rotations ρ and σ [shown in Fig. 25(a)]. First note that, after choosing any triangle operator to be the fundamental domain, each triangle operator is now identified by a unique element in $G_{r,s}^{H+}$, and we denote by T_g the triangle operator identified by $g \in G_{r,s}^{H+}$. A labeling of the triangle operators is then defined by a function $h: G_{r,s}^{H+} \to \mathbb{Z}_4$. Note that, for the labeling of the subsystem toric code in Fig. 25(b), applying either a ρ or σ rotation to *any* triangle operator adds one (modulo 4) to the label. Using similar arguments to those given in Appendix C for valid colorings, we see that the function h must extend to a homomorphism $h: G_{r,s}^{H+} \to \mathbb{Z}_4$ with

$$h(\rho) = h(\sigma) = 1. \tag{D1}$$

We can generalize a translationally symmetric schedule of the subsystem toric code to subsystem hyperbolic codes by first labeling the triangle operators of a subsystem hyperbolic code in such a way that the neighborhood of each triangle operator is the same as it would be in the subsystem toric code, and then apply the same schedule to all triangle operators with the same label in the subsystem hyperbolic code. The *neighborhood* of a triangle operator T is the relative position and label of the triangle operators that overlap with T on at least one qubit (each of which we call a *neighbor*). We see from Fig. 25(b) that each triangle operator T_g in the subsystem toric code has seven neighbors: $T_{g\sigma}$, $T_{g\sigma^2}$, $T_{g\sigma^3}$, $T_{g\rho}$, $T_{g\rho\sigma}$, $T_{g\rho^{-1}}$, and $T_{g\rho^{-1}\sigma^{-1}}$. In the toric code, exactly three of these neighbors overlap on a vertex of the $\{4, 4\}$ tessellation. To ensure this remains the case for the hyperbolic tessellations, it is necessary to require that $s = 4$, which is by definition a property of our subsystem hyperbolic codes. Setting $s = 4$ alone is not sufficient, since we must now also ensure that the entire neighborhood (all *seven* neighbors) of each triangle operator with a given label in the lattice remains identical to that of a triangle operator with the same label in the toric code. The relative labels of these seven neighbors is determined by the homomorphism $h: G_{r,s}^{H+} \to \mathbb{Z}_4$ defined in Eq. (D1).

Therefore, a hyperbolic tessellation is schedulable if its proper symmetry group admits the homomorphism h as defined in Eq. (D1), which is the case if and only if $h(r_i) = 0$ for each relation r_i in the presentation of $G_{r,s}^{H+}$. This condition is met for subsystem hyperbolic codes derived from the subset of closed $\{4c, 4\}$ tessellations (where $c \in \mathbb{Z}^+$), for which the generators g_i of the normal subgroup H defining the compactification satisfies $h(g_i) = 0$. As an example, consider the quotient group for a distance L toric code which has presentation

$$G_{4,4}^{H+} = \langle \rho, \sigma | \rho^4 = \sigma^4 = (\rho\sigma)^2 = (\rho\sigma^{-1})^L = (\sigma^{-1}\rho)^L = e \rangle,$$

$$(D2)$$

from which it is clear that each relation r_i satisfies $h(r_i) = 0$.

For schedulable subsystem hyperbolic codes, we can use the very efficient measurement schedule of Ref. [34] (which is translationally invariant for the subsystem toric code) for each triangle operator, which requires only four time steps (one time step is the duration of a CNOT gate) to measure all X- and Z-check operators. Note that subsystem hyperbolic codes which do *not* satisfy these constraints will still admit a measurement schedule, but such a schedule may be considerably less efficient and also more difficult to construct.

Given the map $m: \mathbb{Z}_4 \to \mathbb{Z}_2$ defined by $m(x) = x \mod 2$ assigning a color to a label, we see that $f(g) = m(h(g)) \ \forall g \in G_{r,s}^{H+}$, where f is defined in Eq. (C1), and hence, every schedulable code is colorable [but not vice versa, as exemplified by the $\{6,4\}$ tessellation for which ρ^6 is a relation yet $h(\rho^6) \neq 0$].

There is another way of interpreting the scheduling: Consider the graph which is generated by the rotation subgroup $\langle \rho, \sigma \rangle$. This group acts regularly between the triangles of the subsystem code, so there is a one-to-one map between them. The labeling is a coloring of the Cayley graph of this group (each vertex of this Cayley graph corresponds to a triangle). This coloring is achieved by a "covering" of the cycle graph with four vertices (Cayley graph of $\mathbb{Z}_4$) since this is clearly four-colorable. More generally, we can consider normal subgroups N of the group as long as this normal subgroup does not contain ρ or σ. The number of colors in this case is the index of N in G.

The dual semihyperbolic tessellations used for constructing the subsystem semihyperbolic codes in Sec. IV C do not have a group structure, so they cannot be labeled using the homomorphism of Eq. (D1) alone. However, we now show that, given a schedulable $\{4c, 4\}$ tessellation, the corresponding dual semihyperbolic tessellation derived from it is also schedulable. Take a schedulable $\{4c, 4\}$ tessellation V, where we already label each corner in the tessellation with an element of $\mathbb{Z}/4\mathbb{Z}$. Now consider its dual tessellation V^* constructed by exchanging vertices and faces in the Hasse diagram of the tessellation [59]. Each corner in V is identified by a face and vertex, and so each corner in V is in one-to-one correspondence with a corner in V^* (where the face and vertex are exchanged). We give each corner in V^* the same label as the corner in V that it is in one-to-one correspondence with. This constitutes a valid labeling of V^*, since each pair of corners related by ρ (σ) in V are related by σ (ρ) in V^*, and $h(\rho) = h(\sigma)$ in Eq. (D1). We now construct a semihyperbolic tessellation V_l^* by tiling each face of V^* with an $l \times l$ square lattice. Note that the corners of each face in V^* are already labeled, so we can label V_l^* just by labeling the new corners introduced by the

$l \times l$ square tiling of each face. Corners related by a σ rotation in V^* are still related by a σ rotation in V_l^*. Corners related by a ρ rotation in V^* are now related by a $(\rho\sigma^{-1})^{l-1}\rho$ translation in V_l^*. However, now treating h as a function not a homomorphism, note that $h(\rho) = h((\rho\sigma^{-1})^{l-1}\rho)$, so the original labels retained from V^* remain valid. We can therefore label the new corners in the square $l \times l$ tilings in V_l^* in a way that is consistent with the corners already labeled. We now take the dual of V_l^* to obtain V_l, preserving the labels of each corner when taking the dual as before. The tessellation V_l is now used to derive a subsystem semihyperbolic code, and we demonstrate that V_l is schedulable if V is schedulable.

APPENDIX E: SUBSYSTEM SEMIHYPERBOLIC AND SUBSYSTEM TORIC CODE COMPARISON

A quantum code derived from an $\{r, s\}$ tessellation satisfies [6]

$$\frac{k}{n} = 1 - \frac{2}{s} - \frac{2}{r} + \frac{2}{n}, \tag{E1}$$

where n is the number of physical qubits and k is the number of logical qubits. A semihyperbolic code derived from such a code has $l^2 n$ qubits, where l is the dimension of the lattice tiling each face in the semihyperbolic code. Therefore, the number of data qubits (excluding ancillas) in a subsystem $\{8, 4\}$-semihyperbolic code with k logical qubits is $6(k-2)l^2$. To compare the performance of subsystem semihyperbolic codes with subsystem toric codes, we compare each semihyperbolic code to multiple independent copies of a toric code with the same rate k/n, such that we can compare the performance keeping k and n fixed. Since the rate of a subsystem toric code with distance L is $2/(3L^2)$, we compare our subsystem semihyperbolic $\{8, 4\}$ codes with copies of a toric code with distance close to

$$L = 2l\sqrt{1 - \frac{2}{k}}, \tag{E2}$$

where k is the number of qubits in the $\{8, 4\}$-semihyperbolic code, and l is the dimension of the lattice tiling each face in the semihyperbolic code. Note that the total number of qubits including ancillas $(1 + 4n_a/3)n$ is proportional to the number of data qubits n with the same constant of proportionality for the subsystem toric, hyperbolic, and semihyperbolic codes. Here, n_a is the number of ancilla qubits used per triangle operator (we can always set $n_a = 1$, but for some schedules setting $n_a = 2$ can improve performance by parallelizing the measurement schedule). Therefore, Eq. (E2) still holds once ancillas are taken into account.

APPENDIX F: DISTANCE OF SUBSYSTEM HYPERBOLIC CODES

We can determine the distance of the subsystem hyperbolic and semihyperbolic codes by considering their matching graphs. Each vertex in the X-type matching graph corresponds to an X stabilizer, and there is an edge between each pair of stabilizers u and v for which a single Z error on a data qubit anticommutes with both u and v. Each face in the X-type matching graph corresponds to a Z-type triangle operator. Each noncontractible closed loop in the X-type matching graph corresponds to a logical Z operator. Therefore, the Z distance of the code is determined by the shortest noncontractible closed loop in the X-type matching graph. A Z-type matching graph can be defined analogously for Z-type stabilizers, and so the X distance of the code is determined by the shortest noncontractible closed loop in the Z-type matching graph.

For the subsystem toric, hyperbolic, and semihyperbolic codes we construct, the X-type matching graph is isomorphic to the Z-type matching graph, since the Z-type matching graph can be obtained from the X-type matching graph (and vice versa) by a single rotation that is also a symmetry of the tessellation from which the code is derived. Therefore, the Z and X distances are the same for these codes.

We now consider how the distance of a subsystem hyperbolic or semihyperbolic code compares to the distance of the subspace CSS (surface) code derived from the same tessellation. To do so, we consider the structure of the matching graph for both codes. The solid red lines in Fig. 27 form the edges of the Z-type matching graph for the toric code, and so the length of the shortest noncontractible loop in that graph is the X distance of the toric code. We can obtain the X-type matching graph for the *subsystem* toric code derived from the same tessellation by adding in the green edges, also shown in Fig. 27, and keeping the same set of vertices. Each green edge in the subsystem toric code X-type matching graph is equivalent (up to a triangle operator) to a *pair* of red edges. Therefore, the distance between two vertices in the matching graph consisting only of red edges can *at most* be reduced by half by the inclusion of the green edges (and inclusion of the green edges cannot increase the distance between vertices).

For the subsystem hyperbolic and semihyperbolic codes, we again find that both the Z-type and X-type matching graphs can be constructed by adding additional edges to the Z-type matching graph V_Z of the subspace codes derived from the tessellation, where each of these additional edges is equivalent to a pair of edges in V_Z. Therefore, the shortest noncontractible loop in either the Z-type or X-type matching graph for a subsystem hyperbolic or semihyperbolic code is between 1 and 2 times smaller than the shortest noncontractible loop in the Z-type matching graph of the subspace code derived from the same tessellation. Consequently, given a hyperbolic or semihyperbolic code with X distance d_X, the distance d of the subsystem hyperbolic or semihyperbolic code derived from the same tessellation is bounded by $d_X/2 \le d \le d_X$. Furthermore, the X distance of hyperbolic codes we consider is always less than or equal to their Z distance. Both the subsystem toric code and standard toric code have distance $d = L$, but for the subsystem hyperbolic and semihyperbolic codes we construct, the subsystem codes do have a reduced distance compared to surface codes derived from the same tessellation. This is shown in Fig. 28, which compares the distance
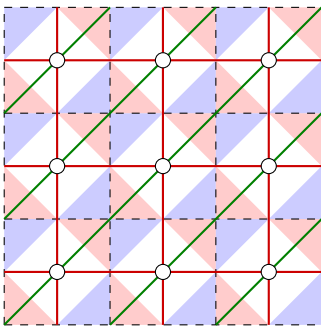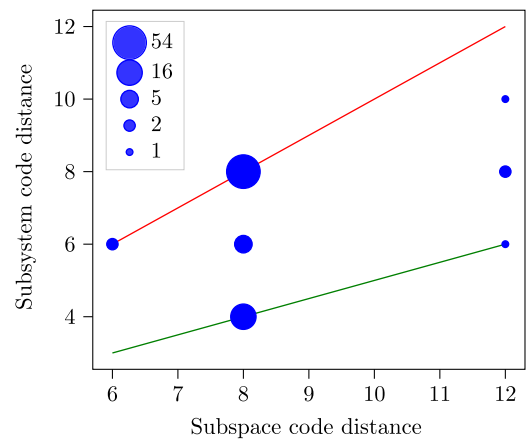

FIG. 27. The subsystem toric code. The black dashed lines are edges of the $\{4,4\}$ tessellation from which the subsystem toric code is derived. The edges in the X-type matching graph are the union of the solid red and green lines, and vertices in the matching graph are denoted by circles. Each edge in the X-type matching graph corresponds to a data qubit, and each face corresponds to a Z-type triangle operator. The solid red lines are the edges of the matching graph for the standard surface code derived from the same $\{4,4\}$ tessellation. Opposite sides are identified.

FIG. 28. For all $l = 2$ $\{8,4\}$ subsystem semihyperbolic codes we construct, here we plot the distance of each code (y axis) against the distance of the (subspace) semihyperbolic surface code derived from the same tessellation (x axis) computed using the method in Ref. [60]. The size of each blue circle corresponds to the number of codes we find with the same (x,y) coordinate on the figure, and the number of codes for each size of circle is given in the legend.
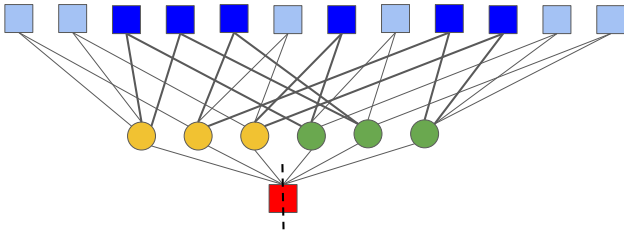
FIG. 29. Neighborhood of a Z stabilizer check in the Tanner graph of a CSS quantum code. Circles represent qubits, blue boxes in the top row represent X checks which are in the neighborhood of a Z check (red box at the bottom). Bold lines represent a cut set which induces a partition of the qubits into two sets (yellow and green) and cut on the Z check. X checks in dark blue belong to the set $\mathcal{I}$ and are merged.

of $l = 2$, $\{8,4\}$ subsystem semihyperbolic codes to the distance of the subspace semihyperbolic codes derived from the same tessellations. We see that the distance of each subsystem code can be reduced by up to $2\times$ relative to the subspace code derived from the same tessellation as expected, with some subsystem codes not suffering any reduction in distance.

APPENDIX G: THE CONSTRUCTION FOR GENERAL LDPC QUANTUM CODES

The ideas from Secs. III and IV of the main text readily apply to the more general class of CSS stabilizer codes. In a CSS stabilizer code, the stabilizer checks operate exclusively as either Pauli X or Pauli Z on all of the qubits in its support. The *Tanner graph* associated with a CSS code is the graph with vertices corresponding to qubits, X checks, and Z checks. There is an edge between two vertices if and only if one vertex corresponds to a check and the other to a qubit in its support. In order to define the merging and splitting for a CSS LDPC code, let us pick a Z check s_Z and consider the subgraph T' of the Tanner graph consisting of all qubits in the support of s_Z, their connected X checks, as well as all edges connecting them. We call a set of X checks in T' a local cut set if removing them and their incident edges from T' renders it disconnected (see Fig. 29). We call a local cut set *independent* if the checks contained are linearly independent. Let $\mathcal{I}$ be the labels of a local cut set. The checks belonging to $\mathcal{I}$ are *merged* by defining a new Tanner graph in which all of the vertices of $\mathcal{I}$ are identified. This procedure is also known as *vertex contraction* in the graph theory literature.

Merging an independent local cut set reduces the number of X checks by $|\mathcal{I}| - 1$. Since the number of physical qubits is not changed and the checks are independent, there must be $|\mathcal{I}| - 1$ new logical degrees of freedom. However, the operator algebra of these degrees of freedom will be supported on at most $|\text{supp}(s_Z)| \in O(1)$ physical qubits. Therefore, they do not offer protection and we consider them to act on gauge qubits.

We now describe a basis for the operator algebra acting on the gauge qubits. Let $C_1, \ldots, C_l$ be the qubits belonging to the connected components of T' induced by the cut set. We define Z-type operators Z_i^g which act on all qubits in C_i. Note that each Z_i^g commutes with all X checks in the code, because all X checks not belonging to the local cut set must overlap with s_Z on either the empty set or any of the C_i. Since all X checks commute with s_Z they must have even support on C_i and hence commute with Z_i^g. A set of operators which anticommute with the Z_i^g are the X checks in the local cut set.

The merging and splitting procedure reduces the number of linearly independent stabilizers r and increases the number of gauge operators g. The number of physical qubits n and logical qubits k are unaffected, so that Eq. (2) is satisfied.

We note that this procedure will generally affect the distance, as it does for the surface, toric, and hyperbolic codes. An extreme example is the surface code defined on a square lattice, where merging the top left and bottom right X checks of each face (Z check) leads to the code distance turning constant. Demonstrating that the procedure gives a subsystem code with growing distance therefore has to be informed by the structure of the code.

APPENDIX H: SCHEDULING FROM GROUP HOMOMORPHISMS

In Appendix D, we show that an efficient syndrome measurement schedule for subsystem hyperbolic codes could be constructed if the orientation-preserving symmetry group $G_{r,s}^{H+}$ of the tessellation (generated by rotations ρ and σ) admits a homomorphism $f : G_{r,s}^{H+} \to \mathbb{Z}_4$ to the cyclic group $\mathbb{Z}_4$, with f defined by $f(\rho) = f(\sigma) = 1$. This homomorphism is a useful tool for scheduling subsystem hyperbolic codes for the same reason that translation invariance is useful for scheduling Euclidean surface codes: The problem of scheduling the entire code reduces to the problem of scheduling only a small number of stabilizers in a region of the tessellation.

While the homomorphism $f : G_{r,s}^{H+} \to \mathbb{Z}_4$ is a useful tool for scheduling the subsystem hyperbolic codes, such a homomorphism exists only for a subset of $\{r, s\}$ tessellations (for which four divides both r and s). In this section, we look for homomorphisms from $G_{r,s}^{H+}$ to *any* cyclic group, in the hope that these homomorphisms will be a useful tool for scheduling subspace hyperbolic codes based on a wider range of tessellations, where each Z stabilizer (plaquette) and X stabilizer (site) is measured using a circuit with a single ancilla qubit. Each corner C_g of a face of the tessellation is identified with an element $g \in G_{r,s}^{H+}$. By finding a homomorphism $f : G_{r,s}^{H+} \to \mathbb{Z}_n$ to a cyclic group $\mathbb{Z}_n$, we can label each corner uniquely with an element in $\mathbb{Z}_n$. The function f is a homomorphism if and only if $f(r_i) = 0$ for each relation r_i in the presentation of $G_{r,s}^{H+}$. The tessellation group $G_{r,s}^{H+}$ has presentation

$$G_{r,s}^{H+} := \langle \rho, \sigma | (\rho\sigma)^2 = \rho^r = \sigma^s = e \rangle, \qquad \text{(H1)}$$

from which we see that $(\rho\sigma)^2$ is always a relation, and hence, f must always satisfy $f((\rho\sigma)^2) = 0$.

For the homomorphism $f: G_{r,s}^{H+} \to \mathbb{Z}_n$ to be useful for scheduling, we require that it must satisfy additional properties. First, the homomorphism should not be defined by $f(\rho) = f(\sigma) = 0$, since this homomorphism does not give us any additional information. Second, the label of each corner C_g should be different from the corner $C_{g\rho\sigma}$. This is because C_g and $C_{g\rho\sigma}$ overlap on an edge e in such a way that, if both corners had the same schedule, two CNOT gates applied to the qubit at e would occupy the same time step.

We assume that can have more than one ancilla for each stabilizer to parallelize the measurement circuits. If we instead insist that only a single ancilla be used, then we must require that all corners belonging to the same vertex must have different labels. This is because these corners share an ancilla qubit on the vertex, and two CNOT gates cannot be applied to the ancilla qubit within the same time step. Furthermore, we would also require that all corners belonging to a face must have a different label, since only a single CNOT gate can be applied to the ancilla qubit in the center of each face in each time step.

Therefore, for each tessellation $\{r, s\}$, we seek to find a cyclic group order n and elements $x, y \in \mathbb{Z}_n$ such that the function defined by $f(\rho) = x, f(\sigma) = y$ extends to a homomorphism $f: G_{r,s}^{H+} \to \mathbb{Z}_n$. The restrictions on the relations in the presentation of $G_{r,s}^{H+}$, along with the additional three properties we impose, correspond to the following constraints on x, y, n:

$$rx = 0 \quad \mathrm{mod}\ n,$$

$$sy = 0 \quad \mathrm{mod}\ n,$$

$$2(x + y) = 0 \quad \mathrm{mod}\ n,$$

$$x + y \neq 0 \quad \mathrm{mod}\ n, \qquad \text{(H2)}$$

and if we could use only a single ancilla per stabilizer, then we would additionally have the constraints

$$\mathrm{lcm}(x, n) = rx,$$

$$\mathrm{lcm}(y, n) = sy. \qquad \text{(H3)}$$

For all $r, s \leq 10$ we search for all n, x, y satisfying Eq. (H2) [for $n < 5 \max(r, s)$] and list all the tessellations we find which admit at least one such homomorphism in Table IV.

While we find homomorphisms to cyclic groups for many tessellations, we do not find any for the $\{5, 5\}$ code, which has the desirable properties of being self-dual and having low stabilizer weights. Therefore, an interesting question is whether there exist homomorphisms to groups

TABLE IV. Solutions to Eq. (H2) for all $r, s \leq 10, r \leq s$. By symmetry, solutions for $r \geq s$ can be found by exchanging column r with s and column x with y. For each tessellation $\{r, s\}$, we give the parameters n, x, y defining only one homomorphism $f: G_{r,s}^{H+} \to Z_n$ (the homomorphism which minimizes both n and x). There are at least two solutions for each tessellation.

r	s	n	x	y
3	6	6	2	1
4	4	4	1	1
4	8	4	1	1
5	10	10	2	3
6	6	6	1	2
6	9	6	1	2
8	8	8	1	3
10	10	10	1	4

that are not cyclic, and which contain a small number of elements, but otherwise satisfy the constraints of Eq. (H2). If such a homomorphism exists for tessellations such as $\{4, 5\}$ and $\{5, 5\}$, the trade-off of circuit-level threshold and encoding rate for these codes may be very favorable.

APPENDIX I: ADDITIONAL NUMERICAL RESULTS

In this section, we give some additional numerical results from simulations of the subsystem toric and semihyperbolic codes. In Table V, we give thresholds for the subsystem toric code under a circuit-level depolarizing noise model using gauge fixing with balanced schedules. In Table VI, we give thresholds for the subsystem toric code under an independent circuit-level noise model using both balanced and unbalanced schedules. In Fig. 30, we plot the threshold for the subsystem surface code with a

TABLE V. Thresholds (in %) for the subsystem toric code for some balanced homogeneous schedules under the circuit-level depolarizing noise model, each computed using the critical exponent method of Ref. [61] to analyze results from Monte Carlo simulations using subsystem toric codes with distances $L = 26$, 30, 34, 38, 42, 46. Numbers in brackets are the 1σ uncertainties in the last digit. For each threshold, we keep the number of syndrome extraction rounds constant for all codes, always using at least 92 rounds to ensure boundary effects (in time) are small even for the largest codes. For the column with an asterisk, gauge fixing is not used when decoding.

Schedule	$p_{\mathrm{depol}}^{\mathrm{th}}$	$p_{\mathrm{depol}}^{\mathrm{th},*}$
ZX	0.666(1)	0.666(1)
$Z^2 X^2$	0.757(1)	0.6587(9)
$Z^3 X^3$	0.810(2)	0.676(1)
$Z^4 X^4$	0.811(2)	0.669(2)
$Z^5 X^5$	0.792(2)	0.652(2)
$Z^{10} X^{10}$	0.522(2)	0.493(1)

TABLE VI.   Thresholds (in %) for the subsystem toric code for various homogeneous schedules under the independent circuit-level noise model, each computed using the critical exponent method of Ref. [61] to analyze results from Monte Carlo simulations using subsystem toric codes with distances $L = 26, 30, 34, 38, 42, 46$. Numbers in brackets are the 1σ uncertainties in the last digit. For each threshold, we keep the number of syndrome extraction rounds constant for all codes, always using at least 92 rounds to ensure boundary effects (in time) are small even for the largest codes. For the final two columns (with asterisks in the title), gauge fixing is not used even when possible.

Schedule	p_X^{th}	$p_X^{th,*}$	p_Z^{th}	$p_Z^{th,*}$
ZX	0.515(1)		0.515(1)	
Z^2X^2	0.5863(9)		0.5863(9)	
Z^3X^3	0.628(1)		0.628(1)	
Z^4X^4	0.631(2)		0.631(2)	
Z^5X^5	0.619(2)		0.619(2)	
ZX^2	0.3928(8)	0.3928(8)	0.749(1)	0.625(3)
ZX^3	0.3236(9)	0.3236(9)	0.931(1)	0.7234(9)
ZX^5	0.2449(5)	0.2449(5)	1.160(2)	0.816(2)
ZX^{10}	0.1595(4)	0.1595(4)	1.430(3)	0.902(2)
Z^2X^{10}	0.2394(5)	0.2259(5)	1.197(3)	0.821(2)
X	0	0	2.2231(1)	1.029(2)

phenomenological noise model, which we find to be 0.02004(2) using the critical exponent method of Ref. [61]. In Fig. 31, we plot the threshold of the $l = 2$ $\{8, 4\}$ subsystem semihyperbolic codes *without* adjusting for the number of logical qubits, unlike in the text. This is helpful to better understand the logical error rates of the codes themselves, but less so for understanding the threshold for the logical error rate per logical qubit, for which multiple independent copies of the smaller codes should be taken, as done in the main text. In Fig. 32, we compare the performance of a [[384,66,4]] $\{8, 4\}$ subsystem hyperbolic code with 33 copies of $L = 3$ and $L = 4$ subsystem toric codes, all encoding 66 logical qubits. Since this hyperbolic

FIG. 31.   Performance of the extremal subsystem $\{8, 4\}$ $l = 2$ semihyperbolic codes under a circuit-level depolarizing noise model. A homogeneous $(ZX)^{20}$ schedule is used for all codes, and the y axis is the probability that at least one logical Z error occurs. Dashed lines are the probability of a Z error occurring on at least one of k physical qubits without error correction under the same error model and for the same duration (80 time steps), with $k = 4$ (orange), $k = 8$ (purple), and $k = 10$ (pink).

code is quite small, its overhead $n/(kd^2) \approx 0.36$ is less favorable than that of the much larger [[8064,338,10]] subsystem semihyperbolic code analyzed in Sec. VI B of the main text, for which $n/(kd^2) \approx 0.24$. However, as can be seen from Fig. 32, the [[384,66,4]] subsystem hyperbolic code still uses $2.3\times$ fewer physical qubits than the subsystem toric code to achieve the same logical error rate per logical qubit below a circuit-level depolarizing physical error rate of 0.1%. Furthermore, it requires only 896 physical qubits to implement this subsystem hyperbolic code including ancillas, compared to the 18816 needed for the [[8064,338,10]] subsystem semihyperbolic code.
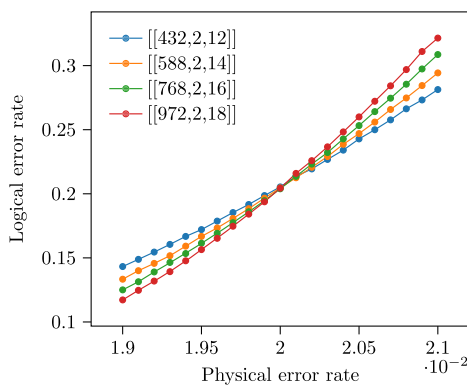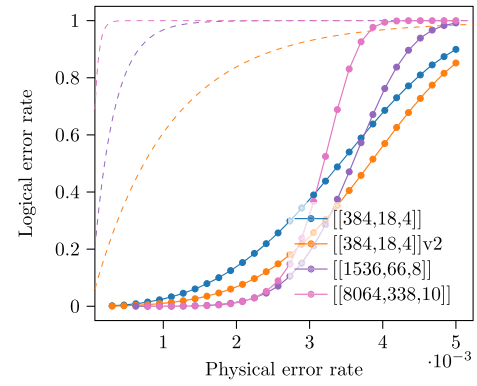

FIG. 30.   Subsystem toric code threshold with a phenomenological noise model, and without using gauge fixing (triangular lattice-matching graph). Using the critical exponent method, we find a threshold of 0.02004(2).
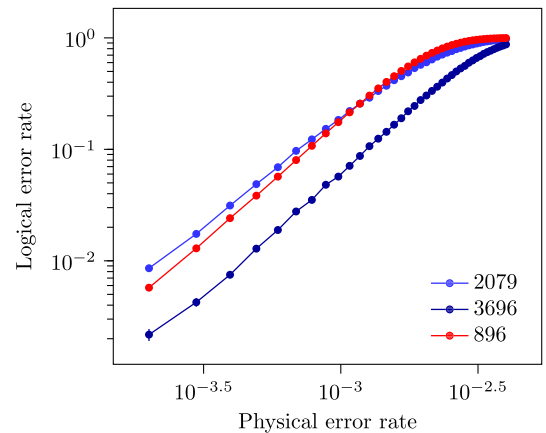

FIG. 32.   Performance of a [[384,66,4]] $\{8, 4\}$ subsystem hyperbolic code (red) compared to the $L = 3$ and $L = 4$ subsystem toric codes (shades of blue) using a $(ZX)^{10}$ schedule with the circuit-level depolarizing error model. We use 33 independent copies of the subsystem toric codes to fix the number of logical qubits at $k = 66$. In the legend, we give the number of physical qubits used, including ancillas.

[1] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological Quantum Memory*, J. Math. Phys. (N.Y.) **43**, 4452 (2002).

[2] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Surface Codes: Towards Practical Large-Scale Quantum Computation*, Phys. Rev. A **86**, 032324 (2012).

[3] N. P. Breuckmann and J. N. Eberhardt, *LDPC Quantum Codes*, arXiv:2103.06309.

[4] J.-P. Tillich and G. Zémor, *Quantum LDPC Codes with Positive Rate and Minimum Distance Proportional to the Square Root of the Blocklength*, IEEE Trans. Inf. Theory **60**, 1193 (2013).

[5] M. H. Freedman, D. A. Meyer, and F. Luo, in *Mathematics of Quantum Computation* (Chapman and Hall/CRC Press, London, 2002), pp. 287–320.

[6] N. P. Breuckmann and B. M. Terhal, *Constructions and Noise Threshold of Hyperbolic Surface Codes*, IEEE Trans. Inf. Theory **62**, 3731 (2016).

[7] L. Guth and A. Lubotzky, *Quantum Error Correcting Codes and 4-Dimensional Arithmetic Hyperbolic Manifolds*, J. Math. Phys. (N.Y.) **55**, 082202 (2014).

[8] N. P. Breuckmann and V. Londe, *Single-Shot Decoding of Linear Rate LDPC Quantum Codes with High Performance*, arXiv:2001.03568.

[9] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, *Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits*, Phys. Rev. X **10**, 011022 (2020).

[10] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, P. V. Klimov *et al.*, *Exponential Suppression of Bit or Phase Flip Errors with Repetitive Error Correction*, arXiv:2102.06132.

[11] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais *et al.*, *Bias-Preserving Gates with Stabilized Cat Qubits*, Sci. Adv. **6**, eaay5901 (2020).

[12] D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, *Fault-Tolerant Thresholds for the Surface Code in Excess of 5% under Biased Noise*, Phys. Rev. Lett. **124**, 130501 (2020).

[13] J. P. Bonilla-Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, *The XZZX Surface Code*, Nat. Commun. **12**, 2172 (2021).

[14] X. Xu, Q. Zhao, X. Yuan, and S. C. Benjamin, *High-Threshold Code for Modular Hardware with Asymmetric Noise*, Phys. Rev. Applied **12**, 064006 (2019).

[15] A. M. Stephens, W. J. Munro, and K. Nemoto, *High-Threshold Topological Quantum Error Correction against Biased Noise*, Phys. Rev. A **88**, 060301(R) (2013).

[16] D. Gottesman, *Stabilizer Codes and Quantum Error Correction*, arXiv:quant-ph/9705052.

[17] D. Poulin, *Stabilizer Formalism for Operator Quantum Error Correction*, Phys. Rev. Lett. **95**, 230504 (2005).

[18] P. Aliferis and A. W. Cross, *Subsystem Fault Tolerance with the Bacon-Shor Code*, Phys. Rev. Lett. **98**, 220502 (2007).

[19] M. Li, D. Miller, and K. R. Brown, *Direct Measurement of Bacon-Shor Code Stabilizers*, Phys. Rev. A **98**, 050301(R) (2018).

[20] A. Paetznick and B. W. Reichardt, *Universal Fault-Tolerant Quantum Computation with Only Transversal Gates and Error Correction*, Phys. Rev. Lett. **111**, 090505 (2013).

[21] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa *et al.*, *Demonstration of Quantum Volume 64 on a Superconducting Quantum Computing System*, Quantum Sci. Technol. **6**, 025020 (2021).

[22] D. Bacon, *Operator Quantum Error-Correcting Subsystems for Self-Correcting Quantum Memories*, Phys. Rev. A **73**, 012340 (2006).

[23] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina *et al.*, *Fault-Tolerant Operation of a Quantum Error-Correction Code*, arXiv:2009.11482.

[24] M. Li, D. Miller, M. Newman, Y. Wu, and K. R. Brown, *2D Compass Codes*, Phys. Rev. X **9**, 021041 (2019).

[25] C. Gidney and M. Ekerå, *How to Factor 2048 Bit RSA Integers in 8 Hours Using 20 Million Noisy Qubits*, Quantum **5**, 433 (2021).

[26] N. P. Breuckmann, C. Vuillot, E. Campbell, A. Krishna, and B. M. Terhal, *Hyperbolic and Semi-Hyperbolic Surface Codes for Quantum Storage*, Quantum Sci. Technol. **2**, 035007 (2017).

[27] M. Li and T. J. Yoder, *A Numerical Study of Bravyi-Bacon-Shor and Subsystem Hypergraph Product Codes*, arXiv:2002.06257.

[28] J. Conrad, C. Chamberland, N. P. Breuckmann, and B. M. Terhal, *The Small Stellated Dodecahedron Code and Friends*, Phil. Trans. R. Soc. A **376**, 20170323 (2018).

[29] M. Suchara, S. Bravyi, and B. Terhal, *Constructions and Noise Threshold of Topological Subsystem Codes*, J. Phys. A **44**, 155301 (2011).

[30] B. M. Terhal, *Quantum Error Correction for Quantum Memories*, Rev. Mod. Phys. **87**, 307 (2015).

[31] H. Bombin, *Gauge Color Codes: Optimal Transversal Gates and Gauge Fixing in Topological Stabilizer Codes*, New J. Phys. **17**, 083002 (2015).

[32] T. J. Yoder, *Universal Fault-Tolerant Quantum Computation with Bacon-Shor Codes*, arXiv:1705.01686.

[33] C. Vuillot, L. Lao, B. Criger, C. G. Almudéver, K. Bertels, and B. M. Terhal, *Code Deformation and Lattice Surgery Are Gauge Fixing*, New J. Phys. **21**, 033028 (2019).

[34] S. Bravyi, G. Duclos-Cianci, D. Poulin, and M. Suchara, *Subsystem Surface Codes with Three-Qubit Check Operators*, Quantum Inf. Comput. **13**, 963 (2013).

[35] N. C. Brown, M. Newman, and K. R. Brown, *Handling Leakage with Subsystem Codes*, New J. Phys. **21**, 073055 (2019).

[36] H. Bombín and M. A. Martin-Delgado, *Optimal Resources for Topological Two-Dimensional Stabilizer Codes: Comparative Study*, Phys. Rev. A **76**, 012305 (2007).

[37] K. Fujii and Y. Tokunaga, *Error and Loss Tolerances of Surface Codes with General Lattice Structures*, Phys. Rev. A **86**, 020303(R) (2012).

[38] R. Raussendorf, J. Harrington, and K. Goyal, *Topological Fault-Tolerance in Cluster State Quantum Computation*, New J. Phys. **9**, 199 (2007).

[39] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, *Surface Code Quantum Computing with Error Rates over 1%*, Phys. Rev. A **83**, 020302(R) (2011).

[40] A. M. Stephens, *Fault-Tolerant Thresholds for Quantum Error Correction with the Surface Code*, Phys. Rev. A **89**, 022321 (2014).

[41] A. J. Kollár, M. Fitzpatrick, and A. A. Houck, *Hyperbolic Lattices in Circuit Quantum Electrodynamics*, Nature (London) **571**, 45 (2019).

[42] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, *Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links*, Phys. Rev. X **4**, 041041 (2014).

[43] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement Distillation between Solid-State Quantum Network Nodes*, Science **356**, 928 (2017).

[44] B. J. Brown, N. H. Nickerson, and D. E. Browne, *Fault-Tolerant Error Correction with the Gauge Color Code*, Nat. Commun. **7**, 1 (2016).

[45] A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg, *Towards Practical Classical Processing for the Surface Code*, Phys. Rev. Lett. **108**, 180501 (2012).

[46] S. Huang, M. Newman, and K. R. Brown, *Fault-Tolerant Weighted Union-Find Decoding on the Toric Code*, Phys. Rev. A **102**, 012419 (2020).

[47] P. Aliferis and J. Preskill, *Fault-Tolerant Quantum Computation against Biased Noise*, Phys. Rev. A **78**, 052331 (2008).

[48] J. Guillaud and M. Mirrahimi, *Repetition Cat Qubits for Fault-Tolerant Quantum Computation*, Phys. Rev. X **9**, 041053 (2019).

[49] O. Higgott, PYMATCHING: *A Python Package for Decoding Quantum Codes with Minimum-Weight Perfect Matching*, arXiv:2105.13082.

[50] J. Edmonds, *Paths, Trees, and Flowers*, Can. J. Math. **17**, 449 (1965).

[51] V. Kolmogorov, BLOSSOM V: *A New Implementation of a Minimum Cost Perfect Matching Algorithm*, Math. Program. Comput. **1**, 43 (2009).

[52] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, *Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity*, Phys. Rev. X **8**, 041015 (2018).

[53] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, *Surface Code Quantum Computing by Lattice Surgery*, New J. Phys. **14**, 123011 (2012).

[54] H. Bombín and M. A. Martin-Delgado, *Quantum Measurements and Gates by Code Deformation*, J. Phys. A **42**, 095302 (2009).

[55] N. Delfosse and N. H. Nickerson, *Almost-Linear Time Decoding Algorithm for Topological Codes*, arXiv: 1709.06218.

[56] H. Bombín, *Topological Subsystem Codes*, Phys. Rev. A **81**, 032301 (2010).

[57] D. P. DiVincenzo and F. Solgun, *Multi-Qubit Parity Measurement in Circuit Quantum Electrodynamics*, New J. Phys. **15**, 075001 (2013).

[58] L. Mosher, *When Can a Homomorphism Be Determined Entirely by Its Generators, Mathematics Stack Exchange*, https://math.stackexchange.com/q/1402612 (version 08-19-2015).

[59] N. P. Breuckmann, *Homological Quantum Codes beyond the Toric Code*, Ph.D. thesis, RWTH Aachen University, 2017.

[60] J. Erickson and K. Whittlesey, *Greedy Optimal Homotopy and Homology Generators*, in *Proceedings of the 16th Annual ACM-SIAM Symposium of Discrete Algorithms* (2005), Vol. 5, pp. 1038–1046, https://doi.org/10.5555/1070432.1070581.

[61] C. Wang, J. Harrington, and J. Preskill, *Confinement-Higgs Transition in a Disordered Gauge Theory and the Accuracy Threshold for Quantum Memory*, Ann. Phys. (Amsterdam) **303**, 31 (2003).