

Almost linear decoder for optimal geometrically local quantum codes

Quinten Eggerickx ^{1,2,3,*} Adam Wills ^{4,5} Ting-Chun Lin ^{6,5} Kristiaan De Greve ^{1,2} and Min-Hsiu Hsieh ⁵

¹*Imec, Kapeldreef 75, 3001 Leuven, Belgium*

²*Department of Electrical Engineering, KU Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium*

³*Department of Physics and Astronomy, Ghent University, Krijgslaan 281, 9000 Gent, Belgium*

⁴*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

⁵*Hon Hai Research Institute, Taipei, Taiwan*

⁶*Department of Physics, University of California San Diego, San Diego, California 92093, USA*



(Received 16 December 2024; revised 15 April 2025; accepted 11 May 2025; published 24 June 2025)

Geometrically local quantum codes, which are error-correction codes embedded in \mathbb{R}^D with checks acting only on qubits within a fixed spatial distance, have garnered significant interest. Recently, it has been demonstrated how to achieve geometrically local codes that maximize both the dimension and the distance, as well as the energy barrier of the code. In this work, we focus on the constructions involving subdivision, and we show that they have an almost linear time decoder, obtained by combining the decoder of the outer good qLDPC code and a generalized version of the Union-Find decoder. This provides the first decoder for an optimal geometrically local three-dimensional code. We demonstrate the existence of a finite threshold error rate under the code capacity noise model using a minimum weight perfect matching decoder. Furthermore, we argue that this threshold is also applicable to the decoder based on the generalized Union-Find algorithm.

DOI: [10.1103/sx dx-qbcz](https://doi.org/10.1103/sx dx-qbcz)

I. INTRODUCTION

Large-scale fault-tolerant quantum computers hold the promise of revolutionizing our world by being able to solve certain computational problems far beyond the reach of conventional computers. A key element for fault-tolerance is the quantum error-correction (QEC) code, which encodes the logical state of the system over many physical qubits and in that way protects the logical state against noise.

However, over time, errors accumulate and logical errors start to become more and more likely. Encoding on its own is therefore not enough. One also needs a decoder that is capable of predicting which errors occurred, such that these errors can be corrected. Moreover, it is crucial for fault-tolerant computation that this decoding process can be performed efficiently [1,2].

Among the different quantum codes, quantum low-density parity-check (qLDPC) codes have garnered significant attention. Practically, qLDPC codes are preferred because each stabilizer operates on only a limited number of qubits, and each qubit is operated on by only a limited number of stabilizers. This low density feature is advantageous for experiments, as it limits the spread of errors during the extraction of the syndrome. Theoretically, qLDPC codes have become the focus of research since Gottesman [3] showed that fault-tolerant

computation can be achieved with only a constant overhead by using qLDPC codes with good properties. Furthermore, recent developments over the past few years have led to constructions [4–6] of *good* qLDPC codes achieving asymptotically optimal linear dimension and distance.

Although good qLDPC codes accomplish the best performance possible, it is well-established [7,8] that these codes cannot be implemented solely with local interactions, which poses a significant challenge to their practical implementation. Recently, this issue was addressed in [9–12] by taking a good qLDPC code and transforming it into a geometrically local code. Evidently, the resulting code is no longer a *good* code, but does saturate the Bravyi-Poulin-Terhal (BPT) bound [7,8] and achieves an optimal energy barrier.

In this paper, we construct a decoder for these geometrically local quantum codes.

A. Main contribution

Our main contribution is the construction of an almost linear-time serial decoder for the geometrically local codes from subdivision, developed in [10] and later generalized in [12], by extending the Union-Find decoder [13] to the generalized surface code.

Theorem 1.1 (informal). The decoder for the subdivided code, as described in Algorithm I, runs in almost linear time $O(n\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann function, and it can correct an error of weight up to a constant fraction of the distance.

This provides a positive answer to the first question raised in [10] by showing that indeed an efficient decoder, capable of correcting a number of errors up to a constant times the

*Contact author: quinten.eggerickx@imec.be

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

distance, can be obtained by combining the decoder for the surface code and the decoder for the original qLDPC code.

Moreover, our decoder can also be applied at once to the more general optimal geometrically local codes of [12]. Although the local product structure is lacking in this generalized construction, all the elements that are needed for our decoder, i.e., patches of surface codes on top of a qLDPC scaffold, are present.

Our decoder is the first decoder for optimal geometrically local quantum codes in any dimension. In particular, it is the first almost linear time decoder for optimal three-dimensional (3D) geometrically local codes. While previous work has demonstrated fast serial and parallel decoding in three dimensions by combining two-dimensional decoding subproblems [14], these approaches were not applied to optimal codes.

Although much research is being conducted on the 3D integration of superconducting qubits, experimental progress remains in its early stages. Current demonstrations are largely limited to flip-chip integration, where a quantum chip and a control chip are connected using through-silicon vias. In this layout, qubits reside on the quantum chip, while read-out resonators and control lines are housed on the control chip [15–17]. While there is no experimental realization yet, these techniques could, in principle, be extended to stack multiple quantum chips, creating a 3D lattice of qubits. Moreover, advances in 3D integration for classical Complementary Metal-Oxide-Semiconductor (CMOS) technology [18] could inform and potentially accelerate the development of such architectures for superconducting qubits.

For spin qubits, the looped pipeline architecture represents a promising approach to achieving effective 3D local connectivity. While this concept has not yet been implemented at the scale envisioned (e.g., Fig. 4 in [19]), significant milestones have been reached. Coherent spin shuttling through a minimal loop of three germanium quantum dots has been demonstrated with an average fidelity per shuttle of $F \simeq 99.63\%$ [20]. This can be viewed as a foundational step towards building larger looped pipeline architectures. Similarly, recent achievements in silicon quantum dots include spin shuttling over an accumulated distance of $10\ \mu\text{m}$ (100 quantum dots) with an average fidelity per shuttle of $F \simeq 99.99\%$ [21]. Extrapolating these results to the looped pipeline architecture suggests that a 3D code with 100 layers in the third dimension could be feasible.

While these advances are promising, significant technical challenges remain before practical implementation of 3D geometrically local codes becomes a reality. Nonetheless, these developments highlight the potential for further exploration and innovation in this area.

Our linear time decoder thus supports the ongoing development of these systems, allowing them to scale to practically useful sizes.

We also considered the subdivided code under the code capacity noise model, and we showed that, using a combination of minimum weight perfect matching and a good qLDPC decoder, there exists a finite threshold error rate.

Theorem 1.2. Geometrically local codes from subdivision, together with a combined minimum weight perfect matching and good qLDPC decoder, have a finite threshold error rate. Moreover, logical errors are suppressed exponentially with the subdivision parameter L .

This is a demonstration of a nontrivial threshold for optimal geometrically local codes in any dimension, thereby strengthening the argument for their practical application.

Additionally, we present a reasoned argument suggesting that this result may also apply to the almost linear-time subdivided decoder outlined in Algorithm I.

B. Construction and proof overview

Since the geometrically local codes from subdivision are formed by taking a good qLDPC code as a scaffold and filling in the gaps with $L \times L$ squares of the surface code, one can identify three distinct types of regions in the code (Fig. 1). The first type are squares of the surface code surrounding an X stabilizer in the original good qLDPC code. Due to the structure of the good qLDPC code, this region will form a patch of a generalized surface code where multiple squares of the surface code are glued together [Fig. 2(a)]. The second type of region are the boundaries of the patches of generalized surface code. Similarly, these form generalized repetition codes [Fig. 2(b)]. The last type of region are the boundaries of the type two region, and they correspond to the Z stabilizers in the good qLDPC code.

The decoder of the subdivided code is constructed by applying a decoder to each of the three types of regions. The decoder on the patches of generalized surface code corrects errors to the best extent possible and pushes any remaining errors to its boundary. The decoder on the sections of generalized repetition code then takes on these additional errors and also corrects them as best as possible and pushes any remaining errors to its boundary. Finally, the decoder of the good qLDPC code takes care of the remaining errors.

The proof of the correctness and the time-complexity of the decoder primarily focuses on dealing with hyperedges and clusters of errors spanning over multiple squares of surface code. The key property that is repeatedly used to solve these issues is the low-density property of the underlying good qLDPC code.

At a high level, our approach relies on the fact that the subdivision code is chain homotopy equivalent to the original good qLDPC code, a strategy also discussed in Appendix A of Ref. [22]. Our decoding procedure essentially consists of two stages: the first maps the subdivided code back to the original code, and the second applies the decoder for the original qLDPC code. Both stages can be naturally extended to the codes in [9,11].

C. Future directions

a. Fast parallel decoder? There are several fast parallel decoders for the surface code under random noise [2,23–27]. Can these be used to lead to a faster parallel decoder for the subdivided code under random noise?

b. Application of parallel window decoding? It was recently demonstrated [2,28] how to parallelize a decoder using parallel window decoding. Can the subdivided decoder, presented in this paper, be modified such that this procedure can be applied? Can it be applied in both the space and time directions?

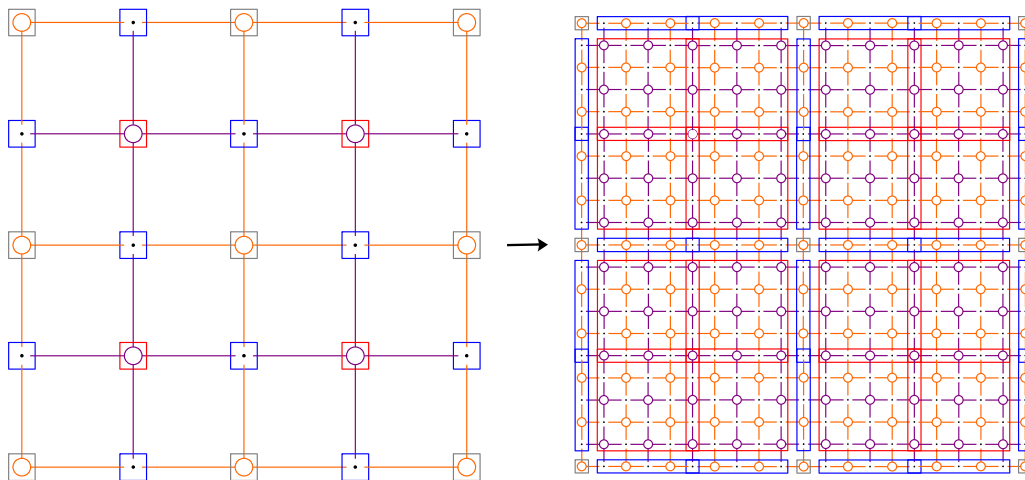


FIG. 1. A scaffold code gets filled in with squares of the surface code. Z stabilizers, qubits, and X stabilizers are represented by orange circles, black dots, and purple circles, respectively. Three distinct types of regions can be identified in the subdivided code. The type 1 regions (patches of generalized surface code) are indicated in red, and the type 2 regions (sections of generalized repetition code) are indicated in blue. The type 3 regions are not specially indicated as they simply correspond to the original Z stabilizers.

II. PRELIMINARY

A. Chain complexes

Definition II.1 (Chain complex). A chain complex X consists of a sequence of vector spaces $\mathbb{F}_2^{X(i)}$ generated by sets $X(i)$, along with linear maps $\delta_i : \mathbb{F}_2^{X(i)} \rightarrow \mathbb{F}_2^{X(i+1)}$ known as coboundary operators, where the coboundary operators satisfy

$$\delta_{i+1}\delta_i = 0.$$

Given a chain complex, one obtains the dual chain complex by reversing the direction of the linear maps. The coboundary operators δ_i then become boundary operators ∂_i . By using the elements of $X(i)$ as a canonical basis of $\mathbb{F}_2^{X(i)}$, the boundary operators $\partial_i : \mathbb{F}_2^{X(i)} \rightarrow \mathbb{F}_2^{X(i-1)}$ can be defined as the matrix transpose of δ_{i-1} , $\partial_i := \delta_{i-1}^T$. Note that the dual chain complex is also a chain complex since

$$\partial_{i-1}\partial_i = 0.$$

We provide some standard terminology. Elements of the kernel of the boundary operators are called cycles, while elements of the kernel of the coboundary operators are called

cocycles,

$$Z_i := \ker \partial_i = \{c_i \in \mathbb{F}_2^{X(i)} : \partial_i c_i = 0\},$$

$$Z^i := \ker \delta^i = \{c^i \in \mathbb{F}_2^{X(i)} : \delta^i c^i = 0\}.$$

Similarly, elements of the image of the boundary operators are called boundaries, and elements of the image of the coboundary operator are called coboundaries,

$$B_i := \text{im } \partial_{i+1} = \{\partial_{i+1} c_{i+1} : c_{i+1} \in \mathbb{F}_2^{X(i+1)}\},$$

$$B^i := \text{im } \delta^{i-1} = \{\delta^{i-1} c^{i-1} : c^{i-1} \in \mathbb{F}_2^{X(i-1)}\}.$$

Since $\partial_i \partial_{i+1} = 0$, it follows that $B_i \subset Z_i$. When $B_i = Z_i$, the chain complex is said to be *exact* at i .

B. Classical and quantum error-correcting codes

1. Classical error-correcting codes

A classical binary code is a k -dimensional linear subspace $C \subset \mathbb{F}_2^n$, which is defined as the kernel of a parity-check matrix $H : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$. n is called the size and k is called the dimension of the code. The distance d is the minimum Hamming weight of a nontrivial codeword,

$$d = \min_{c \in C - \{0\}} |c|. \tag{1}$$

The energy barrier \mathcal{E} is defined as the minimum energy among all walks γ from 0 to a nontrivial codeword,

$$\mathcal{E} = \min_{\gamma_{0 \rightarrow c}, c \in C - \{0\}} \epsilon(\gamma_{0 \rightarrow c}), \tag{2}$$

where a walk $\gamma_{0 \rightarrow c}$ is a sequence of vectors ($C_0 = 0, c_1, \dots, c_t = c$) with $|c_i - c_{i+1}| = 1$, and the energy of a walk is the maximum energy of the vectors in it: $\epsilon(\gamma) = \max_{c_i \in \gamma} |H c_i|$.

Most often we are interested in the scaling behavior of codes. Therefore, we do not just consider one specific code, but rather a whole family of codes where the size of the code n is taken as a parameter of the family. A code family is called a low-density parity-check (LDPC) code if each row and each column of the parity-check matrix has a bounded number of

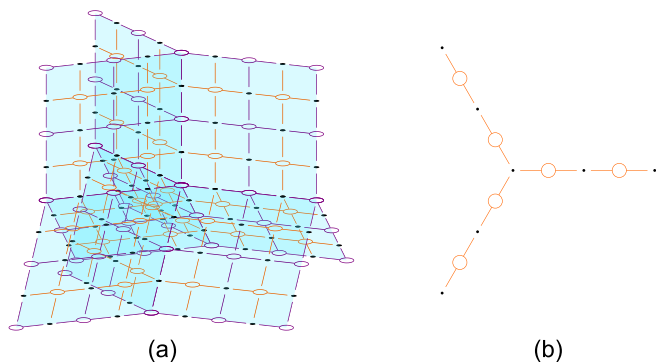


FIG. 2. (a) Generalized surface code and (b) generalized repetition code.

nonzero elements. A code family is said to be *good* if both the dimension and the distance scale linearly with n : $k = \Theta(n)$ and $d = \Theta(n)$.

2. Quantum error-correcting codes

In general, a quantum code is defined to be a linear subspace of the Hilbert space of the appropriate dimension. For a system of n qubits, this Hilbert space is \mathbb{C}^{2^n} .

An important class of quantum codes are the so-called *stabilizer codes*. For these codes, the codespace (subspace) is determined as the simultaneous +1 eigenspace of an Abelian subgroup of the n -fold Pauli group, called the *stabilizer group* S . Elements of S are called *stabilizers*.

The correction of Pauli errors in a stabilizer code is based on the measurement of the stabilizers. The result of these measurements is called the *syndrome* of the error. The task of the *decoder* is then to take the syndrome as an input and determine which correction needs to be applied to avoid a *logical error*. Logical errors commute with all the stabilizers, but they are not stabilizers themselves. As such, they act nontrivially on the codespace and must therefore be avoided.

3. Quantum erasures

Quantum erasures are a special type of error where we know which qubits have suffered an error, but we do not know what the error was. We say that some qubits have been erased. The set of erased qubits is called the *erasure* and is denoted by ε .

For erasures, the decoding task is greatly simplified since we only need to find an error $\tilde{P} \subset \varepsilon$ that is consistent with the syndrome to find the most likely correction. Since operators of the stabilizer S act trivially on the code space, the most likely correction is a coset $\tilde{P}S$ that maximizes the conditional probability $\mathbb{P}(PS|\varepsilon, \sigma)$.

Theorem II.1 [Lemma 1 of [29] (generalized to any stabilizer code)]. Given an erasure ε for a stabilizer code with stabilizer S and a measured syndrome σ , any coset $\tilde{P}S$ of a Pauli error $\tilde{P} \in \varepsilon$ of syndrome σ is a most likely coset.

4. Quantum CSS codes

A quantum CSS code [30,31] is a stabilizer code, specified by two classical codes C_z, C_x represented by their parity-check matrices $H_z : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{m_z}, H_x : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{m_x}$ which satisfy $H_x H_z^T = 0$. This condition allows us to associate a three-term

chain complex to the quantum code,

$$X : \mathbb{F}_2^{m_x} \xrightarrow{\delta_0=H_x^T} \mathbb{F}_2^n \xrightarrow{\delta_1=H_z} \mathbb{F}_2^{m_z}. \tag{3}$$

Here, n is the number of qubits, m_x is the number of X checks, and m_z is the number of Z checks. The dimension $k = \dim C_x - \dim C_z$ is the number of logical qubits. The distance $d = \min(d_x, d_z)$, where

$$d_z = \min_{c_z \in C_z - C_x^\perp} |c_z|, \quad d_x = \min_{c_x \in C_x - C_z^\perp} |c_x| \tag{4}$$

are the Z and X distances.

C. Kitaev surface codes

Kitaev surfaces codes [32,33] are quantum CSS codes where the qubits are placed on a two-dimensional surface and where all parity checks are local, i.e., they only involve nearest-neighbor interactions. More specifically, given a surface and a cellulation of that surface described by $G = (V, E, F)$, with V the set of vertices, E the set of edges, and F the set of faces, the code is defined by placing a qubit on each edge $e \in E$, a Z -parity-check $Z(v)$ on each vertex $v \in V$, and an X -parity-check $X(f)$ on each face $f \in F$ [34].

Since $Z(v) = \prod_{e \in e_v} X_e$ only acts on the qubits of edges incident to v , and $X(f) = \prod_{e \in f} Z_e$ only acts on the qubits of edges on the boundary of f , the code is local. Moreover, the code is geometrically local for geometrically local tessellations (see Definition II.2).

Note that both surfaces with and without boundaries can be considered.

D. Subdivided codes

Definition II.2 (Geometrically local in D dimensions). A code family $\{C(n)\}$ is called geometrically local in D dimensions if there exist a locality parameter $a \in \mathbb{R}_{>0}$, a density $b \in \mathbb{R}_{>0}$, and an embedding in \mathbb{Z}^D such that for each $C(n)$

- (i) the Euclidean distance between each check and the qubit it interacts with is $\leq a$,
- (ii) the number of qubits and checks located at each lattice point in \mathbb{Z}^D is $\leq b$.

Several good quantum LDPC codes have been identified [4–6]. Despite these codes exhibiting optimal scaling behavior, it is well established that they cannot be realized using only local interactions [7]. Restricting oneself to geometrically local codes, the distance and dimension of the code are




FIG. 3. The subdivision of a square of the code.

upper bounded by the BPT bound [7,8]:

$$d = O(n^{\frac{D-1}{D}}), \quad k = O(n^{\frac{D-2}{D}}) \quad (5)$$

with D the dimension of the space in which the code is embedded.

Recently, it was shown [9–12] how to convert a good quantum LDPC code into a geometrically local code that saturates the BPT bound. Focusing on [10], the key idea is to use the local product structure that most of the good quantum LDPC codes possess to represent the code as a square complex (a set of vertices, edges, and squares) and apply a subdivision process to locally embed it in \mathbb{Z}^D . Finally, to obtain the geometrically local code, the subdivided square complex is again represented as a chain complex. Intuitively, the code is made local by forcing an embedding in \mathbb{Z}^D and removing any long-range connections by subdividing them with additional qubits and checks.

The subdivision works as follows. Pick an odd number L and subdivide each square of the square complex into an $L \times L$ grid. Place a Cartesian coordinate system (i, j) on the grid, oriented such that the X -check corresponds to $(0, 0)$ and the Z -check corresponds to (L, L) . The subdivided chain complex, $X_L : \mathbb{F}_2^{X_L(0)} \rightarrow \mathbb{F}_2^{X_L(1)} \rightarrow \mathbb{F}_2^{X_L(2)}$, is then defined by letting

- (i) $X_L(0)$ (i.e., the X -checks) be the set of vertices with i, j both being even.
- (ii) $X_L(1)$ (i.e., the qubits) be the set of vertices with i, j one being even, one being odd.
- (iii) $X_L(2)$ (i.e., the Z -checks) be the set of vertices with i, j both being odd.

In other words, each square of the code is filled in with a patch of the standard, rectangular surface code. The subdivision of one square of the code is shown in Fig. 3.

The geometrically local codes from [10] saturate¹ the BPT bound [Eq. (5)] and possess an optimal energy barrier $\mathcal{E} = \Omega(n^{\frac{D-2}{D}})$. In terms of the parameters of the underlying good quantum LDPC (qLDPC) code $[[\tilde{n}, \tilde{k}, \tilde{d}]]$, the parameters of the subdivided code are given by

$$n = \Theta(L^2 \tilde{n}), \quad (6a)$$

$$k = \tilde{k}, \quad (6b)$$

$$d = \Theta(L \tilde{d}), \quad (6c)$$

$$\mathcal{E} = \Theta(\tilde{\mathcal{E}}). \quad (6d)$$

E. Generalized repetition code and generalized surface code

The repetition code consists of a 1D chain of bits where all the bits are required to take the same value. The generalized repetition code is similar, but it may include a branching point in the center [10], see Fig. 4. We say that a generalized repetition code has length L if there are L bits on a path from one boundary to another.




FIG. 4. (a) The standard repetition code. (b) The generalized repetition code.

Just as the surface code can be defined as the tensor product of two repetition codes, the generalized surface code can be defined as the tensor product of two generalized repetition codes [10], as shown in Fig. 5. What we obtain is a collection of squares of the surface code that are joined together along the branches of the two repetition codes. We will refer to these joining branches of the generalized surface code as *seams*, and to the squares of the surface code (without the seams) as *squares*. We say that a generalized surface code has length L if the two generalized repetition codes have length L .

Note that the lengths of the generalized repetition code and the generalized surface code are defined in such a way that the length of these codes in the subdivision procedure is given by the subdivision parameter L .

F. Structure of the code

The geometrically local code constructed in [10] is obtained by taking a good qLDPC code which has a balanced product code formulation, and applying a subdivision process to it. In particular, each square in the balanced product code (formed from two qubits and a neighboring X - and Z -check) is filled in by an $L \times L$ square patch of the surface code, Fig. 6.

By taking L sufficiently large, namely $L = V^{\frac{1}{D-2}} \log^{D+1} V$, the code becomes geometrically local in dimension D while retaining its optimal parameters. Here $V = |X(0)| + |X(1)| + |X(2)|$ is the total number of code elements, i.e., the sum of the number of X -checks, the number of qubits, and the number of Z -checks.

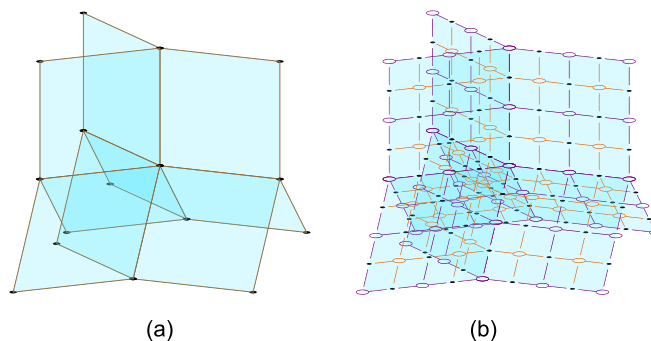


FIG. 5. (a) The Cartesian product of two Y-shaped structures with degree 3 branching. (b) The generalized surface code.

¹The saturation of the BPT bound was demonstrated up to a polylog factor. However, we are aware of an upcoming result by Lin and Portnoy [35] that removes the polylog factor in the embedding of [9] and consequently in [10]. We will therefore omit the polylog factors.




FIG. 6. Filling of a square of the balanced product code with an $L \times L$ square patch of the surface code ($L = 5$). The regions S , T , and U in the subdivided square correspond to the X -check, qubits, and Z -check in the original square, respectively.

Because a balanced product code is locally a product, the resulting code is a collection of generalized surface code patches connected to each other through the T -regions, which correspond to the qubits in the original code. A very simple 2D example is shown in Fig. 1. In general, the patches are generalized surface codes, and their boundary (the T -region) is a generalized repetition code.

Logical errors

Although the subdivided code can be seen as the combination of a good qLDPC outer code and a surface inner code, this can be misleading since it is not a concatenation, and the patches of generalized surface codes do not contain any logical information on their own. This follows from the exactness of the chain complex of the generalized surface code (Corollary 5.5 in Ref. [10]).

Even though there are no nontrivial logical operations within a single patch, a string² of errors crossing many patches is able to form a logical operation. It is therefore important to avoid the creation of such error strings during the decoding process.

Furthermore, since $d = \Theta(L\tilde{d}) = \Theta(L\tilde{n}) = \Theta[L \text{ (no. patches)}]$ [Eq. (6)], a decoder that works up to a constant fraction of the distance must be able to handle a number of errors per patch, which is linear in L .

III. DECODER OF SUBDIVIDED CODE

The strategy to decode the subdivided code is to first push the errors from within each patch to the boundary T -regions, and then to push the errors from the T -regions to the corner U vertices. In this way, the syndrome is fully supported on the U vertices, which allows the use of the decoder of the underlying good qLDPC code [6,36,37]. The reason for this is the one-to-one correspondence of the S , T , and U regions in the subdivided code, and the X -checks, qubits, and Z -checks in the qLDPC code, respectively. The correction predicted by the qLDPC decoder is therefore applied to the T -regions of the corresponding erroneous qubits.

The full decoder of the subdivided code consists of three sub-decoders, each acting on a different region:

- (1) A Union-Find decoder for the generalized surface code, acting on the S regions.
- (2) A minimum weight perfect matching decoder for the generalized repetition code, acting on the T regions.
- (3) The decoder of the underlying good quantum LDPC code, acting on the U regions.

An overview of the decoder of the subdivided code is given in Algorithm I. The serial decoder of the subdivided code runs in almost linear time and can correct errors up to a constant fraction of the distance.

Theorem III.1. The decoder for the subdivided code, as described in Algorithm I, runs in almost linear time $O(n\alpha(n))$ and can correct an X -error of weight s if $s \leq \frac{(L-1)r}{4}$.

Remark III.1. A simple parallel version of the subdivided decoder (for adversarial noise) is obtained by doing the following:

- (1) Applying the Union-Find decoder to each patch of the generalized surface code in parallel: $O(n^{2/D})$.
- (2) Applying the MWPM decoder to each T -region in parallel: $O(n^{1/D})$.
- (3) Applying the parallel version of qLDPC decoder [38,39]: $O(\log(n))$.

Hence the time-complexity of this parallel decoder is $O(n^{2/D})$. Moreover, since there is no known decoder for the surface code with a sublinear worst-case time-complexity for adversarial noise, this appears to be the most one can achieve.

IV. MINIMUM WEIGHT PERFECT MATCHING OF GENERALIZED REPETITION CODE

In this section, we describe the minimum weight perfect matching decoder that is used on the T -regions of the subdivided code in Algorithm I.

Since the generalized repetition code has a 1D structure, the state of a syndrome bit is fully determined by the error on the qubit to its left and the error on the qubit to its right. Fixing the error on the central qubit, therefore, fixes the error on all qubits. To obtain the minimum weight perfect matching, we simply select the configuration (error or no error on the central qubit) with the minimum number of errors.

²A string must be interpreted here as a 1D structure that may include branching points or loops.

ALGORITHM I. Decoder for subdivided code.

-
-
- Input:** The syndrome σ of an error E_X .
Output: An estimation \mathcal{C} of E_X up to a stabilizer.
- 1: **for** every patch of generalized surface code, \mathcal{S}_i **do**
 - 2: Apply the Union-Find generalized surface code decoder, as given in Algorithm V.
 - 3: **for** every boundary T -region, \mathcal{T}_i **do**
 - 4: Apply the minimum weight perfect matching decoder for the generalized repetition code, as given in Algorithm II.
 - 5: Use the decoder of the underlying good qLDPC code to the U vertices, and apply the correction to the corresponding T -regions.
-
-

Algorithm II describes this procedure.

Theorem IV.1. The minimum weight perfect matching decoder for the generalized repetition code, as described in Algorithm II, runs in linear time and can correct an X -error of weight s if $s \leq \frac{(L-1)\Delta}{4}$, with Δ the degree of the branching point.

Proof. It is clear that Algorithm II runs in linear time since we only need to run over the (hyper)edges and vertices once. Since the errors on both sides of a vertex differ if and only if the vertex belongs to the syndrome, the algorithm produces an error with the given syndrome σ . Because the placement of an error on the central (hyper)edge determines the error on all the other edges, there are only two possible error configurations. Since these configurations differ on each (hyper)edge, their weights add up to the total number of (hyper)edges, which is also the distance of the code: $d = \frac{(L-1)\Delta}{2} + 1$. Finally, because we pick the error configuration with the smallest weight, the decoding works for X -errors with a weight of up to $\frac{d-1}{2} = \frac{(L-1)\Delta}{4}$. ■

V. GENERALIZED ERASURE DECODER

In this section we present an adaptation of the erasure decoder [29] such that it can be applied to the generalized

ALGORITHM II. Minimum weight perfect matching decoder of the generalized repetition code.

- Input:** A generalized repetition code with a branching point of degree Δ , and the syndrome σ of an error E_X .
Output: A minimum weight error \mathcal{C} producing the syndrome σ .
- 1: Place an X -error on the central (hyper)edge.
 - 2: **for** all arms of the code **do**
 - 3: **for** all edges i (i increases away from the center) **do**
 - 4: **if** edge i is preceded by a syndrome vertex **then**
 - 5: Place the same error on i as on $i - 1$
 - 6: **else**
 - 7: Place the same error on i as on $i - 1$ plus an additional X -error
 - 8: **if** the number of placed errors is \leq half the number of (hyper)edges **then**
 - 9: **return** the placed errors.
 - 10: **else**
 - 11: **return** the placed errors with an additional X -error on each (hyper)edge.
-
-

surface code. This generalized erasure decoder will be needed in Sec. VI to construct a Union-Find decoder for the generalized surface code. As is the case for the regular Union-Find decoder [13], the Union-Find decoder for the generalized surface code consists of two parts. First the error is reduced to an erasure error, and second, the erasure error is decoded. In this section, we will thus focus on the second part.

To apply the erasure decoder to the generalized surface code, few adaptations need to be made. First, the cluster erasure must be split into subclusters, each of which is only contained within a single square of the generalized surface code. In this way, the regular surface code erasure decoder with boundary may be applied to each subcluster. However, this only works if after splitting, each subcluster is *satisfiable*, that is, it supports an error that is able to produce the syndrome on the subcluster.

To achieve satisfiable subclusters, we do the following. First, we determine for each square s that supports some part of the cluster whether or not the restriction of the cluster to that square $C|_s$ is connected to the boundary. If it is connected to the boundary, $C|_s$ forms a satisfiable subcluster. If it is not connected to the boundary, $C|_s$ is satisfiable if and only if the parity of the syndrome on $C|_s$ is even.

To make the subclusters with an odd parity satisfiable, we give them an additional syndrome vertex by placing errors on some of the (hyper)edges on the seams of the generalized surface code. Note that we have to place at most one error on each seam to achieve the change in parity. We must also be careful not to alter the parity of the even subclusters (not connected to the boundary) during this process.

We are therefore looking for a solution x of the equation

$$\sum_i x_i G_{is} = p_s \quad (7)$$

for all squares s that are not connected to the boundary. The components of Eq. (7) are defined as follows:

(i) x is a row vector in \mathbb{F}_2 with $x_i = 1$ ($x_i = 0$) indicating that seam i has an error (no error).

(ii) G is a matrix in \mathbb{F}_2 with $G_{is} = 1$ if and only if seam i neighbors square s .

(iii) p is a row vector in \mathbb{F}_2 with $p_s = 1$ if and only if the parity of the syndrome on $C|_s$ is odd.

Because the errors on the (hyper)edges on the seams need to be contained within the cluster erasure, we should only look for a solution x that is supported on the seams with some overlap with the cluster erasure. Because the original cluster erasure is by definition satisfiable, there will always be a solution x of Eq. (7) that is contained within the cluster erasure. We refer to a set of (hyper)edges that implement x as a *satisfying configuration*.

Splitting the cluster erasure into subclusters is then done by removing all the (hyper)edges on the seams and flipping the vertices of the (hyper)edges of the satisfying configuration in σ . Finally, the error of the cluster erasure is obtained by taking the union of the errors on the (hyper)edges of the satisfying configuration and the errors we get by applying the regular erasure decoder to each square.

An overview of the full erasure decoder for the generalized surface code is presented in Algorithm III (see Appendix A

ALGORITHM III. Erasure decoder for generalized surface code.

Input: A generalized surface code with boundary, a cluster erasure $\varepsilon \subset \hat{E}$, and the syndrome $\sigma \subset \hat{V}$ of an X error.

Output: An X -error P such that $P \subset \varepsilon$ and $\sigma(P) = \sigma$.

- 1: Initialize A by $A = \emptyset$.
- 2: Find a satisfying configuration x of the cluster.
- 3: Add the satisfying hyperedges to A and flip their vertices in σ .
- 4: **for** each square s of the cluster **do**
- 5: **if** $C|_s$ is connected to the boundary **then**
- 6: Grow a spanning tree $F_{\varepsilon|_s}$ in s from a vertex on the boundary.
- 7: **if** $C|_s$ is not connected to the boundary **then**
- 8: Grow a spanning tree $F_{\varepsilon|_s}$ from any vertex.
- 9: **while** $F_{\varepsilon|_s} \neq \emptyset$ **do**
- 10: pick a leaf edge $e = \{u, v\}$ with pendant vertex $u \in \hat{V}$,
remove e from $F_{\varepsilon|_s}$
- 11: **if** $u \in \sigma$ **then**
- 12: add e to A , remove u from σ , and flip v in σ .
- 13: **if** $u \notin \sigma$ **then**
- 14: do nothing.
- 15: **Return** $P = \prod_{e \in A} Z_e$.

for the relevant definitions), and Fig. 7 illustrates the different steps of the algorithm. Finally, Theorem V.1 shows that the algorithm decodes correctly and that it runs in linear time.

Theorem V.1. The erasure decoder for the generalized surface code, as described in Algorithm III, provides for a given cluster erasure ε and syndrome σ an X -error P such that $P \subset \varepsilon$ and $\sigma(P) = \sigma$. Moreover, the decoder runs in linear time.

Proof. We show that the set A satisfies the following properties: $A \subset \varepsilon$ and $\partial(A) = \sigma$, where $\partial(A)$ denotes the set of vertices that A encounters an odd number of times. The statement $A \subset \varepsilon$ is immediate since we only add elements of ε to A .

To show that $\partial(A) = \sigma$, first note that removing all hyperedges, except those in x , does not change the satisfiability of the cluster. Furthermore, removing the hyperedges in x while flipping their vertices in σ also does not change the satisfiability of the cluster. Call σ' the modified syndrome, i.e., $\sigma'(v) = \sigma(v)$ for every vertex v not neighboring an edge in x , and $\sigma'(v) = [\sigma(v) + j] \bmod 2$ for every vertex v neighboring j hyperedges in x .

Since each connected component of the cluster is now fully contained within a square s , the satisfiability of the cluster implies that for each square s , $C|_s$ is satisfiable. The regular erasure decoder [29] can therefore be applied to each subcluster $C|_s$ with syndrome $\sigma'|_s$, which produces $A'|_s$ with $A'|_s \subset \varepsilon|_s$ and $\partial(A'|_s) = \sigma'|_s$.

Because $A = \bigcup_{e \in x} \{e\} \cup \bigcup_s A'|_s$, its boundary evaluated in a vertex v is given by

$$\partial(A)(v) = \left[\partial\left(\bigcup_{e \in x} \{e\}\right)(v) + \sum_s \partial(A'|_s)(v) \right] \bmod 2. \quad (8)$$

If $v \in \tilde{s}$ does not neighbor any hyperedge in x , we find $\partial(A)(v) = \partial(A'|_{\tilde{s}})(v) = \sigma'|_{\tilde{s}}(v) = \sigma(v)$. If $v \in \tilde{s}$ neighbors some hyperedges $\tilde{e}_1, \dots, \tilde{e}_j$, we find $\partial(A)(v) = [\partial(\bigcup_{e \in \{\tilde{e}_1, \dots, \tilde{e}_j\}} \{e\})(v) + \partial(A'|_{\tilde{s}})(v)] \bmod 2 = [j + \sigma'|_{\tilde{s}}(v)] \bmod 2 = [j + \sigma(v) + j] \bmod 2 = \sigma(v)$. A therefore satisfies the required properties.

We now show that the algorithm runs in linear time. To find a configuration x that satisfies the cluster, we first check every vertex and (hyper)edge to find the squares included in the cluster, the parity of each square, and whether or not a square is connected to the boundary. Finding x from this can then be done in constant time since the number of squares is bounded. Steps (1–3) can therefore be performed in linear time. The rest of the algorithm consists of applying the regular erasure decoder—which runs in linear time—to a bounded number of squares. The overall complexity is therefore linear. ■

VI. GENERALIZED UNION-FIND DECODER

In this section, we present a generalized version of the Union-Find decoder [13] that can be applied to the generalized surface code. In the overall decoder of the subdivided code (Sec. III), this generalized Union-Find decoder is applied to every patch of generalized surface code.

To apply the Union-Find decoder to the generalized surface code, some modifications need to be made. As with the generalized erasure decoder, we must upgrade the notion of an odd cluster to an unsatisfiable cluster, i.e., a cluster for which Eq. (7) does not have a solution. The growing and fusing of edges must also be generalized to include hyperedges. Because we grow a (hyper)edge from a particular vertex towards all the other vertices of the (hyper)edge, growing twice by half a (hyper)edge adds the full (hyper)edge. Similarly, when two half-grown (hyper)edges meet, the full (hyper)edge gets added.

Algorithm IV gives an overview of the Union-Find decoder for the generalized surface code. Note that the output of the algorithm is an estimation of the error up to a stabilizer, but excluding stabilizers that turn a non-patch-crossing path into a patch-crossing path. This difference with the Union-Find decoder for the surface code is due to the fact that the generalized surface code on the S region does not contain any logical information (Sec. II F 1). Although the logical state of the generalized surface cannot change by crossing the patch, it can affect the logical state of the subdivided code. The output of Algorithm IV is therefore exactly what is needed.

Theorem VI.1 is the analog of Theorem B.1; it shows that the generalized Union-Find decoder can correct t erased qubits and s Pauli errors if $t + 2s < L$.

Theorem VI.1. The Union-Find decoder for the generalized surface code, as described in Algorithm IV, can correct any combination of t erased qubits and s X -errors if $t + 2s < L$.

Proof. With every growing step, an additional half-(hyper)edge of the error is covered by the clusters. This means that there can be at most $2s$ growing steps, leading to a combined diameter growth of at most $2s$ (hyper)edges. In case a boundary is hit, the growth of the clusters can stop early, say after \tilde{s} steps, without the clusters covering the error E_X . The

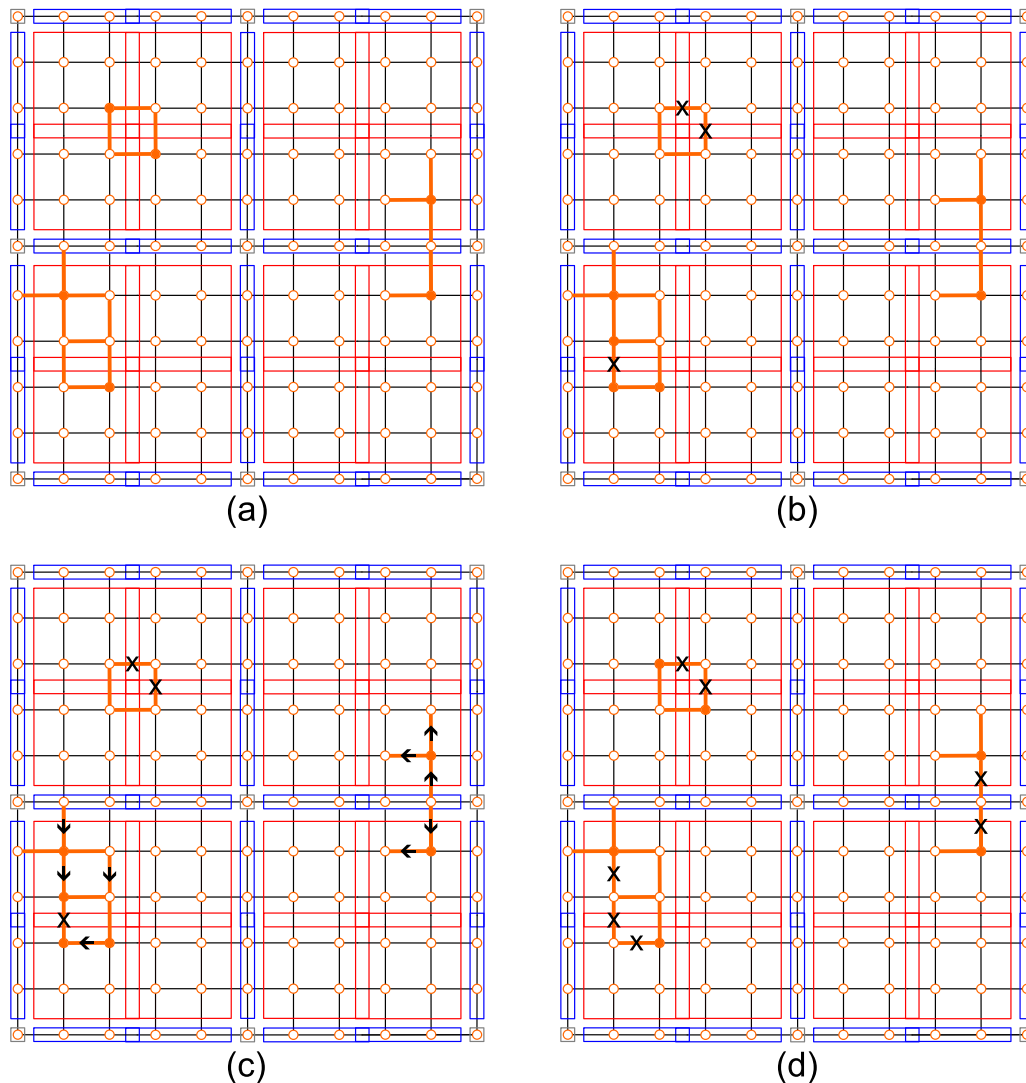


FIG. 7. Generalized erasure decoder applied to four patches of the generalized surface code. The algorithm runs from a to d. (a) The erasure (hyper)edges and the syndrome vertices are marked in orange. (b) X -errors are placed on the satisfying (hyper)edges and their vertices are flipped in σ . (c) Spanning trees are constructed in each square. The arrows indicate the growing direction of the tree. (d) The trees are peeled away from their leaves and the X -error is constructed. The original syndrome vertices are again indicated in orange to show that the estimated X -error produces the syndrome.

largest possible diameter³ of an error string (E_X and estimated correction) is then $2\tilde{s} + (s - \tilde{s}) = s + \tilde{s} \leq 2s$. Adding the initially t erased qubits bounds the diameter of any error string by $2s + t < L$. This is not enough to cross the patch. ■

We now show that the almost-linear time implementation of the Union-Find decoder can also be adapted to the generalized surface code without increasing its time-complexity.

Data structures

The data structures in Algorithm V are largely the same as those in [13]. However, a few modifications are needed

to keep track of the additional structure of the clusters in the generalized surface code.

a. Cluster-tree. A cluster-tree is a tree whose vertices correspond to the vertices of a cluster C . It is initialized by taking an arbitrary vertex u of C as the root of the tree, and adding all other vertices of C at depth 1.

The root stores for each square s :

- (i) the *square inclusion*: does s contain vertices of C
- (ii) the *parity* of $\sigma|_{s \cap C}$
- (iii) the *boundary status*: does C connect to the boundary in s .

The root also stores for each seam f :

- (i) the *seam inclusion*: if f contains hyperedge(s) of C , the seam inclusion is any such hyperedge, otherwise the seam inclusion of f is the empty set.

b. Support. Support is a look-up table that stores for each (hyper)edge its *filled status* and *boundary status*. The filled

³The diameter of an error is the minimum sum of diameters of clusters needed to cover the error.

ALGORITHM IV. Union-Find generalized surface code decoder - Naive version.

Input: The set of erased positions $\varepsilon \subset E$ and the syndrome σ of an error E_X .

Output: An estimation \mathcal{C} of E_X up to a stabilizer.^a

- 1: Create the list of all unsatisfiable clusters C_1, \dots, C_m , and initialize the modified erasure $\varepsilon' = \varepsilon$.
- 2: **while** there exists an unsatisfiable cluster **do**
- 3: **for** all unsatisfiable clusters C_i **do**
- 4: Grow C_i by increasing its radius by one half-(hyper)edge.
- 5: **if** C_i meets another cluster **then**
- 6: fuse and update satisfiability.
- 7: **if** C_i meets the boundary **then**
- 8: update satisfiability.
- 9: **if** C_i is satisfiable **then**
- 10: remove it from the unsatisfiable cluster list.
- 11: Add full (hyper)edges that are in the grown clusters to ε' .
- 12: **for** all satisfiable clusters **do**
- 13: apply the erasure decoder for the generalized surface code to find \mathcal{C} .

^aExcluding stabilizers that turn a non-patch-crossing path into a patch-crossing path.

status can take one of three values: *Unoccupied*, *Half-grown*, or *Grown*. While the filled status is updated during the algorithm, the boundary status of a (hyper)edge is fixed and simply indicates whether it connects to the boundary of the patch of the generalized surface code (a *boundary edge*) or not [an *interior (hyper)edge*].

Support is initialized by marking all the erased (hyper)edges as *Grown* and marking all others as *Unoccupied*. Each (hyper)edge is also given its appropriate boundary status.

c. Boundary list. The Boundary List of a cluster C is a list of all the boundary vertices of C , where a *boundary vertex* of C is a vertex of C with at least one of its incident (hyper)edges not in the erasure ε .

Theorem VI.2. The implementation of the generalized Union-Find decoder, as described in Algorithm V, runs in almost linear time.

Proof. Algorithm V can be divided into three blocks. The first block consists of lines 1–2 and initializes the clusters. The second block consists of lines 3–28 and performs the growing and fusion of clusters. Finally, the third block consists of line 29 and applies the erasure decoder for the generalized surface code.

The creation of the cluster-trees can be achieved in linear time by applying the connected components algorithm from [40] to find the connected components of the sublattice induced by the erased (hyper)edges. This algorithm can directly be applied if we replace an l -hyperedge, discovered by the algorithm from vertex u , by l edges emanating from u , i.e., $\{u, v_1, \dots, v_l\} \rightarrow \{u, v_1\}, \dots, \{u, v_l\}$. During this exploration of the connected components of the erasure, we can also calculate the square inclusion, the parity, and boundary status of the cluster at a constant overhead.

ALGORITHM V. Union-Find generalized surface code decoder-Almost linear time version.

Input: The set of erased positions $\varepsilon \subset E$ and the syndrome σ of an error E_X .

Output: An estimation \mathcal{C} of E_X up to a stabilizer.^a

- 1: Initialize cluster-trees, Support and boundary lists for all clusters.
- 2: Create the list \mathcal{L} of roots of unsatisfiable clusters.
- 3: **while** $\mathcal{L} \neq \emptyset$ **do**
- 4: Initialize fusion list $\mathcal{F} = \emptyset$
- 5: **for** all $u \in \mathcal{L}$ **do**
- 6: Grow the cluster C_u a half (hyper)edge in the Table Support.
- 7: **if** a new grown boundary edge e is added in Support **then**
- 8: Update boundary at root u .
- 9: **if** a new grown interior (hyper)edge e is added in Support **then**
- 10: Add e to the fusion list \mathcal{F} .
- 11: **for** all $e = \{v_1, \dots, v_l\} \in \mathcal{F}$ **do**
- 12: Determine $\text{Find}(v_i)$.
- 13: **if** there are at least two distinct roots **then**
- 14: Read the sizes at the distinct roots.
- 15: Union the smaller clusters to the biggest cluster.
- 16: Append the boundary lists of the smaller clusters at the end of the boundary list of the biggest cluster.
- 17: **if** all roots are identical **then**
- 18: Remove e from the list \mathcal{F} .
- 19: **for** each $u \in \mathcal{L}$ **do**
- 20: **if** $\text{Find}(u) \neq u$ **then**
- 21: remove u from \mathcal{L} .
- 22: **for** all $u \in \mathcal{L}$ **do**
- 23: Remove the vertices of the boundary lists of u that are not boundary vertices.
- 24: **for** all $u \in \mathcal{L}$ **do**
- 25: Calculate the satisfiability of C_u .
- 26: **if** C_u is satisfiable **then**
- 27: Remove u from \mathcal{L} .
- 28: Erase all the (hyper)edges that are fully grown in Support.
- 29: Apply the erasure decoder for the generalized surface code.

^aExcluding stabilizers that turn a non-patch-crossing path into a patch-crossing path.

The remaining clusters are those that consist of a single syndrome vertex. The creation of the cluster-tree for each of these types of clusters is achieved in constant time. The overall time-complexity of the initialization of the cluster-trees is thus linear.

Support and the Boundary Lists can be initialized in linear time since we only have to check each (hyper)edge and vertex once, and each check takes a constant time.

The list \mathcal{L} of unsatisfiable clusters is created by calculating for each cluster its satisfiability. Since the number of squares in the generalized surface code is bounded, the satisfiability of a cluster can be calculated from its square inclusion, parity, and boundary status in a constant time. Since the number of clusters is in the worst case linear in the number of qubits, the time-complexity of the creation of \mathcal{L} is also linear.

Block 1 of the algorithm can therefore be performed in linear time.

Growing the clusters (lines 6–10) is done by running over the vertices in the boundary lists, and for each vertex growing its bounded number of adjacent (hyper)edges. Because each vertex can be a boundary vertex for at most two rounds of growth, this takes at most linear time.

Lines 11–18 mainly describe the Union-Find procedure, which is known to have a time complexity of $O(n\alpha(n))$. The only difference is that we are now also appending the boundary lists of the smaller clusters at the end of the boundary list of the biggest cluster. However, since each vertex can be a boundary vertex for at most two rounds of growth, this does not change the complexity.

The linearity of lines 19–27 is again due to the fact that a vertex is a boundary vertex for at most two rounds of growth, and that a root $u \in \mathcal{L}$ has a nonempty boundary list after each round of growth. This may need some more explanation: The only way a root $u \in \mathcal{L}$ can have an empty boundary list is if C_u covers the full patch of the generalized surface code, such that no further growth is possible. However, such a cluster will be satisfiable and u will therefore be removed from \mathcal{L} . After each round of growth, each root $u \in \mathcal{L}$ therefore has a nonempty boundary list. This also means that (after each round of growth) the number of roots is smaller than, or equal to, the number of boundary vertices.

Because a vertex can be a boundary vertex for at most two rounds of growth, the number of roots u that have to be checked in the full algorithm is smaller than three times the number of vertices. (The factor of 3 accounts for the case in which a cluster loses its last boundary vertex during growth.)

Evidently, adding the fully grown (hyper)edges to the erasure (line 28) is done in linear time, which shows that block 2 of the algorithm can be performed in almost linear time $O(n\alpha(n))$.

By Theorem V.1, the erasure decoder for the generalized surface code runs in linear time.

The complete algorithm has therefore an almost linear time-complexity $O(n\alpha(n))$. ■

VII. PROOF OF THE MAIN THEOREM

Proof of Theorem III.1. Since the algorithm consists of applying three consecutive decoders, which all run in almost linear time or linear time, the overall time-complexity is also almost linear.

We now show that the decoder works for error weights up to $\frac{(L-1)r}{4}$, where r is the maximum number of errors that the decoder of the underlying good qLDPC code is guaranteed to correct.

There are two main reasons why the decoder may fail:

- (1) The minimum weight perfect matching decoder for the generalized repetition code fails.
- (2) The Union-Find generalized surface code decoder fails.

We first consider the failure of the minimum weight perfect matching decoder for the generalized repetition code.

Note that the largest diameter of the total error (initial error plus decoding error) that can be achieved by applying the Union-Find generalized surface code decoder to an initial error of weight s is at most $2s$ (see the proof of Theorem VI.1). Let B be a set of clusters (with minimal total diameter) cov-

ering the total error. Then, assuming that $s < \frac{L}{2}$, no cluster in B can neighbor more than two T -regions. In this case, the diameter of a cluster forms an upper-bound for the number of errors that can be pushed from that cluster to the T -regions.

Therefore, to get more than $\frac{L-1}{2}$ errors in a T -region, more than $\frac{L-1}{4}$ errors are needed from the patches. Of course, the T -region itself may contain initial errors as well. However, since initial errors in the T -region count as singular, and those in the patches count as double, the worst-case scenario is where all the initial errors are in the patches.

Now, if the total number of initial errors in the subdivided code is at most $\frac{(L-1)r}{4}$, there are fewer than r T -regions with more than $\frac{L-1}{2}$ errors, and they may not be decoded correctly. This leads to fewer than r errors on the good qLDPC decoding level, which is therefore guaranteed to succeed.

Finally, we consider the failure of the Union-Find generalized surface code decoder.

By Theorem VI.1, the Union-Find generalized surface code decoder can only fail when a patch contains more than $\frac{L-1}{2}$ initial errors. However, if the total number of initial errors in the subdivided code is at most $\frac{(L-1)r}{4}$, there are fewer than $\frac{r}{2}$ patches with more than $\frac{L-1}{2}$ initial errors, and they may not be decoded correctly. This leads to fewer than $\frac{r}{2}$ errors on the good qLDPC decoding level, which is therefore guaranteed to succeed. ■

VIII. THRESHOLD

In this section, we consider the effect of random noise. More specifically, we look at the code under the code capacity noise model.

A. Threshold of the good qLDPC code

Given a good qLDPC code with a decoder that decodes correctly up to r errors, the probability of a logical error is given by

$$P_L \leq P(\text{no. errors} \geq r + 1) \quad (9a)$$

$$= \sum_{i=r+1}^{n_{\text{qLDPC}}} \binom{n_{\text{qLDPC}}}{i} p^i (1-p)^{n_{\text{qLDPC}}-i} \quad (9b)$$

$$= F(n_{\text{qLDPC}} - (r + 1); n_{\text{qLDPC}}, 1 - p) \quad (9c)$$

with $F(r + 1; n_{\text{qLDPC}}, p)$ the cumulative distribution function of the binomial distribution. By taking $n_{\text{qLDPC}} p \leq r + 1$, Hoeffding's inequality yields the bound

$$P_L \leq \exp \left[-2n_{\text{qLDPC}} \left(\frac{r+1}{n_{\text{qLDPC}}} - p \right)^2 \right]. \quad (10)$$

This upper bound can be made arbitrarily small by increasing $d \propto n_{\text{qLDPC}}$, provided $p \leq \frac{r+1}{n_{\text{qLDPC}}}$.

This shows that a good qLDPC code always has a threshold under the code capacity noise model.

B. Threshold of the generalized surface code and generalized repetition code (MWPM+MWPM)

We now look at the “threshold” of the combined generalized surface code and generalized repetition code, using

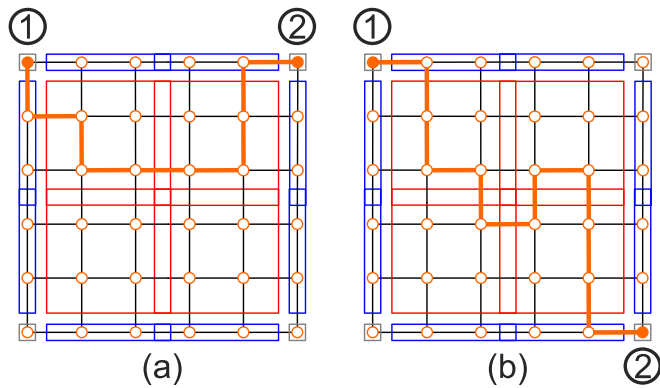


FIG. 8. “Logical” errors of the combined generalized surface code and generalized repetition code. (a) A path connecting the U -vertices at the boundary of a single (top) T -region. (b) A path connecting the U -vertices at the boundary of two (top and right, or, bottom and left) T -regions.

minimum weight perfect matching on both the generalized surface code and the generalized repetition code. Since these codes do not contain any logical information themselves, it is technically impossible to have a logical error. However, since the T -regions correspond to the qubits in the good qLDPC code (see Fig. 6), we can identify the logical errors with paths that connect the U -vertices at the boundary of one or multiple T -regions. Two examples of such paths are shown in Fig. 8.

Consider such a path between U -vertices. This path can consist of both initial errors and decoding errors, and it can have sections in T -regions as well as in S -regions. Observe that while a segment of the path within a T -region can be entirely composed of decoding errors, a segment within an S -region can contain decoding errors in at most half of its (hyper)edges (MWPM in S). The ratio of decoding errors in the path is therefore maximized by maximizing the number of (hyper)edges of the path within T -regions.

Since MWPM is also performed on the T -regions, each T -region can contain decoding errors in at most half of its (hyper)edges. To connect to the next T -region (or simply to the U -vertices at the end points of the path), the remaining errors must be allocated to the S -regions. Consequently, the length of the path within S must be at least equal to the length within T . Therefore, the maximum proportion of decoding errors in the entire path is $\leq \frac{1}{2} + \frac{1}{4} = \frac{3}{4}$.

Focusing on a logical error associated with a single T -region, there must be at least one path of $m \geq \frac{(L-1)\Delta}{2} + 1$ (hyper)edges connecting the corner U -vertices adjacent to this T -region, and containing at least $\lceil m/4 \rceil$ (hyper)edges associated with errors.

Since the end points of the path are fixed, we have to make at most m choices for the next (hyper)edges. Because each vertex is connected to at most Δ (hyper)edges and because we do not allow backtracking, the total number of paths of length m (with fixed end points) is upper-bounded by $(\Delta - 1)^m$.

Following [41], we find that for a path of length m , the probability of at least $\lceil m/4 \rceil$ of its (hyper)edges being associated with errors is at most $2^m \epsilon^{\lceil m/4 \rceil}$, with ϵ being the error probability at a single (hyper)edge.

The probability of having a “logical error” associated with a single T -region is therefore

$$P_L \leq \sum_{m=\frac{(L-1)\Delta}{2}+1}^{\infty} a^m 2^m \epsilon^{\lceil m/4 \rceil} \quad (11a)$$

$$\leq \sum_{m=\frac{(L-1)\Delta}{2}+1}^{\infty} (2a^4 \sqrt{\epsilon})^m \quad (11b)$$

$$\leq (2a^4 \sqrt{\epsilon})^{\frac{(L-1)\Delta}{2}+1} \sum_{m=0}^{\infty} (2a^4 \sqrt{\epsilon})^m \quad (11c)$$

$$= (2a^4 \sqrt{\epsilon})^{\frac{(L-1)\Delta}{2}+1} \frac{1}{1 - 2a^4 \sqrt{\epsilon}}. \quad (11d)$$

Provided that $\epsilon < 1/(2a^4)$, the probability can be made arbitrarily small by increasing n and, consequently, L . This establishes the existence of a threshold for the generalized surface and repetition code. When combined with the threshold of the outer good qLDPC code, this further demonstrates the existence of a threshold for the total subdivided code.

We have therefore proven the following.

Theorem VIII.1. Geometrically local codes from subdivision, together with a combined minimum weight perfect matching and good qLDPC decoder, have a finite threshold error rate. Moreover, logical errors are suppressed exponentially with the subdivision parameter L .

C. Threshold of the generalized surface code and generalized repetition code (generalized UF+MWPM)

In this section, we discuss the “threshold” of the combined generalized surface code and generalized repetition code, using the generalized Union-Find decoder and the minimum weight perfect matching decoder, respectively.

While a formal proof of the existence of a threshold is beyond the scope of this work, we offer an intuitive argument to suggest why this might hold true, noting that a similar level of rigor is also absent for the union-find decoder of the regular surface code.

One argument that can be made is the following. Looking back at the proof presented in Sec. VIII B, we see that the minimum weight perfect matching on the T -regions essentially halves the length required in the S -regions to create a logical error. Alternatively, we can interpret this by maintaining the same path length in the S -regions (treating the contribution of the T -region as part of S) and rescaling the number of initial errors needed to construct the path. Since the T -region can take up to half the (hyper)edges of the path, the probability of a particular path increases by up to the square root of the original probability, i.e., $\mathbb{P}(\text{path}) = p \rightarrow \mathbb{P}(\text{path}) = \sqrt{p}$.

Given that numerical evidence suggests that the surface code exhibits similar thresholds for MWPM and the Union-Find decoder [13], it is reasonable to infer that this square root reduction in path probability also applies to the Union-Find decoder.

Finally, because the generalized Union-Find decoder can be understood as applying the Union-Find decoder to each square s , the numerical threshold of the Union-Find decoder for the surface code [13], combined with the square root




FIG. 9. Gridlike representation.

probability reduction (which correspondingly rescales the threshold by the square root), leads to a threshold for the generalized Union-Find and MWPM decoders.

ACKNOWLEDGMENT

Q.E. acknowledges the support of the Research Foundation-Flanders through the Fundamental Research PhD programme (Grant No. 11Q4A24N), as well as the EOS-FWO-FNRS project CHEQS.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX A: ERASURE DECODER FOR THE SURFACE CODE

The goal of the erasure decoder for the surface code is to find a Z -error and an X -error supported on the erasure ε that reproduce the given syndromes σ_z and σ_x . As stated in Sec. II B 3, this corresponds to the most likely correction. From now on, and for the rest of the paper, we will focus only on the X -errors since the Z -errors can be handled in a completely analogous way. The syndrome σ should therefore always be understood as the X -part of the syndrome, σ_x .

To describe the algorithm, it is best to use a gridlike representation of the surface code whereby the vertices V correspond to the Z checks, the faces correspond to the X checks, the edges E correspond to qubits, and two vertices are connected by an edge if there is a qubit that is acted on by the corresponding checks. An example of such a gridlike representation is shown in Fig. 9.

The strategy [29] is to translate the decoding problem (how to pair up the syndrome vertices through paths in the erasure) into a graphical language. To this end, the *boundary* of a set of edges A , denoted as ∂A , is defined as the set of vertices that A encounters an odd number of times. Decoding is therefore equivalent to finding a set A such that $A \subset \varepsilon$ and $\partial(A) = \sigma$.

Finding such a set A in linear-time is made difficult by the many cycles in ε . To avoid this issue, a *spanning forest* F_ε is created inside ε , which is defined as the maximal subset of edges of ε that contains no cycles and spans all the vertices of ε . If the erasure ε together with its neighboring vertices

forms a connected graph, F_ε is also connected and is called a *spanning tree*. A spanning forest can always be found in linear time [42].

The next step is to peel away the forest F_ε from its leaves and construct the set A in the process. A *leaf* is defined as an edge $e = \{u, v\}$ that is only connected to the forest by one of its vertices, say v . The other vertex, u , is called the *pendant vertex*.

Starting from the empty set, A is built up by recursively applying the following rules:

- (1) Pick a leaf edge $e = \{u, v\}$ with pendant vertex u and remove e from F_ε .
- (2) If $u \in \sigma$, then add e to A , remove u from σ , and flip v in σ .
- (3) If $u \notin \sigma$, then do nothing.

Flipping a vertex v in σ means that if $v \in \sigma$, v gets removed from σ , and if $v \notin \sigma$, v gets added to σ .

Once the forest F_ε is fully peeled, A is complete and satisfies the requirements $A \subset \varepsilon$ and $\partial(A) = \sigma$. The reason for this is the following: since F_ε spans all the vertices of ε , and F_ε gets fully peeled, each vertex is a pendant vertex at some step in the peeling process. If a vertex $u \in \sigma$, that vertex will have a single incident edge of A (rule 2) such that $u \in \partial A$. Alternatively, if $u \notin \sigma$, u will have no incident edges of A (rule 3) such that $u \notin \partial A$. For a full proof, we refer to [29].

Algorithm VI summarizes the decoding procedure for a general surface code $G = (V, E, F)$ without boundary.

ALGORITHM VI. Erasure decoder for the surface code.

Input: A surface $G = (V, E, F)$, an erasure $\varepsilon \subset E$, and the syndrome $\sigma \subset V$ of an X -error.

Output: An X -error P such that $P \subset \varepsilon$ and $\sigma(P) = \sigma$.

- 1: Initialize A by $A = \emptyset$.
 - 2: Construct a spanning forest F_ε of ε .
 - 3: **while** $F_\varepsilon \neq \emptyset$ **do**
 - 4: Pick a leaf edge $e = \{u, v\}$ with pendant vertex u and remove e from F_ε .
 - 5: **if** $u \in \sigma$ **then**
 - 6: Add e to A , remove u from σ , and flip v in σ .
 - 7: **if** $u \notin \sigma$ **then**
 - 8: Do nothing.
 - 9: Return $P = \prod_{e \in A} X_e$.
-

1. Surfaces with boundaries

To deal with surface codes with boundaries, that is, surface codes whose grid representations have edges belonging to a unique face, some modifications need to be made to the algorithm. The reason for this is that the vertices on the boundary do not contain enough information to reconstruct the error starting from those vertices. It is therefore important to make the distinction between the edges and vertices at the boundary and those not at the boundary.

To be precise, let \check{E} be the set of *open edges*, meaning those edges that belong to a unique face, and let \check{V} be the set of *open vertices*, where a vertex is defined to be open if it neighbors an open edge. The set of nonopen edges is denoted by $\dot{E} := E \setminus \check{E}$ and the set of nonopen vertices is denoted by $\dot{V} := V \setminus \check{V}$.

To prevent the algorithm from getting stuck before the whole forest is peeled away, each connected component of the erasure, referred to as a *cluster erasure*, that contains a vertex $v \in \check{V}$ is grown from v . Furthermore, during the growing process, edges are only added to the tree if they reach new vertices in \check{V} . This results in trees that contain at most one open vertex, located at their root. The whole spanning forest can therefore be peeled away without any problem.

Algorithm VII summarizes the decoding procedure for a cluster erasure of a general surface code $G = (V, E, F)$ with boundary. The full decoder consists of applying this procedure to each cluster erasure.

It is known that Algorithm VII produces the most likely correction and that its time-complexity is linear in the number of qubits.

Theorem A.1. (Theorem 2 of [29]). For general surface codes with bounded degree and faces of bounded size, applying Algorithm VII to the graph and to its dual produces a linear-time maximum-likelihood decoder, i.e., a decoder that finds the most likely correction.

ALGORITHM VII. Erasure decoder for the surface code (with boundary).

Input: A surface $G = (V, E, F)$ with boundary, a cluster erasure $C \subset \varepsilon \subset \dot{E}$, and the syndrome $\sigma \subset \dot{V}$ of an X -error.
Output: An X -error P such that $P \subset \varepsilon$ and $\sigma(P) = \sigma$.

- 1: Initialize A by $A = \emptyset$.
- 2: **if** C is connected to the boundary **then**
- 3: Grow a spanning tree F_C of C from a vertex on the boundary.
- 4: **if** C is not connected to the boundary **then**
- 5: Grow a spanning tree F_C of C from any vertex.
- 6: **while** $F_C \neq \emptyset$ **do**
- 7: Pick a leaf edge $e = \{u, v\}$ with pendant vertex $u \in \check{V}$ and remove e from F_C .
- 8: **if** $u \in \sigma$ **then**
- 9: Add e to A , remove u from σ , and flip v in σ .
- 10: **if** $u \notin \sigma$ **then**
- 11: Do nothing.
- 12: Return $P = \prod_{e \in A} X_e$.

APPENDIX B: UNION-FIND DECODER FOR THE SURFACE CODE

The goal of the Union-Find decoder for the surface code [13] is to find an estimation of the X -error up to a stabilizer, given an erasure ε and a syndrome σ . To achieve this, the decoder works in two stages. The first stage consists of converting the original erasure ε and any possible Pauli errors into a modified erasure ε' . In other words, we look for a set of qubits that supports an error that is consistent with the syndrome. This stage is therefore called *syndrome validation*. The second stage is then to simply apply the erasure decoder for the surface code.

Syndrome validation works by identifying all the clusters of erasures which are “invalid,” and growing them until they become “valid.” Similar to a cluster erasure, a *cluster* is defined to be either a connected component of the erased qubits in the subgraph (V, ε) , or a single vertex outside the erasure supporting a syndrome bit. A cluster is said to be “invalid” or *unsatisfiable* if it does not support an error that produces the syndrome on the cluster. A cluster that does support such an error is called *satisfiable*.

For the surface code (without boundary), a cluster is satisfiable if and only if the cardinality of the syndrome on the cluster is even (Lemma 1 of [13]). Odd clusters must therefore be grown by adding edges until they connect with another odd cluster. The result of the merger of two odd clusters is an even cluster, which is correctable by the erasure decoder.

Algorithm VIII describes this growing and merging procedure.

1. Surfaces with boundaries

For surface codes with boundaries, a cluster is satisfiable if and only if it is an even cluster *or* it connects to the boundary. The reason for this is that any remaining syndrome bit can always be canceled by an error coming from the boundary.

For surface codes with boundaries, we must therefore keep track of the *boundary status* of each cluster, in addition to

ALGORITHM VIII. Union-Find surface code decoder-Naive version.

Input: The set of erased positions $\varepsilon \subset E$ and the syndrome σ of an error E_X .
Output: An estimation C of E_X up to a stabilizer.

- 1: Create the list of all odd clusters C_1, \dots, C_m , and initialize the modified erasure $\varepsilon' = \varepsilon$.
- 2: **while** there exists an odd cluster **do**
- 3: **for** every odd cluster C_i **do**
- 4: Grow C_i by increasing its radius by one half-edge.
- 5: **if** C_i meets another cluster **then**
- 6: fuse and update parity.
- 7: **if** C_i is even **then**
- 8: remove it from the odd cluster list.
- 9: Add full edges that are in the grown clusters to ε' .
- 10: **for** all even clusters **do**
- 11: apply the erasure decoder for the surface code to find C .

ALGORITHM IX. Union-Find surface code (with boundary) decoder - Naive version.

Input: The set of erased positions $\varepsilon \subset \hat{E}$ and the syndrome $\sigma \subset \hat{V}$ of an error E_X .
Output: An estimation \mathcal{C} of E_X up to a stabilizer.

- 1: Create the list of all unsatisfiable clusters C_1, \dots, C_m , and initialize the modified erasure $\varepsilon' = \varepsilon$.
- 2: **while** there exists an unsatisfiable cluster **do**
- 3: **for** all unsatisfiable clusters C_i **do**
- 4: Grow C_i by increasing its radius by one half-edge.
- 5: **if** C_i meets another cluster **then**
- 6: fuse and update satisfiability.
- 7: **if** C_i meets the boundary **then**
- 8: update satisfiability.
- 9: **if** C_i is satisfiable **then**
- 10: remove it from the unsatisfiable cluster list.
- 11: Add full edges that are in the grown clusters to ε' .
- 12: **for** all satisfiable clusters **do**
- 13: apply the erasure decoder for the surface code to find \mathcal{C} .

its *parity*. We must also use the appropriate erasure decoder, capable of dealing with boundaries.

Algorithm IX shows the modified procedure.

It has been shown that below the minimum distance d of the surface code, the Union-Find decoder performs just as well as the most likely error (MLE) decoder.

Theorem B.1 (Theorem 1 of [13]). If $t + 2s < d$, Algorithm IX can correct any combination of t erased qubits and s X -errors.

2. Almost-linear time-complexity

The decoder can be implemented in an almost-linear time $O(n\alpha(n))$, where n is the number of qubits and α is the inverse of Ackermann's function. To achieve this almost-linear time-complexity, a Union-Find data-structure and algorithm [43,44] are utilized, hence the name of the decoder.

Each cluster is represented by a tree, a *cluster tree*, that consists of all the vertices⁴ of the cluster. An arbitrary vertex of the cluster is chosen as the root of the tree and serves as a representative of the cluster. To determine if two vertices belong to the same or to a different cluster, each vertex recursively looks at its parent until the root of the tree is found. The two vertices belong to the same cluster if and only if the roots are identical. Finding the root of a vertex is done with the `Find()` procedure.

The root of the cluster tree also stores the size of the cluster (the number of vertices of the cluster), its parity, and its boundary status. The parity and boundary status are needed to determine the satisfiability of the cluster, while the size is needed to efficiently merge two clusters.

To merge two clusters, first the smallest cluster-tree is added as a subtree of the root of the largest cluster-tree, and then the size, parity, and boundary status of the resulting tree

ALGORITHM X. Union-Find surface code (with boundary) decoder-Almost linear time version.

Input: The set of erased positions $\varepsilon \subset \hat{E}$ and the syndrome $\sigma \subset \hat{V}$ of an error E_X .
Output: An estimation \mathcal{C} of E_X up to a stabilizer.

- 1: Initialize cluster-trees, Support, and boundary lists for all clusters.
- 2: Create the list \mathcal{L} of roots of unsatisfiable clusters.
- 3: **while** $\mathcal{L} \neq \emptyset$ **do**
- 4: Initialize fusion list $\mathcal{F} = \emptyset$
- 5: **for** all $u \in \mathcal{L}$ **do**
- 6: Grow the cluster C_u a half-edge in the Table Support.
- 7: **if** a new grown boundary edge e is added in Support **then**
- 8: Update boundary at root u .
- 9: **if** a new grown interior edge e is added in Support **then**
- 10: Add e to the fusion list \mathcal{F} .
- 11: **for** all $e = \{u, v\} \in \mathcal{F}$ **do**
- 12: Determine `Find(u)` and `Find(v)`.
- 13: **if** `Find(u) \neq Find(v)` **then**
- 14: Read the sizes at the roots.
- 15: Union the smaller cluster to the bigger cluster.
- 16: Append the boundary lists of the smaller cluster to the boundary list of the bigger cluster.
- 17: **if** `Find(u) = Find(v)` **then**
- 18: Remove e from the list \mathcal{F} .
- 19: **for** each $u \in \mathcal{L}$ **do**
- 20: **if** `Find(u) \neq u` **then**
- 21: remove u from \mathcal{L} .
- 22: **for** all $u \in \mathcal{L}$ **do**
- 23: Remove the vertices of the boundary lists of u that are not boundary vertices.
- 24: **for** all $u \in \mathcal{L}$ **do**
- 25: Calculate the satisfiability of C_u .
- 26: **if** C_u is satisfiable **then**
- 27: Remove u from \mathcal{L} .
- 28: Erase all the edges that are fully grown in Support.
- 29: Apply the erasure decoder for the surface code.

are updated. Merging two clusters is done with the `Union()` procedure.

A look-up table, called *Support*, is used to keep track of the status of the edges of the code. An edge is attributed both a *filled status*: indicating whether it is *Unoccupied*, *Half-Grown*, or *Grown*, and a *boundary status*: indicating whether it neighbors an open vertex or not.

To prevent having to update the whole Support table at each round of growth, which would lead to a quadratic time-complexity, a list of boundary vertices for each cluster is maintained. A *boundary vertex* of a cluster is a vertex of the cluster for which at least one of the incident edges is not in the erasure ε . It is clear that an edge can only grow if it is incident to a boundary vertex. Only the filled status of these edges must therefore be updated at each round of growth.

Algorithm X describes the full almost-linear implementation of the Union-Find decoder for a surface code with boundary.

⁴The vertices of a cluster include all the nonopen vertices that neighbor at least one edge of the cluster.

- [1] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
- [2] L. Skoric, D. E. Browne, K. M. Barnes, N. I. Gillespie, and E. T. Campbell, Parallel window decoding enables scalable fault tolerant quantum computation, *Nat. Commun.* **14**, 7040 (2023).
- [3] D. Gottesman, Fault-tolerant quantum computation with constant overhead, [arXiv:1310.2984](https://arxiv.org/abs/1310.2984).
- [4] P. Panteleev and G. Kalachev, Asymptotically good quantum and locally testable classical ldpc codes, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (ACM Press, New York, NY, 2022), pp. 375–388.
- [5] A. Leverrier and G. Zémor, Quantum tanner codes, in *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2022), pp. 872–883.
- [6] I. Dinur, M.-H. Hsieh, T.-C. Lin, and T. Vidick, Good quantum ldpc codes with linear time decoders, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (ACM Press, New York, NY, 2023), pp. 905–918.
- [7] S. Bravyi and B. Terhal, A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes, *New J. Phys.* **11**, 043029 (2009).
- [8] S. Bravyi, D. Poulin, and B. Terhal, Tradeoffs for reliable quantum information storage in 2D systems, *Phys. Rev. Lett.* **104**, 050503 (2010).
- [9] E. Portnoy, Local quantum codes from subdivided manifolds, [arXiv:2303.06755](https://arxiv.org/abs/2303.06755).
- [10] T.-C. Lin, A. Wills, and M.-H. Hsieh, Geometrically local quantum and classical codes from subdivision, [arXiv:2309.16104](https://arxiv.org/abs/2309.16104).
- [11] D. J. Williamson and N. Baspin, Layer codes, [arXiv:2309.16503](https://arxiv.org/abs/2309.16503).
- [12] X. Li, T.-C. Lin, and M.-H. Hsieh, Transform arbitrary good quantum ldpc codes into good geometrically local codes in any dimension, [arXiv:2408.01769](https://arxiv.org/abs/2408.01769).
- [13] N. Delfosse and N. H. Nickerson, Almost-linear time decoding algorithm for topological codes, *Quantum* **5**, 595 (2021).
- [14] B. J. Brown and D. J. Williamson, Parallelized quantum error correction with fracton topological codes, *Phys. Rev. Res.* **2**, 013303 (2020).
- [15] S. Kosen, H.-X. Li, M. Rommel, D. Shiri, C. Warren, L. Grönberg, J. Salonen, T. Abad, J. Biznárová, M. Caputo *et al.*, Building blocks of a flip-chip integrated superconducting quantum processor, *Quantum Sci. Technol.* **7**, 035018 (2022).
- [16] D. Rosenberg, D. Kim, R. Das, D. Yost, S. Gustavsson, D. Hover, P. Krantz, A. Melville, L. Racz, G. O. Samach *et al.*, 3D integrated superconducting qubits, *npj Quantum Inf.* **3**, 42 (2017).
- [17] D.-R. W. Yost, M. E. Schwartz, J. Mallek, D. Rosenberg, C. Stull, J. L. Yoder, G. Calusine, M. Cook, R. Das, A. L. Day *et al.*, Solid-state qubits integrated with superconducting through-silicon vias, *npj Quantum Inf.* **6**, 59 (2020).
- [18] S. Zhang, Z. Li, H. Zhou, R. Li, S. Wang, K.-W. Paik, and P. He, Challenges and recent perspectives of 3d heterogeneous integration, *e-Prime* **2**, 100052 (2022).
- [19] Z. Cai, A. Siegel, and S. Benjamin, Looped pipelines enabling effective 3D qubit lattices in a strictly 2D device, *PRX Quantum* **4**, 020345 (2023).
- [20] F. van Riggelen-Doelman, C.-A. Wang, S. L. de Snoo, W. I. L. Lawrie, N. W. Hendrickx, M. Rimbach-Russ, A. Sammak, G. Scappucci, C. Déprez, and M. Veldhorst, Coherent spin qubit shuttling through germanium quantum dots, *Nat. Commun.* **15**, 5716 (2024).
- [21] M. De Smet, Y. Matsumoto, A.-M. J. Zwerver, L. Trypuzen, S. L. de Snoo, S. V. Amitonov, A. Sammak, N. Samkharadze, Ö. Gül, R. N. M. Wasserman *et al.*, High-fidelity single-spin shuttling in silicon, [arXiv:2406.07267](https://arxiv.org/abs/2406.07267).
- [22] M. B. Hastings, J. Haah, and R. O’Donnell, Fiber bundle codes: breaking the $n^{1/2}$ polylog (n) barrier for quantum ldpc codes, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (ACM Press, New York, NY, 2021), pp. 1276–1288.
- [23] A. G. Fowler, Minimum weight perfect matching of fault-tolerant topological quantum error correction in average $O(1)$ parallel time, *Quantum Inf. Comput.* **15**, 145 (2015).
- [24] Y. Wu and L. Zhong, Fusion blossom: Fast mwpm decoders for qec, in *Proceedings of the 2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, New York, 2023), Vol. 1, pp. 928–938.
- [25] P. Das, C. A. Pattison, S. Manne, D. M. Carmean, K. M. Svore, M. Qureshi, and N. Delfosse, Afs: Accurate, fast, and scalable error-decoding for fault-tolerant quantum computers, in *Proceedings of the 2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, New York, 2022), pp. 259–273.
- [26] X. Tan, F. Zhang, R. Chao, Y. Shi, and J. Chen, Scalable surface-code decoders with parallelization in time, *PRX Quantum* **4**, 040344 (2023).
- [27] N. Liyanage, Y. Wu, A. Deters, and L. Zhong, Scalable quantum error correction for surface codes using FPGA, in *Proceedings of the 2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, New York, 2023), Vol. 1, pp. 916–927.
- [28] H. Bombín, C. Dawson, Y.-H. Liu, N. Nickerson, F. Pastawski, and S. Roberts, Modular decoding: Parallelizable real-time decoding for quantum computers, [arXiv:2303.04846](https://arxiv.org/abs/2303.04846).
- [29] N. Delfosse and G. Zémor, Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel, *Phys. Rev. Res.* **2**, 033042 (2020).
- [30] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, *Phys. Rev. A* **54**, 1098 (1996).
- [31] A. M. Steane, Error correcting codes in quantum theory, *Phys. Rev. Lett.* **77**, 793 (1996).
- [32] A. Yu. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys.* **303**, 2 (2003).
- [33] Kitaev surface code, in *The Error Correction Zoo*, edited by V. V. Albert and P. Faist (2025), <https://errorcorrectionzoo.org/c/surface>.
- [34] N. Delfosse, P. Iyer, and D. Poulin, Generalized surface codes and packing of logical qubits, [arXiv:1606.07116](https://arxiv.org/abs/1606.07116).
- [35] T.-C. Lin and E. Portnoy (unpublished).
- [36] S. Gu, C. A. Pattison, and E. Tang, An efficient decoder for a linear distance quantum ldpc code, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (ACM Press, New York, NY, 2023), pp. 919–932.
- [37] A. Leverrier and G. Zémor, Efficient decoding up to a constant fraction of the code length for asymptotically good quantum codes, *ACM Trans. Alg.* (2024), doi:10.1145/3663763.
- [38] A. Leverrier and G. Zémor, Efficient decoding up to a constant fraction of the code length for asymptotically good quantum

- codes, in *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA 2023), Florence, Italy* (ACM, New York, 2023), Vol. 1, pp. 1216–1244.
- [39] S. Gu, E. Tang, L. Caha, S. H. Choe, Z. He, and A. Kubica, Single-shot decoding of good quantum ldpc codes, *Commun. Math. Phys.* **405**, 85 (2024).
- [40] J. Hopcroft and R. Tarjan, Algorithm 447: Efficient algorithms for graph manipulation, *Commun. ACM* **16**, 372 (1973).
- [41] A. G. Fowler, Proof of finite surface code threshold for matching, *Phys. Rev. Lett.* **109**, 180502 (2012).
- [42] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. (MIT Press, Cambridge, MA, 2009).
- [43] B. A. Galler and M. J. Fisher, An improved equivalence algorithm, *Commun. ACM* **7**, 301 (1964).
- [44] R. E. Tarjan, Efficiency of a good but not linear set union algorithm, *J. ACM* **22**, 215 (1975).