

Almost-linear time decoding algorithm for topological codes

Nicolas Delfosse^{1,2,3} and Naomi H. Nickerson⁴

¹IQIM, California Institute of Technology, Pasadena, CA, USA

²Department of Physics and Astronomy, University of California, Riverside, CA, USA

³Station Q Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA

⁴Quantum Optics and Laser Science, Blackett Laboratory, Imperial College London, Prince Consort Road, London SW7 2AZ, United Kingdom

In order to build a large scale quantum computer, one must be able to correct errors extremely fast. We design a fast decoding algorithm for topological codes to correct for Pauli errors and erasure and combination of both errors and erasure. Our algorithm has a worst case complexity of $O(n\alpha(n))$, where n is the number of physical qubits and α is the inverse of Ackermann’s function, which is very slowly growing. For all practical purposes, $\alpha(n) \leq 3$. We prove that our algorithm performs optimally for errors of weight up to $(d-1)/2$ and for loss of up to $d-1$ qubits, where d is the minimum distance of the code. Numerically, we obtain a threshold of 9.9% for the 2d-toric code with perfect syndrome measurements and 2.6% with faulty measurements.

Introduction

The main obstacle to the construction of a quantum computer is the unavoidable presence of errors, which left unchecked quickly destroy quantum information. Error correction will therefore be essential to perform meaningful quantum computation. Experimental efforts have made rapid progress in recent years [2, 3, 14, 15, 33, 35, 39, 40, 47, 48], and may soon have the capabilities to demonstrate small-scale error correction. Topological codes, in particular Kitaev’s surface code [42], are currently expected to form the core architecture of this first generation of quantum computers, due to their high thresholds and their locality. However, to use these codes, we also require a classical decoding algorithm, which must process measurement information fast enough to keep pace with the clock-speed of the quantum device. While the question of which codes will be the first to be realized seems to be answered, no existing decoder is yet fast enough to match the speeds that the first generation of quantum processors will require [24].

Many decoding algorithms that run in polynomial time have developed [1, 5, 7, 9, 10, 17, 19–

Nicolas Delfosse: nidelfos@microsoft.com

23, 25–27, 32, 34, 37, 38, 45, 46, 50, 53, 55, 59–61, 63, 64, 66, 67] but although this is considered efficient, in practice quadratic or cubic complexity is likely too slow to correct errors faster than they accumulate in a quantum device. Minimum weight perfect matching (MWPM) decoder [20] is currently the most standard decoder for topological codes, and has a worst case complexity of between $O(n^3)$ and $O(n^7)$ depending on the implementation [44]. Significant efforts have been made to optimize its performance [27], and extend it to more general noise models [4, 54, 56, 65]. Most notably, Fowler has achieved large speed improvements [25]. But despite this, further speed-up is required if the decoder is to be practical in a real device.

In this work, we design a decoding algorithm for topological codes that runs in the worst case in almost-linear time in the number of physical qubits n , with a high threshold (See Table 1). We focus on the worst case complexity and not only on the average case complexity since it is the maximum running time of the decoder that will determine the clock-time of the quantum computer. Our key insight is the use of the Union-Find data-structure algorithm [31, 57] that allows us to dynamically keep track of and update the estimation of the error as the decoder runs. We obtain a threshold of 9.9% for correction of phase-flip or bit-flip Pauli error and 2.6% with faulty syndrome measurement. By *almost-linear* complexity, we do not mean $O(n \log n)$ but even lower. Our decoder has a worst-case complexity $O(n\alpha(n))$, where α is the inverse of Ackermann’s function [57]. Although it is not formally linear, α is so slowly growing that it can be considered as a constant. If the number of physical qubits used is smaller than the number of atoms in the universe, then $\alpha(n) \leq 3$.

Table 1: Comparison of the Union-Find and MWPM decoders’ thresholds under phase-flip error.

	UF decoder	MWPM decoder
2d-toric code	9.9%	10.3% [20]
2+1d-toric code	2.6%	2.9% [62]
2d-hexagonal color code	8.4%	8.7% [17]

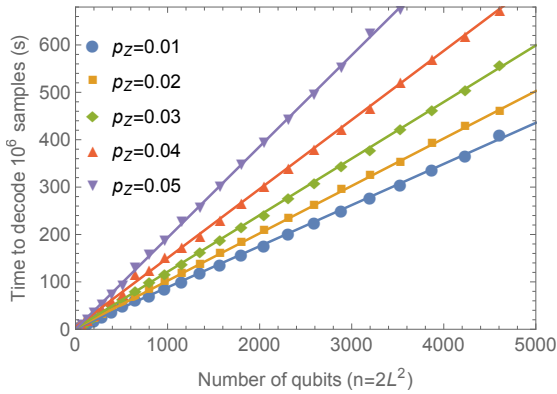


Figure 1: **Average running time.** We show the time taken to perform 10^6 Montecarlo trials of decoding for a 2d toric lattice, under Pauli error and erasure error, with $p_e = 0.1$. The average decoder running time increases linearly with the number of qubits. Our decoder was implemented in C, and run on a single 2.9 GHz Intel Core i5 CPU.

We begin by introducing the surface code in Section 1. In Section 2 we introduce an outline of the decoding algorithm, and give analytic arguments about its performance in Section 3. Sections 4 and 5 contains our main result, the description of the implementation that can be used to implement our decoding algorithm in almost-linear time. Finally we present our numerical results, and discuss the application of the decoder beyond the surface code in Section 6. Further discussion of the complexity scaling, and numerical simulations is given in the Appendices.

1 Background - the surface code

Our decoding algorithm generalizes to any surface code with any genus, with or without boundaries [8, 29] as well as to color codes [6]. For simplicity, we choose to describe the implementation for the surface code without periodic boundary conditions.

The surface code: The surface code, introduced by Kitaev [42], is a topological code, defined on a square lattice of the torus, where a qubit is placed on each edge. Denote respectively by V, E, F the set of vertices, edges and faces of the lattice. The code is defined to be the ground space of the Hamiltonian,

$$H = - \sum_{v \in V} X_v - \sum_{f \in F} Z_f$$

There is an operator X_v associated with each vertex v of the lattice and a plaquette operator Z_f associated with each face f . X_v is the product of the Pauli- X matrices acting on the edges incident to v , *i.e.* $X_v = \prod_{e \in v} X_e$, and $Z_f = \prod_{e \in f} Z_e$ is the product of the Pauli- Z acting on all edge of the face f . The code space is defined as the simultaneous ‘+1’ eigenstate of these operators X_v and Z_f . These operators, and




Figure 2: **The two stages of decoding.** The erasure decoder proceeds in two stages. We begin with a syndrome, σ , caused by an error, E_Z , that is supported only partially in ε . In the first stage, syndrome validation modifies the erasure to ε' in such a way that there is an error E'_Z , entirely supported within ε' , whose syndrome is also σ . In the second stage σ and ε' are decoded using the erasure decoder.

any product of them, are called the *stabilizers* of the code, and form the stabilizer group, S .

Error model: For simplicity, we consider only i.i.d. phase-flip errors, where each qubit is subjected to a Z -error with probability p_Z . The X -part of a Pauli error can be corrected identically. In addition to Pauli errors, a qubit may be *erased*, with probability p_e . We call the set of all erased qubits, *the erasure*, ε . We use the term erasure to describe the error channel through which a qubit at a *known location* is subjected to a Pauli Z error with probability $1/2$. The terms *erasure*, *loss* and *leakage* are sometimes used interchangeably, but physically loss and leakage are two separate mechanisms through which a qubit can be erased. If we detect that a qubit has left the computational subspace (leakage), or that it is physically missing (lost), it can be reinitialized or replaced, which corresponds to a random Pauli error after measurement of the stabilizers. However, it is worth noting that other physical processes could cause an error at a known location, and these can be treated identically as an erasure.

Error correction: Error correction proceeds by measuring the stabilizer operators X_v . When an error $E_Z \in \{I, Z\}^{\otimes n}$ has affected the qubits of the code, any stabilizer X_v that anticommutes with the error returns a ‘-1’ outcome. The subset of vertices, v , with -1 measurement outcomes is called the *syndrome*, σ . Given a syndrome, σ , and an erasure, ε , the role of the decoder is to find a correction operator, $\mathcal{C}(\sigma, \varepsilon)$, such that $\mathcal{C}(\sigma, \varepsilon)E_Z \in S$. That is, when the correction operator is applied to the code, the error is corrected up to a stabilizer.

2 Union-Find decoder for surface codes

We begin by introducing an outline of the decoding procedure, without the details of its implementation that are required to achieve low complexity.

Our algorithm is motivated by the fact that erasure errors are much simpler to decode than Pauli errors. Simply put, we can describe an erasure error as a Pauli error at a known location. When erasure is the only

source of error the decoding problem is significantly simplified, as errors can only be present within the erasure. The peeling decoder for erasure introduced in [18] uses this fact to find a correction in the case that only erasure errors are present. When Pauli errors can also occur this algorithm will necessarily fail, as the correction now has support outside the erasure.

The decoder is divided into two stages, as shown in Figure 2. The goal of the first stage is to take a syndrome generated by both Pauli error and erasure, ε , and from this generate a modified erasure, ε' , such that there is a valid correction operator supported entirely in ε' . We call this stage *syndrome validation*. After the syndrome is validated, we can apply the peeling decoder to find a correction.

To perform syndrome validation, we identify ‘invalid’ clusters of erasures, and iteratively grow them until the updated state is correctable by the erasure decoder. This idea of growing and merging clusters in order to correct them as locally as possible has been previously explored [19, 21, 32], but in order to make such a decoder fast one needs to be able rapidly update the clusters as they grow. This dynamical update gives the dominant contribution to the decoding complexity.

To define how we can identify invalid clusters of erasures, we state our first lemma.

Lemma 1 (even vs odd clusters). *Let ε be a connected subset of edges and let σ be a set of syndrome vertices included in ε . There exists a Z -error $E_Z \subset \varepsilon$ of syndrome σ if and only if the cardinality of σ is even.*

This is straightforward to prove. We define a cluster to be a connected component of erased qubits in the subgraph (V, ε) . These clusters can be either a connected subgraph induced by a subset of erased edges, or an isolated vertex. If no Pauli error is present, it must be the case that the cluster supports an even number of syndrome bits. If the cluster supports an odd number of syndrome vertices then we identify it as ‘invalid’, at least one error chain must terminate in this cluster. The erasure decoder cannot be applied directly to odd clusters

Since we cannot apply the erasure decoder to odd clusters directly, instead we grow these clusters by adding edges to the erasure until they connect with another odd cluster. When two odd clusters merge the resulting cluster is even, and can therefore be corrected. Note that a single vertex outside the erasure supporting a syndrome bit is an odd cluster.

Algorithm 1 describes this procedure. Steps 1-7 perform syndrome validation, and in step 8 the erasure decoder is applied. Lemma 1 proves that we can apply the erasure decoding at the end of Algorithm 1 and that this final step will return a correction \mathcal{C} that

is consistent with the syndrome. Figure 3 shows an example of the growing and merging procedure.

We will later discuss the detail of the implementation required to achieve a low complexity, but first we focus on the error tolerance of the algorithm.

Algorithm 1: Union-Find decoder – Naive version

input : The set of erased positions $\varepsilon \subset E$ and the syndrome $\sigma \subset V$ of an error E_Z .

output: An estimation \mathcal{C} of E_Z 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 all** odd cluster C_i **do**
 - 4 Grow C_i by increasing its radius by one half-edge.
 - 5 If C_i meets another cluster, fuse and update parity.
 - 6 If C_i is even, remove it from the odd cluster list.
 - 7 Add full edges that are in the grown clusters to ε' .
 - 8 Apply the peeling decoder to the erasure to find \mathcal{C} .
-

3 Decoder Performance

We now discuss the performance of the decoder in the low error regime, and show that it performs just as well as the most likely error (MLE) decoder below the minimum distance. MLE decoder returns the most likely error configuration given the observed syndrome. The MWPM-decoder is an implementation of an MLE decoder.

For any code, the MLE decoder can correct any error configuration, E_Z , with a weight of up to $(d-1)/2$, where d is the minimum distance of the code. For the erasure channel, any erasure pattern of up to $d-1$ qubits can be corrected. Moreover, both of these bounds are tight. Reaching these values with an efficient decoding algorithm is evidence of good performance.

In our mixed noise model, MLE decoder can correct any combination of t erased qubits and s Z -errors (outside the erased set) as long as $t+2s < d$. This bound, generalizing both previous cases, is also tight.

Theorem 1. *If $t+2s < d$, Algorithm 1 can correct any combination of t erased qubits and s Z -error.*

This algorithm performs well for both Pauli errors, erasures and combinations of the two.

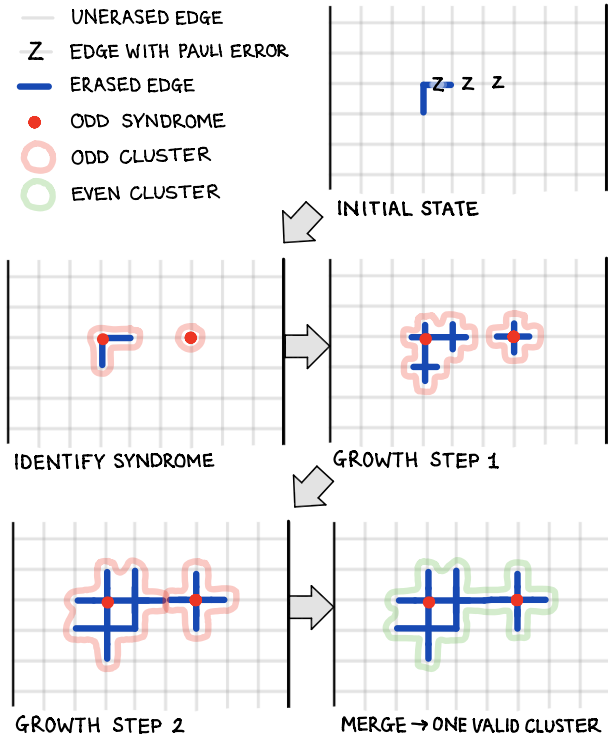


Figure 3: Schematic representation of syndrome validation.

Proof. For the case of erasures the proof is straightforward since the erasure decoder can be called immediately without going through steps 1. to 7. Consider next the other extreme where only Pauli errors occur and denote the error by E_Z . A cluster C grows in Algorithm 1 (Step 4.), when there are an odd number of syndrome vertices contained within it. This implies that there exists at least one path in the support of E_Z connecting a vertex of C to a vertex outside of C . Therefore, when a cluster grows, at least one new half-edge of E_Z is covered. After at most $2s$ rounds of growth, the grown cluster covers the entire error E_Z and there can be no more odd clusters left to grow. At the end of the growing procedure, the diameter of the largest erased cluster is at most $2s$ edges ($4s$ half-edges). By erasing this cluster, we obtain an erasure pattern that covers E_Z and that does not cover a non-trivial logical error since $2s < d$. When the peeling decoder is run in the final step it must therefore succeed at identifying E_Z up to a stabilizer. This argument relies on the optimality of the peeling decoder proven in [18]

The general argument for a combination of s errors and t erasures is similar. Growing the clusters increases the diameter of the largest cluster by at most $2s$. It is then upper bounded by $2s + t < d$. Just as in the case of only Pauli errors, the final step returns an error equivalent to E_Z , up to a stabilizer. \square

4 Achieving almost-linear complexity

Our next goal is to show that our decoder can be implemented in almost-linear time in the number of qubits, n , by exploiting Union-Find data-structure algorithms [31, 57].

4.1 Union-Find algorithm for cluster growth

The key function of the decoder is to grow clusters, and fuse them when they meet. For this we need two things: a function $\text{Union}(u, v)$ that performs the fusion operation on two clusters C_u and C_v , and a function $\text{Find}(v)$ that identifies the cluster to which vertex v belongs. The $\text{Find}()$ function allows us to distinguish clusters from one another, since we only wish to fuse clusters when they are distinct. When an edge, (u, v) is added to a growing cluster, we call the function $\text{Find}(u)$ and $\text{Find}(v)$ on its two endpoints. If the endpoints belong to the same cluster, $\text{Find}(u) = \text{Find}(v)$, in which case we do nothing. If $\text{Find}(u) \neq \text{Find}(v)$, they belong to different clusters and we must fuse C_u and C_v . The complexity of this subroutine provides the leading order in the complexity of our decoder.

A naive algorithm: We first describe a naive implementation of the $\text{Union}()$ and $\text{Find}()$ functions. Let us store an index $\text{Find}(v)$ for each vertex v in a look-up table of size $|V|$. This makes the cluster identification trivial. However, when $\text{Union}(u, v)$ is called one must update the cluster indices $\text{Find}(w)$ for all the vertices w of one of the two clusters C_u or C_v in order to correctly update the state. This might require a number of updates of cluster indices $\text{Find}(v)$, which is itself linear in n . In the worst case, we call Union up to $n - 1$ times, yielding a quadratic overall complexity, $O(n^2)$.

4.2 Implementation

We now describe the data structure, and steps required in the implementation to achieve an almost-linear complexity.

Weighted Union: In order to reduce the complexity due to updating the cluster index after merging, we can choose to always select the smallest component of the two clusters, C_u and C_v , to update. The size of the small cluster at least doubles at each call of $\text{Union}()$, which means that every vertex index is updated at most $O(\log(n))$ times, reducing the complexity of the Union-Find update to $O(n \log n)$. To do this it is also necessary to store the size of each component, but this does not affect the complexity, as we can simply add them to our look-up table.

Tree representation: We now consider how the clusters can be stored in memory in order to speed up the index update in $\text{Union}()$. The representation we




Figure 4: **A cluster-tree data structure.** The tree is initialized with depth one but this may increase during the decoding procedure. The size and parity are stored at the root of the cluster-tree.

choose increases the cost of the $\text{Find}()$ function, but overall the complexity is reduced. We represent each cluster C as a tree that we call a *cluster tree*. The vertices of the cluster-tree correspond to the vertices of the cluster C of the lattice, however, the cluster-tree is an arbitrary tree and does not have to respect the lattice structure. An example of this data structure is shown in Figure 4. An arbitrary vertex u of C is chosen as the *root* of the cluster-tree and is used to identify the cluster. The size of the cluster and its parity are also stored at its root. $\text{Find}(v)$ returns the root of u , for which we must traverse the tree from u to its root. The cost of $\text{Find}()$ is therefore given by the depth of the tree. To minimize this cost, the cluster-trees are initialized with depth 1 at the beginning of the Union-Find decoder. When $\text{Union}(u, v)$ is applied, one must first check whether u and v belong to the same cluster. To do this we call $\text{Find}(u)$ and $\text{Find}(v)$ and walk through the trees from u and v to their respective roots. If these components are distinct, the trees are merged in two steps. First, the smallest cluster-tree (say C_v) is added as a subtree of the root $\text{Find}(u)$ of the largest cluster-tree (C_u). Then, the size and parity stored at the root is updated. Path compression can then be applied to minimize the depth of the new clusters.

Path compression: When $\text{Find}(u)$ is called, we step through the tree until we reach its root. Once the root of u is reached we compress the by pointing all the vertices encountered along the path directly to the root. This only doubles the cost $\text{Find}()$ but speeds up any future call to the function that passes through any part of the compressed path. Using path compression alone, without weighted union also leads to a complexity of $O(n \log n)$. However, when combining both procedures together there is a further reduction of the complexity. The analysis of Union-Find algorithms with both weighted union and path compression is quite involved, Tarjan proved that remarkably the worst-case complexity is $O(n\alpha(n))$ [57]. It has also been shown that this result is optimal [28, 58].

Support of the growing clusters: The cluster-trees encode the set of vertices of the clusters. The

edges are represented by a look-up table, Support , of size $|E|$ that stores the state of an edge, which can take one of three values: *Unoccupied*, *Grown* or *Half-grown*. When an edge is half-grown also store the vertex of the lattice from which it was grown. Support is initialized with all erased edges marked as *Grown* and all other edges being *Unoccupied*. As the algorithm proceeds, half-edges are grown, growing by one half-edge twice produces a grown edge.

Boundary representation: In order to preserve the low complexity of the dynamical fusion using Union-Find, we must be careful when updating Support . In each round of growth we must identify the boundary vertices of a cluster from which we will grow new edges. By a *boundary vertex*, we mean a vertex of the lattice such that at least one of the incident edges is not in the erasure, ε . One way to do this would be to explore the cluster to identify the boundary, but this would induce a linear cost in the size of the cluster in each round, resulting in a quadratic complexity overall. We now describe a more efficient strategy.

We can avoid recomputing the boundary of the cluster by storing a list of boundary vertices for each root. To grow a cluster, we must then simply iterate over this list and grow the incident edges. When clusters are fused, it is necessary to update the boundary lists. We now describe a subroutine $\text{Grow}(\mathcal{L})$ that can achieve this in linear time. This function takes a list $\mathcal{L} = \{u_1, u_2, \dots, u_m\}$ of cluster roots as an input, each with its associated boundary list. The function $\text{Grow}(\mathcal{L})$ increases the radius of these clusters by one half-edge, fuses the clusters that meet and updates the clusters and their boundary lists. The function runs in five steps.

- (i) *Growth:* Grow all the clusters of root u_i by one half-edge. This is done by running over the lists of boundary vertices and by growing the incident half-edges in the table Support . This step returns a list of all the newly grown edges that connect two distinct clusters. Call these edges *fusion edges*.
- (ii) *Fusion of clusters:* Run over fusion edges $e = \{u, v\}$ and if u and v belong to distinct clusters then merge them with $\text{Union}(u, v)$, otherwise remove e from the list of fusion edges.
- (iii) *Fusion of boundary lists:* Run over fusion edges (updated in (ii)) and for each $e = \{u, v\}$, append the boundary list of the smallest cluster C_u or C_v (before fusion) at the end of the boundary list of the largest one (before fusion).
- (iv) *Update roots:* Replace each element u_i of the root list \mathcal{L} by $\text{Find}(u_i)$.

- (v) *Update Boundary lists*: Run once more over each boundary list and remove the vertices that are no longer boundary vertices.

We remark that new boundary vertices are added when merging the boundary list with the boundary list of a neighbor cluster. We only need to remove vertices from those lists in the last step to update them.

Avoiding growth duplication: In order to simplify the description of the algorithm, we have so far omitted a detail of the implementation. If two clusters of the list are fused during step (ii), then updating their roots in (iv) results in a duplicated root in the list. We can avoid this issue by also storing in a look-up table an indicator that marks the vertices of the odd-root list at the beginning of each round of growth. Maintaining this extra look-up table (of linear-size) does not increase the complexity of the algorithm. This modification allows us to detect the presence of $Find(u)$ in the list before replacing u by $Find(u)$. If it is already present, we simply remove u from the list, and avoid creating a duplicate.

4.3 Summary of data structure and algorithm

Each cluster C in the lattice is encoded using the following data:

- **Cluster-tree:** A tree whose vertices encode the vertices of C with an arbitrary root.
- **Size and Parity:** The size of the cluster and the parity of C are stored at the root of the cluster-tree.
- **Support:** A look-up table that stores the state of each edge in the growing clusters.
- **Boundary List:** The list of all the boundary vertices of C .

A full version of the algorithm we have described is given in Algorithm 2. Let us now summarize the contribution of each part of the algorithm to the overall complexity.

Complexity of the full algorithm: Algorithm 2 can be decomposed into three blocks. The first block contains lines 1-2 and initializes the clusters. Block 2 contains lines 3-11 and relies on a Union-Find algorithm that we have just described. The complexity of this subroutine provides the leading order in the complexity of our decoder. In the third block, line 12, the erasure decoder is applied.

Creating the list of clusters can be achieved in linear time by exploring the connected components of the sublattice of erased edges [36]. During this same exploration, we can compute the size and syndrome parity of the cluster, and store it at the root, which

does not affect the complexity. The table Support is also initialized in linear time. The second block grows and merges odd clusters until they disappear and runs from line 3 to 11. We have discussed how the Union-Find algorithm can be used to achieve this in $O(n\alpha(n))$ time. Each of the n vertices can be at a boundary during at most two rounds of growth, and during the growing procedure each list is iterated over $O(1)$ times. The complexity of growing Support and updating the boundary lists is therefore $O(n)$. The support of the clusters and the boundary lists can be updated in linear time $O(n)$ using Grow(). The last two instructions are simply the erasure decoder and are already known to have linear complexity [18]. The dominant contribution to the complexity is therefore the $O(n\alpha(n))$ due to the Union-Find algorithm used to keep track of the cluster-trees during syndrome validation.

Algorithm 2: Union-Find decoder – Almost-linear time version

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

- 1 Initialize cluster-trees, Support and boundary lists for all clusters.
 - 2 Create the lists \mathcal{L} of roots of odd clusters.
 - 3 **while** \mathcal{L} is not empty **do**
 - 4 (o) Initialize the fusion list \mathcal{F} as an empty list.
 - 5 (i) For all $u \in \mathcal{L}$, grow the cluster C_u of a half-edge in the Table Support. If a new grown edge e is added in Support then add e to the fusion list \mathcal{F} .
 - 6 (ii) For all $e = \{u, v\} \in \mathcal{F}$, if $Find(u) \neq Find(v)$ then applies Union(u, v) to merge the cluster C_u and C_v . If $Find(u) = Find(v)$ then remove e from the list \mathcal{F} .
 - 7 (iii) For all $e = \{u, v\} \in \mathcal{F}$, read the sizes of the clusters C_u and C_v stored at the roots (uses $Find()$) and append the boundary list of the smallest cluster at the end of the boundary list of the largest one.
 - 8 (iv) Replace each root $u \in \mathcal{L}$ by $Find(u_i)$. (in a way that does not create duplicated elements)
 - 9 (v) For all $u \in \mathcal{L}$, remove the vertices of the boundary list of u that are not boundary vertices.
 - 10 (vi) For all $u \in \mathcal{L}$, if C_u is an even cluster then remove it from \mathcal{L} .
 - 11 Erase all the edges that are fully grown in *Support*.
 - 12 Apply the peeling decoder to the erasure.
-

5 Weighted growth version of the Union-Find decoder

The algorithm we have described so far grows all (odd) clusters uniformly in each round. We now discuss one simple way in which this growth strategy can be altered to improve the performance of the decoder. We recall that our basic argument in the proof of Theorem 1 was that if a cluster is odd there exists at least one path of errors connecting this cluster to a vertex outside the cluster. When the error rate is small, it is likely that only one error chain terminates in the cluster. When the cluster grows, we add some number of edges that is proportional to the number of boundary vertices. If we add b new edges, then only $1/b$ of those edges correctly cover the error. The larger the boundary of the cluster, the more 'incorrect' edges we are adding. By growing smaller clusters first, fewer erasures will be added in total, which increases the chance success in the the final erasure decoding step .

The *Weighted Growth* version of the Union-Find decoder is achieved by always growing the cluster with the smallest boundary size. This small modification improves the threshold of the 2d toric code from 9.2% to 9.9% for phase-flip errors and from 2.4% to 2.6% with faulty-syndrome measurements.

The Weighted Growth Union-Find decoder can also be implemented in almost-linear time. To achieve this complexity, we need to sort the clusters by boundary size. This can be done in linear time in n because clusters boundary size are integer with values between 1 and $O(n)$.

6 Application to Quantum Computing

We have so far described the decoder only for the 2d surface code. To be of use for quantum computing applications, where measurements may be faulty, we must be able to solve the (2+1)-dimensional variant of the decoding problem [20]. Multiple rounds of syndrome measurement on the surface code produce a three-dimensional cubic lattice of syndrome outcomes, where the x-y directions correspond to the physical code, and the third dimension represents time. Space-like edges in this syndrome lattice correspond to physical errors on the qubits of the code, while time-like edges correspond to measurement errors. The decoder as we have described it requires no adaption to be applied to the 3d lattice.

Conceptually we can understand the decoder as adding virtual erasures to the syndrome information that include the support of the error operator, and we can extend this interpretation to the 3d case. Erasing the space-like edges of the syndrome graph corresponds to erasure of a physical qubit, while erasing a time-like edge of the syndrome graph corresponds

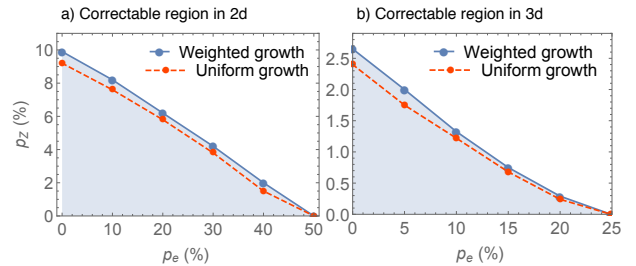


Figure 5: **Performance results.** a) Results of numerical threshold simulations on the 2d toric code under an error model of erasure and Pauli error at independent error rates p_e and p_Z . The shaded area indicates the correctable region, which is below threshold. Example threshold plots can be seen in the Appendix. b) Results of numerical threshold simulations of the 2+1d toric code under an error model of erasure and phenomenological error at independent error rates p_e and p_Z .

to the erasure of a measurement outcome - a classical erasure. Physically such a measurement erasure corresponds to the case that no stabilizer outcome was recorded.

We performed numerical simulations to quantify the performance of the decoder in both 2d and 3d. In 2d each qubit is treated as erased with a probability p_e and acquires a Pauli Z error with a probability p_Z . In 3d we consider making L rounds of repeated syndrome measurements, for a lattice of dimension L . Each measurement outcome is erased with probability p_e and returns the incorrect value with probability p_Z . Between rounds of measurement each qubit is erased with probability p_e and acquires a Pauli Z error with probability p_Z . This corresponds to the phenomenological error model in the case that $p_e = 0$. More details of the numerics are described in Appendix B. The results of threshold calculations are shown in Figure 5a) and b), defining a correctable region in the space of erasure and Pauli error. In 2d, with no erasure, we find $p_{th} = 9.9\%$ and in 3d we find $p_{th} = 2.6\%$. This is only a small reduction in the decoder performance compared with MWPM (see Table-1). Below threshold the logical error rate is well behaved, showing a clear exponential suppression with increasing lattice size.

We measured the average running time to perform decoding as a function of the number of qubits, which gives numerical evidence of a linear average complexity. Figure 1 shows the running time, for a C implementation of the decoder running on a single 2.9 GHz Intel Core i5 CPU.

6.1 Application beyond the surface code

The functionality of the Union-Find decoder extends beyond Pauli errors, as the decoder naturally handles a mixed noise model of erasure and Pauli error. Bar-

rett and Stace [4, 54] first introduced an algorithm to decode the Raussendorf lattice [50–52] in the presence of erasure, which combines MWPM with geometric deformations of the lattice around erased qubits, and runs in $O(n^3)$ time. This noise model is particularly relevant for photonic measurement based quantum computing [13, 41, 43, 49], where qubit erasure is likely to be a dominant source of error. However, in these architectures the complexity of the Barrett-Stace algorithm would make it challenging to implement fast enough to keep up with the natural clock-speed of a quantum computer. The Union-Find decoder provides a significant decrease in computational requirements for decoding in a photonic architecture, with only a small decrease in threshold.

As well as its speed, one advantage of the Union-Find decoder is its immediate flexibility to other geometries and dimensions. The Union-Find decoder requires only an underlying graph structure, and therefore requires no alteration to run in any geometry as long as the errors form string-like objects. This makes it a simple tool for use in complex geometries, such as hyperbolic codes [11, 12, 16, 30, 68]. In contrast, MWPM must in general be combined with costly path finding algorithms to compute the distance between syndrome measurements. Here we have only simulated the performance on a torus, but non periodic boundary conditions can be handled by allowing clusters to become valid by merging with a rough boundary. Furthermore, the decoder can also be used on a syndrome graph of any dimension, although we expect that the threshold performance of the Union-Find algorithm, at least in its basic form, will decrease relative to perfect matching as the dimension of the space increases.

Beyond surface codes, the Union-Find decoder can also be used to decode the 2D color code, when used in combination with the method introduced in [17] for projecting color codes onto surface codes. This procedure does not affect the complexity of the algorithm. By numerically simulating decoding on a hexagonal lattice, we find a threshold for the [6,6,6] color code of 8.4% under i.i.d Pauli Z error.

7 Conclusion

We have presented a decoder with a high threshold which has a considerably lower worst-case complexity than any other existing practical decoder. Indeed, we are very close to the the best possible complexity for any decoder that is not parallelized. Since one must at a minimum iterate at least once over the syndrome, the lowest possible time is $O(n)$. But we are interested in practical applications, and complexity is not the only figure of merit. It is important that our algorithm also has a small constant overhead, making it fast in practice as well as in theory.

Instead of designing a fast decoder which approximates the MWPM decoder as we did in this work, one could consider a parallel MWPM decoder [25] which achieves a linear average-case complexity. It would be interesting to investigate the parallelization of the Union-Find decoder.

Many questions remain to be answered. Our implementation achieved linear complexity, but was not heavily optimized for speed. Work should be done to understand how much this can be improved, in particular by implementing the algorithm in hardware. In terms of error tolerance, we have studied only simple noise models, and the performance under circuit-level error will need to be studied to draw a more meaningful comparison with the thresholds of other decoders.

There are many challenges still to solve before an error corrected quantum computer can be realized, and one of these is decoding. Our algorithm has important practical implications in achieving this goal, and is a significant step towards overcoming the hurdle of fast decoding in real quantum devices.

Acknowledgement – The authors would like to thank Eric Johnson and Chris Dawson for valuable discussions, and Terry Rudolph for first introducing them to the question of error correction in photonic devices. The authors would like to thank Aleksander Kubica for his comments on a preliminary version of this paper. ND acknowledges funding provided by the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-2644).

References

- [1] Hussain Anwar, Benjamin J Brown, Earl T Campbell, and Dan E Browne. Fast decoders for qudit topological codes. *New Journal of Physics*, 16(6):063038, 2014. DOI: [10.1088/1367-2630/16/6/063038](https://doi.org/10.1088/1367-2630/16/6/063038).
- [2] CJ Ballance, TP Harty, NM Linke, and DM Lucas. High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions. *arXiv preprint arXiv:1406.5473*, 2014. DOI: [10.1103/PhysRevLett.117.060504](https://doi.org/10.1103/PhysRevLett.117.060504).
- [3] Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C White, Josh Mutus, Austin G Fowler, Brooks Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014. DOI: [10.1038/nature13171](https://doi.org/10.1038/nature13171).
- [4] Sean D Barrett and Thomas M Stace. Fault tolerant quantum computation with very high threshold for loss errors. *Physical review letters*, 105(20):200502, 2010. DOI: [10.1103/PhysRevLett.105.200502](https://doi.org/10.1103/PhysRevLett.105.200502).

- [5] Sean D Barrett and Thomas M Stace. Fault tolerant quantum computation with very high threshold for loss errors. *Physical review letters*, 105(20):200502, 2010. DOI: [10.1103/PhysRevLett.105.200502](https://doi.org/10.1103/PhysRevLett.105.200502).
- [6] Hector Bombin and Miguel Angel Martin-Delgado. Topological quantum distillation. *Physical review letters*, 97(18):180501, 2006. DOI: [10.1103/PhysRevLett.97.180501](https://doi.org/10.1103/PhysRevLett.97.180501).
- [7] Hector Bombin, Guillaume Duclos-Cianci, and David Poulin. Universal topological phase of two-dimensional stabilizer codes. *New Journal of Physics*, 14(7):073048, 2012. DOI: [10.1088/1367-2630/14/7/073048](https://doi.org/10.1088/1367-2630/14/7/073048).
- [8] S. B. Bravyi and A. Y. Kitaev. Quantum codes on a lattice with boundary. *arXiv preprint arXiv:9811052*, 1998.
- [9] Sergey Bravyi and Jeongwan Haah. Quantum self-correction in the 3d cubic code model. *Physical review letters*, 111(20):200501, 2013. DOI: [10.1103/PhysRevLett.111.200501](https://doi.org/10.1103/PhysRevLett.111.200501).
- [10] Sergey Bravyi, Martin Suchara, and Alexander Vargo. Efficient algorithms for maximum likelihood decoding in the surface code. *Physical Review A*, 90(3):032326, 2014. DOI: [10.1103/PhysRevA.90.032326](https://doi.org/10.1103/PhysRevA.90.032326).
- [11] Nikolas P Breuckmann and Barbara M Terhal. Constructions and noise threshold of hyperbolic surface codes. *IEEE Transactions on Information Theory*, 62(6):3731–3744, 2016. DOI: [10.1109/TIT.2016.2555700](https://doi.org/10.1109/TIT.2016.2555700).
- [12] Nikolas P Breuckmann, Christophe Vuillot, Earl Campbell, Anirudh Krishna, and Barbara M Terhal. Hyperbolic and semi-hyperbolic surface codes for quantum storage. *arXiv preprint arXiv:1703.00590*, 2017. DOI: [10.1088/2058-9565/aa7d3b](https://doi.org/10.1088/2058-9565/aa7d3b).
- [13] Daniel E Browne and Terry Rudolph. Resource-efficient linear optical quantum computation. *Physical Review Letters*, 95(1):010501, 2005. DOI: [10.1103/PhysRevLett.95.010501](https://doi.org/10.1103/PhysRevLett.95.010501).
- [14] Jacques Carolan, Christopher Harrold, Chris Sparrow, Enrique Martín-López, Nicholas J Russell, Joshua W Silverstone, Peter J Shadbolt, Nobuyuki Matsuda, Manabu Oguma, Mikitaka Itoh, et al. Universal linear optics. *Science*, 349(6249):711–716, 2015. DOI: [10.1126/science.aab3642](https://doi.org/10.1126/science.aab3642).
- [15] S Debnath, NM Linke, C Figgatt, KA Landsman, K Wright, and C Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, 2016. DOI: [10.1038/nature18648](https://doi.org/10.1038/nature18648).
- [16] Nicolas Delfosse. Tradeoffs for reliable quantum information storage in surface codes and color codes. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 917–921. IEEE, 2013. DOI: [10.1109/ISIT.2013.6620360](https://doi.org/10.1109/ISIT.2013.6620360).
- [17] Nicolas Delfosse. Decoding color codes by projection onto surface codes. *Physical Review A*, 89(1):012317, 2014. DOI: [10.1103/PhysRevA.89.012317](https://doi.org/10.1103/PhysRevA.89.012317).
- [18] Nicolas Delfosse and Gilles Zémor. Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel. *arXiv preprint arXiv:1703.01517*, 2017. DOI: [10.1103/PhysRevResearch.2.033042](https://doi.org/10.1103/PhysRevResearch.2.033042).
- [19] Eric Dennis. Purifying quantum states: Quantum and classical algorithms. *arXiv preprint quant-ph/0503169*, 2005.
- [20] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002. DOI: [10.1063/1.1499754](https://doi.org/10.1063/1.1499754).
- [21] Guillaume Duclos-Cianci and David Poulin. Fast decoders for topological quantum codes. *Physical review letters*, 104(5):050504, 2010. DOI: [10.1103/PhysRevLett.104.050504](https://doi.org/10.1103/PhysRevLett.104.050504).
- [22] Guillaume Duclos-Cianci and David Poulin. Kitaev’s z d-code threshold estimates. *Physical Review A*, 87(6):062338, 2013. DOI: [10.1103/PhysRevA.87.062338](https://doi.org/10.1103/PhysRevA.87.062338).
- [23] Guillaume Duclos-Cianci and David Poulin. Fault-tolerant renormalization group decoder for abelian topological codes. *Quantum Info. Comput.*, 14(9 & 10):721–740, jul 2014. ISSN 1533-7146. DOI: [10.5555/2638670.2638671](https://doi.org/10.5555/2638670.2638671).
- [24] Austin Fowler. Towards sufficiently fast quantum error correction. *Conference QEC 2017*, 2017.
- [25] Austin G Fowler. Minimum weight perfect matching of fault-tolerant topological quantum error correction in average $o(1)$ parallel time. *Quantum Information and Computation*, 15(1&2):0145–0158, 2015. DOI: [10.5555/2685188.2685197](https://doi.org/10.5555/2685188.2685197).
- [26] Austin G Fowler, Adam C Whiteside, and Lloyd CL Hollenberg. Towards practical classical processing for the surface code. *Physical review letters*, 108(18):180501, 2012. DOI: [10.1103/PhysRevLett.108.180501](https://doi.org/10.1103/PhysRevLett.108.180501).
- [27] Austin G Fowler, Adam C Whiteside, Angus L McInnes, and Alimohammad Rabbani. Topological code autotune. *Physical Review X*, 2(4):041003, 2012. DOI: [10.1103/PhysRevX.2.041003](https://doi.org/10.1103/PhysRevX.2.041003).
- [28] Michael Fredman and Michael Saks. The cell probe complexity of dynamic data structures. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 345–354. ACM, 1989. DOI: [10.1145/73007.73040](https://doi.org/10.1145/73007.73040).
- [29] M. H. Freedman and D. A. Meyer. Projective plane and planar quantum codes. *Foundations of Computational Mathematics*, 1(3):325–332, 2001. DOI: [10.1007/s102080010013](https://doi.org/10.1007/s102080010013).

- [30] M.H. Freedman, D.A. Meyer, F. Luo, and Computer Society (IEEE) Washington DC. *Z(2)-Systolic Freedom and Quantum Codes*. Defense Technical Information Center, 2002. DOI: 10.2140/gtm.1999.2.113. URL <https://books.google.com/books?id=KieaDAEACAAJ>.
- [31] Bernard A Galler and Michael J Fisher. An improved equivalence algorithm. *Communications of the ACM*, 7(5):301–303, 1964. DOI: 10.1145/364099.364331.
- [32] James William Harrington. *Analysis of quantum error-correcting codes: symplectic lattice codes and toric codes*. PhD thesis, California Institute of Technology, 2004.
- [33] Reinier W Heeres, Philip Reinhold, Nissim Ofek, Luigi Frunzio, Liang Jiang, Michel H Devoret, and Robert J Schoelkopf. Implementing a universal gate set on a logical qubit encoded in an oscillator. *Nature Communications*, 8, 2017. DOI: 10.1038/s41467-017-00045-1.
- [34] Michael Herold, Michael J Kastoryano, Earl T Campbell, and Jens Eisert. Fault tolerant dynamical decoders for topological quantum memories. *arXiv preprint arXiv:1511.05579*, 2015. DOI: 10.1088/1367-2630/aa7099.
- [35] Charles D Hill, Eldad Peretz, Samuel J Hile, Matthew G House, Martin Fuechsle, Sven Rogge, Michelle Y Simmons, and Lloyd CL Hollenberg. A surface code quantum computer in silicon. *Science advances*, 1(9):e1500707, 2015. DOI: 10.1038/npjqi.2015.19.
- [36] John Hopcroft and Robert Tarjan. Algorithm 447: efficient algorithms for graph manipulation. *Communications of the ACM*, 16(6):372–378, 1973. DOI: 10.1145/362248.362272.
- [37] Adrian Hutter, James R Wootton, and Daniel Loss. Efficient markov chain monte carlo algorithm for the surface code. *Physical Review A*, 89(2):022326, 2014. DOI: 10.1103/PhysRevA.89.022326.
- [38] Adrian Hutter, Daniel Loss, and James R Wootton. Improved hdrg decoders for qudit and non-abelian quantum error correction. *New Journal of Physics*, 17(3):035017, 2015. DOI: 10.1088/1367-2630/17/3/035017.
- [39] IBM. Quantum experience API. <https://github.com/QISKit>, 2017.
- [40] Norbert Kalb, Andreas A Reiserer, Peter C Humphreys, Jacob JW Bakermans, Sten J Kammerling, Naomi H Nickerson, Simon C Benjamin, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. Entanglement distillation between solid-state quantum network nodes. *Science*, 356(6341):928–932, 2017. DOI: 10.1126/science.aan0070.
- [41] Konrad Kieling, Terry Rudolph, and Jens Eisert. Percolation, renormalization, and quantum computing with nondeterministic gates. *Physical Review Letters*, 99(13):130501, 2007. DOI: 10.1103/PhysRevLett.99.130501.
- [42] A Yu Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003. DOI: 10.1016/S0003-4916(02)00018-0.
- [43] Emanuel Knill, Raymond Laflamme, and Gerald J Milburn. A scheme for efficient quantum computation with linear optics. *nature*, 409(6816):46–52, 2001. DOI: 10.1038/35051009.
- [44] Vladimir Kolmogorov. Blossom v: a new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation*, 1(1):43–67, 2009. DOI: 10.1007/s12532-009-0002-8.
- [45] Aleksander Marek Kubica. *The ABCs of the Color Code: A Study of Topological Quantum Codes as Toy Models for Fault-Tolerant Quantum Computation and Quantum Phases Of Matter*. PhD thesis, California Institute of Technology, 2018.
- [46] Andrew J Landahl, Jonas T Anderson, and Patrick R Rice. Fault-tolerant quantum computing with color codes. *arXiv preprint arXiv:1108.5738*, 2011.
- [47] Bjoern Lekitsch, Sebastian Weidt, Austin G Fowler, Klaus Mølmer, Simon J Devitt, Christof Wunderlich, and Winfried K Hensinger. Blueprint for a microwave trapped ion quantum computer. *Science Advances*, 3(2):e1601540, 2017. DOI: 10.1126/sciadv.1601540.
- [48] Thomas Monz, Daniel Nigg, Esteban A Martinez, Matthias F Brandl, Philipp Schindler, Richard Rines, Shannon X Wang, Isaac L Chuang, and Rainer Blatt. Realization of a scalable shor algorithm. *Science*, 351(6277):1068–1070, 2016. DOI: 10.1126/science.aad9480.
- [49] Michael A Nielsen. Optical quantum computation using cluster states. *Physical review letters*, 93(4):040503, 2004. DOI: 10.1103/PhysRevLett.93.040503.
- [50] R. Raussendorf and J. Harrington. Fault-tolerant quantum computation with high threshold in two dimensions. *Physical Review Letters*, 98(19):190504, 2007. DOI: 10.1103/PhysRevLett.98.190504.
- [51] R. Raussendorf, J. Harrington, and K. Goyal. Topological fault-tolerance in cluster state quantum computation. *New Journal of Physics*, 9: 199, 2007. DOI: 10.1088/1367-2630/9/6/199.
- [52] Robert Raussendorf, Jim Harrington, and Kovid Goyal. A fault-tolerant one-way quantum computer. *Annals of physics*, 321(9):2242–2270, 2006. DOI: 10.1016/j.aop.2006.01.012.
- [53] Pradeep Sarvepalli and Robert Raussendorf. Efficient decoding of topological color codes. *Physical Review A*, 85(2):022317, 2012. DOI: 10.1103/PhysRevA.85.022317.

- [54] Thomas M Stace and Sean D Barrett. Error correction and degeneracy in surface codes suffering loss. *Physical Review A*, 81(2):022317, 2010. DOI: [10.1103/PhysRevA.81.022317](https://doi.org/10.1103/PhysRevA.81.022317).
- [55] Thomas M Stace and Sean D Barrett. Error correction and degeneracy in surface codes suffering loss. *Physical Review A*, 81(2):022317, 2010. DOI: [10.1103/PhysRevA.81.022317](https://doi.org/10.1103/PhysRevA.81.022317).
- [56] T.M. Stace, S.D. Barrett, and A.C. Doherty. Thresholds for topological codes in the presence of loss. *Physical Review Letters*, 102(20):200501, 2009. DOI: [10.1103/PhysRevLett.102.200501](https://doi.org/10.1103/PhysRevLett.102.200501).
- [57] Robert Endre Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM (JACM)*, 22(2):215–225, 1975. DOI: [10.1145/321879.321884](https://doi.org/10.1145/321879.321884).
- [58] Robert Endre Tarjan. A class of algorithms which require nonlinear time to maintain disjoint sets. *Journal of computer and system sciences*, 18(2):110–127, 1979. DOI: [10.1016/0022-0000\(79\)90042-4](https://doi.org/10.1016/0022-0000(79)90042-4).
- [59] Giacomo Torlai and Roger G Melko. Neural decoder for topological codes. *Physical Review Letters*, 119(3):030501, 2017. DOI: [10.1103/PhysRevLett.119.030501](https://doi.org/10.1103/PhysRevLett.119.030501).
- [60] David K Tuckett, Stephen D Bartlett, and Steven T Flammia. Ultra-high error threshold for surface codes with biased noise. *arXiv preprint arXiv:1708.08474*, 2017. DOI: [10.1103/PhysRevLett.120.050505](https://doi.org/10.1103/PhysRevLett.120.050505).
- [61] Savvas Varsamopoulos, Ben Criger, and Koen Bertels. Decoding small surface codes with feedforward neural networks. *arXiv preprint arXiv:1705.00857*, 2017. DOI: [10.1088/2058-9565/aa955a](https://doi.org/10.1088/2058-9565/aa955a).
- [62] Chenyang Wang, Jim Harrington, and John Preskill. Confinement-higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory. *Annals of Physics*, 303(1):31–58, 2003. DOI: [10.1016/S0003-4916\(02\)00019-2](https://doi.org/10.1016/S0003-4916(02)00019-2).
- [63] David S Wang, Austin G Fowler, Charles D Hill, and Lloyd Christopher L Hollenberg. Graphical algorithms and threshold error rates for the 2d colour code. *Quantum Info. Comput.*, 10(9):780–802, sep 2009. ISSN 1533-7146. DOI: [10.5555/2011464.2011469](https://doi.org/10.5555/2011464.2011469).
- [64] Fern HE Watson, Hussain Anwar, and Dan E Browne. Fast fault-tolerant decoder for qubit and qudit surface codes. *Physical Review A*, 92(3):032309, 2015. DOI: [10.1103/PhysRevA.92.032309](https://doi.org/10.1103/PhysRevA.92.032309).
- [65] Adam C. Whiteside and Austin G. Fowler. Upper bound for loss in practical topological-cluster-state quantum computing. *Phys. Rev. A*, 90:052316, Nov 2014. DOI: [10.1103/PhysRevA.90.052316](https://doi.org/10.1103/PhysRevA.90.052316).
- [66] James Wootton. A simple decoder for topological codes. *Entropy*, 17(4):1946–1957, 2015. DOI: [10.3390/e17041946](https://doi.org/10.3390/e17041946).
- [67] James R Wootton and Daniel Loss. High threshold error correction for the surface code. *Physical review letters*, 109(16):160503, 2012. DOI: [10.1103/PhysRevLett.109.160503](https://doi.org/10.1103/PhysRevLett.109.160503).
- [68] Gilles Zémor. On cayley graphs, surface codes, and the limits of homological coding for quantum error correction. In *IWCC*, pages 259–273. Springer, 2009. DOI: [10.1007/978-3-642-01877-0_21](https://doi.org/10.1007/978-3-642-01877-0_21).

A Ackermann’s function

Different versions of Ackermann’s original function exist, though they generally only differ in a constant. In this section, we describe the version used by Tarjan to obtain an optimal upper bound on the complexity of Union-Find algorithms [57]. Ackermann’s function is a two-parameter function defined for any pair (i, j) of non-negative integers by the following relations,

$$\begin{aligned} A(0, i) &= 2i \\ A(i, 0) &= 0 \\ A(i, 1) &= 2 \\ A(i, j) &= A(i - 1, A(i, j - 1)) \end{aligned}$$

From this recursive definition, we find, for instance, that for all $j \geq 1$

$$A(1, j) = A(0, A(1, j - 1)) = 2A(1, j - 1) = 2^j.$$

A similar calculation shows that $A(2, 1) = 2$, $A(2, 2) = 4$ and

$$A(2, j) = 2^{2^{j-2}} \text{ with } j \text{ twos.}$$

We call the inverse of Ackermann’s function, which we denote $\alpha(n)$, the value

$$\alpha(n) = \min\{i \mid A(i, 4) \geq \log_2 n\}$$

Ackermann’s function grows amazingly quickly. The first terms to appear in the definition of α are

$$A(1, 4) = 2^4,$$

and

$$A(2, 4) = 2^{2^2} = 2^{16} = 65536.$$

Then, we obtain $A(3, 4) = A(2, A(3, 3)) = A(2, A(2, A(3, 2))) = A(2, A(2, 4))$ and using $A(3, 2) = 4$, we find

$$A(3, 4) = 2^{2^{2-2}} \text{ with } 65536 \text{ twos.}$$

This number is so large that we never hit $A(3, 4)$ in any practical situation, making $\alpha(n) \leq 3$.




Figure 6: (a) Threshold of the 2d toric code for Z -type error with no erasure. (b) Threshold of the 3d toric code for Z -type error with no erasure.




Figure 7: Below threshold performance for the 2d toric code with $p_e = 0.1$. Below the threshold the logical error rate is exponentially suppressed as the lattice size increases. We perform up to 10^9 Montecarlo trials per data point.

B Numerical results

Here we show an example of the threshold plots for the numerics that were used to generate the data shown in Figure 5. All data points were computed by repeated montecarlo simulations of random erasure and noise on a surface code followed by decoding. The simulations were repeated until at least 10,000 failures had been observed. Figure 6(a) shows the threshold for the 2d toric code, when $p_e = 0$, for which we find a threshold of 9.9%. Figure 6(b) shows the threshold for the 2+1d toric code, when $p_e = 0$, for which we find a threshold of 2.6%.

Figure 7 shows the below threshold scaling of the where $p_e = 0.1$. We see an exponential suppression in the logical error rate with increasing lattice size, strong evidence of the threshold behavior.