



Error-rate-agnostic decoding of topological stabilizer codes

Downloaded from: <https://research.chalmers.se>, 2025-12-04 22:45 UTC

Citation for the original published paper (version of record):

Hammar, K., Orekhov, A., Hybelius, P. et al (2022). Error-rate-agnostic decoding of topological stabilizer codes. Physical Review A, 105(4). <http://dx.doi.org/10.1103/PhysRevA.105.042616>

N.B. When citing this work, cite the original published paper.

Error-rate-agnostic decoding of topological stabilizer codes

Karl Hammar¹, Alexei Orekhov², Patrik Wallin Hybelius¹, Anna Katariina Wisakanto¹,
Basudha Srivastava¹, Anton Frisk Kockum², and Mats Granath^{1,*}

¹*Department of Physics, University of Gothenburg, 41296 Gothenburg, Sweden*

²*Department of Microtechnology and Nanoscience, Chalmers University of Technology, 41296 Gothenburg, Sweden*



(Received 15 December 2021; revised 21 March 2022; accepted 5 April 2022; published 25 April 2022)

Efficient high-performance decoding of topological stabilizer codes has the potential to crucially improve the balance between logical failure rates and the number and individual error rates of the constituent qubits. High-threshold maximum-likelihood decoders require an explicit error model for Pauli errors to decode a specific syndrome, whereas lower-threshold heuristic approaches such as minimum-weight matching are error agnostic. Here we consider an intermediate approach, formulating a decoder that depends on the bias, i.e., the relative probability of phase-flip to bit-flip errors, but is agnostic to error rate. Our decoder is based on counting the number and effective weight of the most likely error chains in each equivalence class of a given syndrome. We use Metropolis-based Monte Carlo sampling to explore the space of error chains and find unique chains that are efficiently identified using a hash table. Using the error-rate invariance, the decoder can sample chains effectively at an error rate which is higher than the physical error rate and without the need for thermalization between chains in different equivalence classes. Applied to the surface code and the $XZZX$ code, the decoder matches maximum-likelihood decoders for moderate code sizes or low error rates. We anticipate that, because of the compressed information content per syndrome, it can be taken full advantage of in combination with machine-learning methods to extrapolate Monte Carlo-generated data.

DOI: [10.1103/PhysRevA.105.042616](https://doi.org/10.1103/PhysRevA.105.042616)

I. INTRODUCTION

Quantum decoherence is one of the major challenges that has to be overcome in building a quantum computer [1–6]. A prominent line of research to address this issue focuses on topological stabilizer codes implemented on low-connectivity lattices of qubits [7–11], with small stabilizer codes presently being realized experimentally [12–24]. The topological codes provide protection against errors by encoding the quantum information into entangled states that are removed from each other by a code distance d , corresponding to the minimal number of local qubit operations required for a logical operation. These logical operators act as effective Pauli operators on the logical code space. Local stabilizers in the form of local Pauli operations commute with the logical operators; measuring these stabilizers provides a syndrome of the possible error configurations, so-called error chains, that are affecting the code. The error chains can be grouped into equivalence classes depending on whether they commute with a given representation of the logical operators or not.

To decide which corrective actions to take given a measured syndrome, there is a range of different decoders. A maximum-likelihood decoder (MLD) aims directly at the core problem of decoding a stabilizer code, which is to identify the most likely equivalence class of error chains that correspond to the syndrome, thus specifying a correction chain (or class of chains) least likely to cause a logical error [9]. Ideally, an MLD will be optimal in the sense of providing the highest possible logical fidelity given the type of stabilizer code and the error rate of individual qubits. Consequently, it will also provide the highest possible threshold for the code [9,25–27], corresponding to the error rate below which the logical failure rate can always be decreased by increasing the code size.

In contrast to maximum-likelihood decoding, there are heuristic decoders that suggest the correction chain based on a rule or algorithm which is not guaranteed to fall in the most likely equivalence class. The standard such algorithm is minimum-weight perfect matching (MWPM) [9,28–36], which is based on finding a most likely error chain by mapping the syndrome to a graph problem with weighted edges. Other examples are based on addressing the syndrome successively from small to large scale, using the renormalization group [37], cellular-automaton [38,39], or union-find algorithm [40,41]. Such heuristic decoders generally have higher logical failure rates and lower error thresholds than MLDs. Nevertheless, because of the computational complexity of MLDs, nonoptimal decoders may be preferable, or even necessary, for practical purposes [11].

Following the recent transformative developments in deep learning [42,43], with applications in quantum physics

*mats.granath@physics.gu.se

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. Funded by [Bibsam](https://www.bibsam.se/).

[44–48], decoders based on machine-learning algorithms have also been developed. Approaches using deep reinforcement learning [49–53] or evolutionary algorithms [54] have the appealing property that they can learn by exploring a model of the errors and measurements of the physical code, but a disadvantage is that they are difficult to train and to scale to large code distances. Deep-learning approaches based on supervised learning [55–62] have the potential for fast and scalable decoding. However, supervised learning is based on training data, which has to be generated using a reference decoder, and the type and quality of the data will be crucial. Part of the objective of the present work is to generate data in a form which is compact enough for deep learning while still containing information that allows for high performance and flexibility with respect to error rate.

Two types of MLDs for general noise models are the Metropolis-based Markov-chain Monte Carlo (MCMC) algorithm by Wootton and co-workers [63,64] and the matrix-product-states (MPS) decoder by Bravyi *et al.* [65]. Apart from the high time complexity, a limitation of these decoders is that they are fixed to the noise model and error rate, i.e., the specific single-qubit error rates p_x , p_y , and p_z for X , Y , and Z Pauli errors, with $p = p_x + p_y + p_z$ the total error rate. Given a syndrome s consistent with a set of chains $\{C\}_s$, the probability P_E that this is caused by an error chain in equivalence class E is proportional to the sum of the (*a priori*) probabilities of the chains π_C , i.e., $P_E \sim \sum_{C \in E} \pi_C$. The MCMC decoder is based on Metropolis sampling of the error chains, while the MPS decoder performs an efficient truncated summation of the probabilities. For both of these methods, a statement about most likely equivalence class for a given syndrome does not generalize beyond the error rates for which the calculation was performed.

In this article, we explore a decoder based on a highly reduced count of the error chains. We only consider the most likely set of chains and their degeneracy in each equivalence class, in order to approximate the probability of the equivalence class. Defining an effective weight w , which takes into account the relative probability of X , Y , and Z errors, the set of most likely chains in each class is independent of the overall error rate. The decoder is in this sense error-rate agnostic. For codes with a single logical qubit and four corresponding equivalence classes of error chains, each syndrome is characterized by an octet of numbers corresponding to weight and degeneracy of the most likely (lightest) chains in each class. In contrast, an ideal MWPM decoder would identify one lightest chain among all the equivalence classes, whereas the MCMC decoder would take into account a large number of error chains with varying weights. We find that for moderate code distances, the limited information of weight and degeneracy is sufficient to provide near-optimal decoding, matching results from the MCMC MLD for different noise bias for two standard topological stabilizer codes.

We anticipate two potential uses for this effective weight and degeneracy (EWD) decoder. The first is as a stand-alone near-optimal decoder. While using Metropolis sampling to find high probability chains, the decoder has comparatively faster convergence than the MCMC decoder as equivalence classes can be explored separately and at a higher error rate than the physical error rate. The second use is to provide

high-information-content data for supervised deep-learning-based decoders that may allow for flexibility with respect to the overall error rate.

This article is organized as follows. We start in Sec. II with a short review of the two stabilizer codes that we will use to benchmark our decoder. In Sec. III we describe the basic formulation of the decoder. After this, in Sec. IV we show some test cases and compare to the MCMC decoder, the MPS decoder, as well as exact analytical results. This is followed by a discussion in Sec. V and a summary in Sec. VI.

II. TOPOLOGICAL STABILIZER CODES

To investigate the performance of the EWD decoder, we consider two different stabilizer codes: the rotated surface code and the $XZZX$ code. The surface code [7–9] is of the Calderbank-Shor-Steane type with stabilizers containing only one type of Pauli operators [4,5]. (Formally, these are the generators of the stabilizer group. Any product of the generators is also a stabilizer.) The rotated configuration reduces the number of qubits without reducing the code distance [66,67]. Recently it was pointed out [68] that a modified version of the code [69,70], with mixed $XZZX$ stabilizers, has a very high threshold for biased noise and also allows for efficient decoding. For depolarizing noise the two codes are equivalent, but for Z - or X -biased noise they have very different properties; whereas the threshold decreases with bias for the surface code, it increases for the $XZZX$ code.

As illustrated in Fig 1, both codes consist of a square grid of $N = d^2$ qubits on which $N - 1$ local (two- or four-qubit) stabilizers act. The stabilizers commute and square to identity, which means that they split the Hilbert space into size-2 blocks specified by the parity (± 1 eigenvalues) with respect to each stabilizer. The logical code space, $|0\rangle_L$ and $|1\rangle_L$, is taken to be the all-even-parity block. Measuring the stabilizers projects the state onto one of the blocks and any odd-parity measurement is an indication that the state has been affected by errors and escaped the code space. The set of odd-parity measurements with their coordinates on the grid is referred to as the syndrome.

Assuming Pauli noise, each qubit may be subject to one of three errors X , Z , or $Y = XZ$ (up to an overall phase), occurring with rates p_x , p_y , and p_z , respectively, such that the total error rate is $p = p_x + p_y + p_z$, defined per measurement cycle of the syndrome. This type of error model gives a standard, approximate, representation of qubit decay and dephasing [71,72]. In addition, in practice, the measurement of the stabilizers requires additional ancilla qubits that are entangled with the (data) qubits. As these ancillas are also noisy and the measurement circuit involves several two-qubit gates, there will be errors in the syndrome readout, which complicates the decoding significantly [29,73]. Also, in standard qubit architectures, leakage out of the qubit space leads to additional complications, as do possible correlated errors affecting two or more qubits [32,74,75]. Nevertheless, to benchmark and discuss the basic features of the decoder we will use the basic error model and assume perfect syndromes.

Logical operators act within the code space, thus commuting with the stabilizers. For the surface code they are a string of Z 's between two edges (Z_L) and a string of X 's between the




FIG. 1. Illustrations of the two stabilizer codes used to benchmark our EWD decoder. (a) Rotated surface code for $d = 5$, showing weight-4 and weight-2 stabilizers acting on a grid of d^2 physical qubits (circles). The logical operators used to define the equivalence classes of error chains are weight-5 strings: X_L , on the left edge, and Z_L , on the top edge (not shown). (b) The XZZX code for $d = 5$. All weight-4 stabilizers are equivalent, whereas the weight-2 boundary stabilizers are complementary to the adjacent weight-4 stabilizer. The logical X_L used to define the equivalence classes is shown on the left edge, with a corresponding mixed Z_L operator (not shown) on the top edge.

other two edges (X_L). For XZZX they are in general mixed operators, except on the diagonals.

An arbitrary configuration of errors, or error chain $C \in \{I, X, Y, Z\}^{\otimes N}$, is specified by an element in $\{I, X, Y, Z\}$ acting on each qubit, where I is the identity operator. Each C has a number of nonidentity operators, quantified by n_x , n_y , and n_z , and gives a unique syndrome $s(C)$. It is a mapping of many to

one: A chain gives a unique syndrome, but a syndrome does not give a unique chain. The *a priori* probability of chain C , i.e., before measuring a syndrome, is given by

$$\pi_C = (1 - p)^N \pi_C^r, \quad (1)$$

where

$$\pi_C^r = \left(\frac{p_z}{1 - p} \right)^{n_z} \left(\frac{p_x}{1 - p} \right)^{n_x} \left(\frac{p_y}{1 - p} \right)^{n_y} \quad (2)$$

is the probability relative to the empty chain.

For an arbitrary syndrome s , consider an arbitrary chain C consistent with that syndrome. The equivalence class $E(C)$ of the chain C is the set of all 2^{N-1} chains generated by acting with all elements of the stabilizer group. (A general element of the stabilizer group consists of a product of stabilizers or the identity operator acting on the $N - 1$ faces of the code.) Three other equivalence classes with an equal number of chains are generated by acting with one of the logical operators X_L , Z_L , or $Y_L = X_L Z_L$. The equivalence class of the chain $E \in \{1, X, Y, Z\}$ is then decided by whether it commutes with X_L and/or Z_L .¹

The product of any two chains with the same syndrome gives a chain with a trivial syndrome, but unless the two chains are in the same equivalence class, the product will correspond to a logical operator. Thus, maximum-likelihood decoding corresponds to providing an arbitrary correction chain in the most likely equivalence class.

III. MODEL AND DECODER

A. Effective weight parametrization

As discussed in Sec. II, we only consider single-qubit Pauli errors with error rates $p_x + p_y + p_z = p$ and assume noise-free stabilizer measurements. We also assume Z-biased noise $p_z \geq p_x, p_y$. After an ideal stabilizer measurement cycle, the probability of a particular error chain with n_x , n_y , and n_z errors is (relative to the empty chain) given by Eq. (2). It is convenient to rewrite this in terms of the most likely error, using the relative error rate $\tilde{p}_z \equiv \frac{p_z}{1-p}$, as

$$\pi^r = (\tilde{p}_z)^w = e^{-\beta w}, \quad (3)$$

with $\beta = -\ln \tilde{p}_z$ taking the role of an inverse temperature. Here we have defined the effective weight of the chain as

$$w = n_z + \alpha_x n_x + \alpha_y n_y, \quad (4)$$

¹In order to uniquely specify the equivalence class, it is necessary to choose a specific representation of the logical operators. As the error chain does not commute with all stabilizers, deforming a logical operator may change the equivalence class classification. For the surface code we choose X_L as the product of X operators on the left edge and Z_L as the product of Z operators on the top edge such that the class of the chain can be identified from the parity of Z errors on the left edge and the parity of X errors on the top edge, and similarly for the XZZX code, where the operators on the edges are mixed.

with α_i given by

$$\left(\frac{p_z}{1-p}\right)^{\alpha_x} = \left(\frac{p_x}{1-p}\right), \quad (5)$$

$$\left(\frac{p_z}{1-p}\right)^{\alpha_y} = \left(\frac{p_y}{1-p}\right). \quad (6)$$

The set of parameters $(\tilde{p}_z, \alpha_x, \alpha_y)$ provides a convenient parametrization of the space of error rates (p_x, p_y, p_z) , with $p_x + p_y + p_z = p \leq 1$. For fixed α_x and α_y , which includes standard error models such as depolarizing, uncorrelated, and pure phase- or bit-flip noise, the single scalar w specifies the probability of an error chain according to Eq. (3). For this reason, the representation is also used for enhanced minimum-weight-matching decoders, where the weight of edges corresponds to their relative probability [36,68].

Limiting to $p_x = p_y$ with $\alpha \equiv \alpha_x = \alpha_y$, the total error rate p is related to \tilde{p}_z by $\tilde{p} \equiv \tilde{p}_z + 2\tilde{p}_z^\alpha$, with $\tilde{p} \equiv \frac{p}{1-p}$. Depolarizing noise corresponds to $\alpha = 1$, such that $w = n_x + n_y + n_z$, while pure phase-flip noise (for $p \leq 0.5$) corresponds to $\alpha = \infty$, such that any chain with an X or Y error would have infinite weight, consistent with their vanishing probability. The parametrization for Z -biased noise in [67,68,76] in terms of fixed $\eta = \frac{p_z}{p_x + p_y}$ and $p_x = p_y$ corresponds to fixed α only for depolarizing noise $\eta = \frac{1}{2}$ and pure phase-flip noise $\eta = \infty$, while in general there is a mapping

$$\alpha = \frac{\ln \tilde{p}_z - \ln 2\eta}{\ln \tilde{p}_z}. \quad (7)$$

Which parametrization to use for a single bias and error rate is a matter of convenience, but the advantage of the α parametrization is that the weight of a chain is conserved under changes of the overall error rate. (A disadvantage could be that the mapping between total and individual error rates is less direct.) In practice, this means that for a given syndrome, the same lightest chains will most likely be independent of error rate.

For reference, uncorrelated noise, i.e., independent bit-flip and phase-flip errors, would correspond to $\alpha_x = 1$ and $\alpha_y = 2$. Note, however, that for uncorrelated noise, p is conventionally defined as the independent rate of X or Z errors. The definition of effective length should still reflect the relative error rates, such that Eq. (6) is replaced by

$$\left[\frac{p(1-p)}{(1-p)^2}\right]^{\alpha_y} = \frac{p^2}{(1-p)^2}. \quad (8)$$

B. Effective weight and degeneracy decoder

1. Calculating the probabilities of equivalence classes

The probability of an equivalence class E of a given syndrome s is simply the probability of all chains in the class normalized with respect to all chains that are consistent with the syndrome

$$P_E = \frac{\sum_{C \in E} \pi_C}{\sum_{C: s(C)=s} \pi_C} = \frac{\sum_w N_E(w) e^{-\beta w}}{\sum_{E', w} N_{E'}(w) e^{-\beta w}}, \quad (9)$$

where $N_E(w)$ is the total number of unique chains in class E with effective weight w . The numbers $N_E(w)$ are combinatorial; they depend on the syndrome, but are independent of

the overall error rate given by β , provided the relative error rates set by α are fixed. In principle, $N_E(w)$ can be calculated by finding one chain, using, e.g., minimum-weight matching, acting on this chain with the full stabilizer group and, to switch equivalence class, the logical operators.

Since the number of elements in the stabilizer group grows exponentially with the number of qubits N , it quickly becomes infeasible to find or store the complete set of values. However, in the case that the partition function Z_E of each equivalence class is dominated by the contribution of the lightest chains, we expect to obtain a good approximation by having an explicit count of only those chains

$$Z_E = \sum_w N_E(w) e^{-\beta w} \approx N_E^* e^{-\beta w_E^*}, \quad (10)$$

where w_E^* is the effective weight of lightest chain(s) in the class and $N_E^* = N_E(w_E^*)$ is the number of such chains. Thus, the probability of an equivalence class $P_E = Z_E / \sum_{E'} Z_{E'}$ according to this approximation requires knowledge of an octet of values

$$\{w_E^*, N_E^*\}_{E=\{1,X,Z,Y\}}. \quad (11)$$

2. Algorithm

The EWD algorithm, which uses the above approximation, is outlined in pseudocode form in Algorithm 1, using a hash table to store and identify unique chains. The Metropolis update rule (for more detail, see Sec. V) that sets the probability of accepting a new error chain is operated at an error rate p_{sample} , which may be different from the physical error rate p . Since the weight of an error chain is set by the bias α , independently of p , chains can be generated and stored at an arbitrary error rate. Sampling chains at an appropriately chosen (typically high) error rate has the advantage of avoiding the algorithm becoming stuck in a low-weight configuration. The sampling error rate is effectively a hyperparameter of the decoder algorithm, which should be optimized such that the logical failure rate is minimized. We remark that in benchmarking the decoder using randomly generated error chains, it is of course important that the starting chains in the four equivalence classes are well randomized to ensure that it is not possible to underestimate the failure rate.

A variation of the EWD algorithm, referred to as “All” in Sec. IV, evaluates the probability based on all unique observed chains, i.e., with $Z_E = \sum_w N_E(w) e^{-\beta w}$. If keeping the information content per syndrome to a minimum is not a priority, this version of the algorithm is more accurate, at least for noninteger α .

3. Limitations of applicability

Formally, the requirement for the approximation of Eq. (10) to hold is that the probability density for chains of length w , $N_E(w) e^{-\beta w}$, is monotonically decreasing with chain weight. Otherwise the probability of each equivalence class will not be dominated by the lightest chains, which invalidates the basic assumption of the decoder. In other words, it should satisfy

$$N_E(w + \delta w) e^{-\beta \delta w} < N_E(w) \quad (12)$$

Algorithm 1. The EWD decoder

Algorithm 1: The EWD decoder

```

1 if Syndrome  $s$  is nontrivial then
2   foreach equivalence class  $E$  do
3     Create empty hash table  $T_E$ ;
4     Get an error chain  $C \in E$  consistent with  $s$ ;
5     Calculate the weight of  $C$ :  $w$ ;
6     Calculate a hash of  $C$ :  $H$ ;
7     Store  $w$  in  $T_E[H]$ ;
8   for  $N$  iterations do
9     Apply a random stabilizer to  $C$  to get a new
       chain  $C'$ ;
10    Calculate the weight of  $C'$ :  $w'$ ;
11    Use  $w$  and  $w'$  to calculate the relative
       probability  $a = \pi_{C'}/\pi_C$  using the error rate
        $p_{\text{sample}}$ ;
12    Generate a uniform random number
        $u \sim U(0, 1)$ ;
13    if  $a > u$  then
14       $C \leftarrow C'$ ;
15       $w \leftarrow w'$ ;
16      Calculate a hash of  $C$ :  $H$ ;
17      Store  $w$  in  $T_E[H]$ ;
18    end
19  end
20  Find lowest weight  $w_E^*$  in  $T_E$  and count its
       multiplicity  $N_E^*$ ;
21 end
22 Output probability of each equivalence class
   according to  $P_E = Z_E / \sum_{E'} Z_{E'}$  with
    $Z_E = N_E^* e^{-\beta w_E^*}$  using the physical error rate  $p$ ;
23 Apply an arbitrary correction chain from the
   equivalence class with the highest  $P_E$  on the
   physical qubits;
24 end

```

for any $w \geq w_E^*$, with δw the shortest difference between effective weights. Correspondingly, for large w , $\frac{d \ln N_E}{dw} \lesssim \beta$. Generally, the number of chains grows exponentially with the length (for $w \ll d^2$) as (averaged over syndromes and equivalence classes) $\sum_E N_E \sim d^{aw}$, where we have found that the parameter $a > 0$ depends on type of code and error model, but only weakly on error rate. The latter dependence follows from the changes in the distribution of syndromes with error rate. This gives a corresponding upper bound on d versus error rate, as $d \lesssim e^{\beta/a}$ for the approximation of estimating the partition function using only the lightest chains.

In practice, going beyond this limit to larger code distances or larger error rates (smaller β), the EWD decoder will perform worse than optimal. In this regime, the distribution $N_E(w)e^{-\beta w}$ will be centered on heavier chains such that evaluating the partition function based on the lightest chains will be an increasingly poor approximation with increasing d or increasing p . We have not explored the limitations in detail, but the relatively sharp threshold presented for depolarizing noise (see Sec. IV) indicates that the decoder gives close to optimal results for $d < 20$. For sufficiently low error rates (depending on d), the decoder will be optimal, as the lightest chains will dominate asymptotically.

For general noninteger α , it is less clear how to motivate the approximation, as δw , the difference between effective weights, can become arbitrarily small. Nevertheless, as discussed in Sec. IV below, even though the approximation may be quantitatively quite poor for specific syndromes, aggregated over many syndromes, the effect of this discrepancy is found to be small.

IV. RESULTS**A. Test cases**

To demonstrate the capacity of the EWD decoder we apply it to several test cases, depolarizing noise ($\alpha = 1$), moderately biased noise ($\alpha = 5$), and highly biased noise ($\alpha = 10\,000$), focusing primarily on the XZZX code. For all these cases, the Metropolis sampling is done at a high error rate ($p_{\text{sample}} = 0.3$) compared to the error rate at which the syndromes are generated. The decoder is run for $25d^5$ Metropolis steps in each equivalence class. Due to the long correlation length, only every fifth step is considered and added to the set of observed chains throughout all tests. Unless otherwise stated, equivalence classes are seeded with a random chain consistent with that syndrome and class. For each configuration of d and p , syndromes are randomly generated and solved in order to estimate the logical failure rate P_f . The latter is the fraction of all generated syndromes for which the randomly generated chain that provides the syndrome is not in the most likely equivalence class predicted by the decoder for that syndrome. For $d \leq 11$, 10 000 syndromes are solved; for $d = 13$, at least 4000; and for $d = 15$, at least 2000. Error bars indicate one standard deviation based on the number of sampled syndromes and the mean logical failure rate.

As shown in Fig. 2, we find an approximate threshold for depolarizing noise which is close to that of the MLDs for perfect syndrome measurements: $p_{\text{th}} \approx 18.5\%$ [63–65]. Note that for depolarizing noise, the XZZX code and the rotated surface code are equivalent. For larger bias $\alpha = 5$, we find good correspondence with the results using the MPS decoder as implemented in Ref. [68,77], with a threshold near $p \approx 30\%$, as shown in Fig. 3(a). At $p = 0.30$ and $\alpha = 5$, corresponding to $\eta = 18.3$, the hashing bound [78] is in fact $p_{\text{hash}} = 0.307$, confirming the exceptional nature of the XZZX code.

For very large bias $\alpha = 10\,000$, results are expected to be very close to those for $\alpha = \infty$, i.e., pure Z noise, with a threshold of 50%. For the latter, the following exact expression follows from the fact that there is only a single pure Z operator (on one diagonal) that commutes with all stabilizers. It is not part of the stabilizer group, i.e., it is a logical operator, which implies that for every syndrome there are only two possible error chains in two different equivalence classes. Consequently, maximum-likelihood decoding implies that the decoder will suggest the lighter of the two and will always fail when the heavier chain occurs. For odd d (and $p \leq 0.5$), the expression is

$$P_{f,\alpha=\infty}(p) = \sum_{w=\lceil d/2 \rceil}^d \binom{d}{w} p^w (1-p)^{d-w}. \quad (13)$$

Figure 3(b) shows good correspondence between the numerical data from the EWD decoder and this expression.

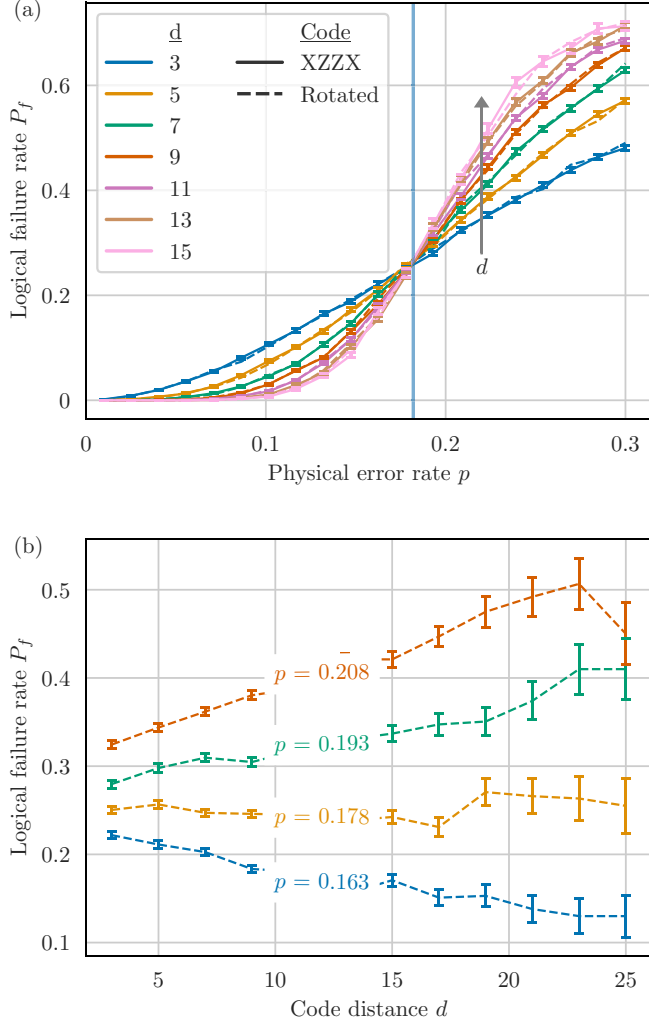


FIG. 2. Performance of the EWD decoder for depolarizing noise. (a) Logical failure rate P_f as a function of physical error rate p for depolarizing noise ($\alpha = 1$) for the XZZX code and the rotated surface code using the EWD decoder on several different code distances d . The vertical blue line marks the approximate threshold at 18.2 %. (b) Logical failure rate as a function of code distance for error rates close to the threshold for depolarizing noise on the XZZX code. The transition from decreasing to increasing failure rate with increasing code distance indicates a threshold near an error rate $p \approx 18\%$. The number of syndromes for $d > 15$ is at least 200.

For the rotated surface code, we show in Fig. 2 results for depolarizing noise that, as expected, coincide with those for the XZZX code. In Fig. 4 we show results for $\alpha = 5$, compared to the MPS decoder. As expected for phase-biased noise on the surface code and contrary to the XZZX code, the threshold is reduced compared to depolarizing noise.

B. Specific syndrome example

1. Integer weights

To better illustrate the workings and performance of the EWD decoder, it is useful to not only consider aggregate results on the logical failure rates over many syndromes, but to also look at individual syndromes. The example syndrome

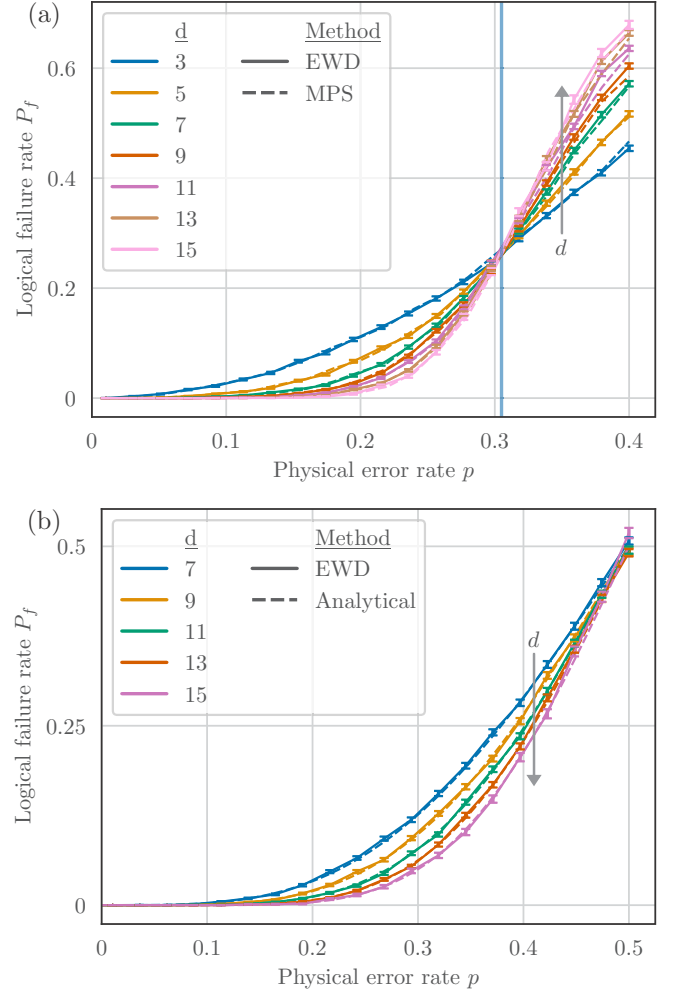


FIG. 3. Performance of the EWD decoder for biased noise on the XZZX code. The logical failure rate P_f is plotted as a function of physical error rate p for (a) moderately biased noise $\alpha = 5$ and (b) large bias $\alpha = 10000$ compared to (a) data obtained using the MPS decoder as implemented in [77], with the vertical line indicating the threshold at 30.5 %, and (b) exact results [Eq. (13)] for pure phase-flip noise $\alpha = \infty$.

shown in the inset of Fig. 5(a) has the interesting property that the most likely equivalence class may (depending on the noise bias) change as the error rate is increased. We consider a noise bias $\alpha = 2$, where the weight of an error chain depends on the number of X, Y, and Z errors according to $w = n_z + 2(n_x + n_y)$.

Figure 5(a) shows the number of unique chains observed through Metropolis sampling at an error rate $p_{\text{sample}} = 0.3$ for 25×5^5 Metropolis steps in each equivalence class starting from a randomized initial chain in each class. There is a single minimum-weight chain in class Z ($w_Z^* = 10, N_Z^* = 1$), which will dominate the probability distribution at low error rates, while at larger error rates the larger number of one-unit heavier chains in class I ($w_I^* = 11, N_I^* = 12$) will be more important to decide most likely equivalence class, as seen in Fig. 5(b). Here we also see that using only the most likely chains in each class gives results that are in close correspondence with those given by evaluating P_E based on all unique




FIG. 4. Performance of the EWD decoder for moderately biased noise $\alpha = 5$ on the rotated surface code. The logical failure rate is plotted as a function of physical error rate, obtained with the EWD and MPS decoders for several code distances. The error rate used for sampling was $p_{\text{sample}} = 0.3$ for all code distances except $d = \{13, 15\}$, which were sampled at $p_{\text{sample}} = 0.15$. The vertical line indicates the approximate threshold at $p = 0.101$.

observed chains, as well as by using the MCMC decoder of Ref. [63].

As is clear from Fig. 5(b), the data allow for the evaluation of the equivalence class probabilities for any error rate p . In this sense, the decoder is error-rate agnostic. Notably, whereas the EWD curves are smooth (being generated from the same data), the MCMC results have large variations with error rate. For the latter, each error rate corresponds to generating a different Markov chain, and even though convergence criteria are applied on the most likely class, we find substantial variations in class probabilities. Even aggregated over 10 000 syndromes, the fluctuations are visible.

2. Noninteger weights

The previous example considered biased noise with $\alpha = 2$, i.e., an integer value, which implies that chain weights come in well-separated integer steps. Given limited knowledge of individual qubit noise biases, it may be reasonable to assume depolarizing noise or some other integer α bias, especially if the bias is large. Nevertheless, it is interesting to explore the consequences of noninteger α . Figure 6(a) shows the distribution of unique chains for the same syndrome [inset of Fig. 5(a)] for an arbitrary fraction, $\alpha = 1.873$, close to $\alpha = 2$. As expected, the bars of equal-weight chains are now fractionalized, i.e., fewer chains now have the same weight.

As shown in Fig. 6(b), the EWD algorithm still gives the correct equivalence class for low error rates and it qualitatively captures the transition between class Z and class I, but at an incorrect error rate (outside of the view of the figure). In fact, for this syndrome, using only the most likely chains in each class, it is clearly a better approximation to round α to $\alpha = 2$, as can be seen by comparing Figs. 5(b) and 6(b). However, as seen from Fig. 7, which shows the logical failure rate averaged over 10 000 randomly generated syndromes at




FIG. 5. Statistics of error chains obtained for a specific syndrome. (a) Number of unique error chains $N_E(w)$ for biased noise $\alpha = 2$, as a function of weight w and equivalence class $E = \{I, X, Y, Z\}$ for the syndrome on the $d = 5$ XZZX code shown in the inset. The data are obtained from Metropolis sampling of error chains at an error rate of $p_{\text{sample}} = 0.3$ for 25×5^5 Metropolis steps. Nonstriped bars show the lowest weights in each equivalence class. Chains with $w > 13$ are not shown. (b) Probability of equivalence class as a function of error rate for the syndrome in (a). Two different methods are used to evaluate Eq. (9) from the observed unique chains $N_E(w)$, using only the most likely chains in each equivalence class (EWD) and using all observed chains (All), compared to the MCMC decoder, which is not based on identifying unique chains.

$\alpha = 1.873$, decoded using the EWD at $\alpha = 1.873$ and 2, together with decoding using all identified chains, the difference at the aggregate level is small. The syndrome considered here as an example is thus an outlier.

As argued in Sec. III, it is only when it is important to have a compact representation of the equivalence class probabilities for a syndrome (such as for possible deep-learning purposes) that the EWD decoder using only the lowest-weight chains serves a potential purpose. If this is not the intent, there is no reason not to use all observed unique chains, regardless of




FIG. 6. Statistics of error chains obtained for the syndrome in the inset of Fig. 5(a) for a noninteger noise bias. (a) Number of unique error chains $N_E(w)$ for biased noise $\alpha = 1.873$, as a function of weight w and equivalence class $E = \{I, X, Y, Z\}$. Nonstriped bars show the lowest weights in each equivalence class. (b) Probability of equivalence class as a function of error rate using data from (a). The methods are the same as in Fig. 5(b).

weight, to evaluate the relative class probabilities using the first expression in Eq. (10).

C. Decoding efficiency

The number of stabilizers grows exponentially with the code size d as 2^{d^2} such that a brute force decoder aiming to generate all error chains of a given syndrome will have an exponential time complexity $t \sim O(2^{d^2})$. As discussed in Refs. [63,64], the time complexity of the MCMC decoder using parallel tempering is superpolynomial in d . It should be noted that this measure is based on convergence of the algorithm, implying close to optimal maximum-likelihood decoding. If the aim only is to improve on heuristic methods such as MWPM, the runtime can scale significantly better with d . The single-temperature (ST) algorithm formulated in




FIG. 7. Comparison of decoding accuracy of different methods for noninteger bias $\alpha = 1.873$ on the $XZZX$ code. The methods used are the EWD decoder (i.e., accounting for only the lightest chains in each class) with $\alpha = 1.873$ and 2, weighing in all observed chains in each class (All), and the parallel-tempering (MCMC) decoder. The inset shows relative variations, which do not display any systematic differences. The data are obtained by averaging over 10 000 syndromes for $d = 5$.

Ref. [64] is based on estimating the partition function (10) in terms of the expectation value of the chain weight over a fixed number of Metropolis steps number in each class, $\langle w \rangle_{\beta, E}$. This was found to significantly improve on MWPM failure rates at the cost of an extra $O(d^2)$ factor for the runtime.

To study the time complexity of the EWD decoder we have considered two different regimes for depolarizing noise, of high error rates $p = 0.15$ and asymptotically low error rates. For high error rates we compare the method to the MCMC decoder using parallel tempering and to the aforementioned ST decoder. As shown in Fig. 8, the convergence for moderate code sizes is significantly more rapid than the MCMC decoder, which we ascribe to the overhead of the latter from parallel tempering. The ST algorithm, although it has similar improved convergence as the EWD decoder, converges to a worse logical failure rate. In fact, the ST algorithm has a suboptimal threshold for depolarizing noise (inset of Fig. 7 in Ref. [64]) of between 15% and 16%.

For low error rates we have studied the time (in terms of the number of Metropolis steps) required to identify a weight $w_{\min} = (d - 1)/2$ chain for syndromes for which these are the shortest possible chains. In the limit $p \rightarrow 0$, the shortest chains that can fail (given the MLD) have weight $(d + 1)/2$, so establishing the time required to find weight- w_{\min} chains gives a good measure of the time complexity for low error rates. Specifically, we generate random error chains of weight $w = (d - 1)/2$, deform them by a large number of random stabilizers and logical operators to produce heavy chains in all four equivalence classes, and then let the decoder find unique chains in the four classes, stopping when a weight- $w = (d - 1)/2$ chain is identified. The results are plotted in Fig. 9, showing an approximate scaling $t \lesssim O(d^5)$. However, there is an indication of superpolynomial behavior as d is




FIG. 8. Convergence of the logical failure rate as a function of the number of Metropolis steps executed on a fixed set of 10 000 syndromes for three different decoders: our EWD decoder, the MCMC decoder [63,64], and the single-temperature Monte Carlo-based decoder of Ref. [64]. The number of attempts is normalized with the number of equivalence classes (four) for the EWD and ST decoders and with the number of parallel tempering levels (seven) for the MCMC decoder. All algorithms are seeded with the enhanced MWPM solutions [64] based on Blossom V [79], giving an initial logical failure rate of approximately 0.23. The syndromes are generated on the standard (nonrotated) surface code with code distance $d = 7$, subject to depolarizing noise ($\alpha = 1$) with error rate $p = 0.15$. The sampling error rate for the EWD decoder is set to $p_{\text{sample}} = 0.25$. The inset shows the early convergence (using a different random set of 20 000 syndromes).

increased, which may be consistent with the complexity discussed in Ref. [63] for the MCMC decoder.

The inset in Fig. 9 gives the distribution of the number of steps required for $d = 9$, as an example, showing that for most syndromes the lightest chain is easily found, while there is a long tail of rarer syndromes for which it takes the decoder significantly longer. The standard implementation of MWPM using the Blossom algorithm on a complete syndrome graph has a time complexity $O(d^6 \log d)$, which can be reduced significantly in practice, for a realistic set of syndrome graphs, and by reducing the number of neighbors, to $O(d^2)$ [80,81].

It is quite clear that the EWD decoder is not an efficient decoder compared to heuristic decoders such as MWPM and union find. Nevertheless, as discussed in Sec. VB below (see Table I), since it is a maximum-likelihood decoder, the EWD decoder will provide lower logical failure rates.

V. DISCUSSION

A. Comparison with parallel-tempering-based decoders

As the implementation of the EWD decoder uses Metropolis sampling of error chains similarly to the MCMC decoder formulated in [63,64], a closer discussion of the similarities to and differences from these decoders is motivated. The probability of an equivalence class is the probability of an error chain in the class given a syndrome s according to Eq. (9). Using the Metropolis algorithm, a Markov chain of




FIG. 9. Estimated time complexity of the EWD decoder at asymptotically low error rate based on the number of Metropolis steps (t) required to identify a weight $w = (d - 1)/2$ error chain for randomly generated syndromes that are guaranteed to have a chain with such weight. Shown is the maximum number of steps required to find 95% and 50% of the lowest weight chains and fitted to single exponents. The inset exemplifies the data for $d = 9$ (containing 2×10^6 syndromes) showing the number of syndromes versus the number of steps the decoder takes to find the weight $(d - 1)/2$ chain.

$M \ll 2^{N+1}$ (4 times the number of chains in each class) error chains $\{C_i\}$, $i = 1, \dots, M$, representing the true distribution

$$P_C = \frac{\pi_C}{\sum_{C: S(C)=s} \pi_C}, \quad (14)$$

is generated. The probability $P_E = \sum_{C \in E} P_C$ of the equivalence class E is thus approximated by the relative number of observed chains in E :

$$P_E \approx \sum_{i: C_i \in E} \frac{1}{M}. \quad (15)$$

The Markov chain is generated by acting with random stabilizers that deform the error chain $C \rightarrow C'$ without changing the syndrome, with the new chain accepted in standard fashion with probability $\min(1, \pi_{C'}/\pi_C)$. However, stabilizers will not change the equivalence class, so in addition it is also necessary to act with logical operators to generate chains in all equivalence classes. Whereas the stabilizers are weight-4 (or

TABLE I. Fraction of failed chains with weight $w = (d + 1)/2$ for the XZZX code and depolarizing noise, comparing MWPM (using [80]) and EWD. The relative fraction corresponds to the relative asymptotic ($p \rightarrow 0$) logical failure rate for the two decoders. Enough chains are sampled such that the digits specified are precise within one standard deviation.

d	MWPM	EWD
5	0.075	0.040
7	0.0086	0.0028
9	0.00073	0.00018

-2) operators, the logical operators are weight d , which makes these exponentially (with d) less likely to be accepted. A naive implementation would thus require exponential computational time to generate a Markov chain which is appropriately balanced between the equivalence classes.

In order to deal with this problem, parallel tempering [82] is employed in [63,64], in which chains are generated and exchanged between several layers of the stabilizer code with different error probabilities. The probability of accepting a swap of chains C and C' between layers with error rate \bar{p} and \bar{p}' is given by

$$\min\left(1, \frac{\pi_C(\bar{p}')\pi_{C'}(\bar{p})}{\pi_C(\bar{p})\pi_{C'}(\bar{p}')}\right). \quad (16)$$

The error rates progress from the physical error rate to the top layer where all error chains are assumed equally likely such that proposed logical operators are always accepted. For depolarizing noise $p_x = p_y = p_z = p/3$, the interpolation is straightforward, changing p from the physical value to $p = 0.75$ in the top layer [63]. For biased noise, we have implemented an interpolation at constant α , with \bar{p}_z ranging from the physical value to $\bar{p}_z = 1$ in the top layer, where the effective weight is irrelevant [83]. An alternative approach to account for the heavy logical operators in the Monte Carlo sampling is to use Bennett's acceptance-ratio method to progress from high to low error rates, as discussed in Ref. [84].

The accuracy of the MCMC decoder depends crucially on appropriately exploring all four equivalence classes through parallel tempering. It is thus quite sensitive to the convergence criteria. In particular, we have found that for low error rates the decoder has quite poor accuracy. We interpret this as being due to an inefficiency in finding the short-effective-length error chains that will dominate the class probabilities. At low error rates, i.e., low temperatures, the exploration effectively becomes frozen into a small nonrepresentative part of the configuration space.

In contrast, the EWD decoder is based on explicitly evaluating the contribution of the most likely chains. In order to find the chains, we use Metropolis sampling, but without parallel tempering. Since the algorithm is based on an explicit identification and count of unique chains, there is no need to find a properly class-balanced set of chains. In addition, even within a class, there is no need that the Markov chain actually follows the distribution corresponding to the physical error rate. As long as the effective weight of the chains is conserved, i.e., fixed α , one may use a different error rate to explore and identify the chains, allowing for more efficient exploration. From a computational perspective, the EWD algorithm can also be run efficiently within memory constraints by using a hash-table data structure for identifying chains as unique, as there is no need to store the actual chains, but only the total number of unique chains according to weight.

B. Comparison with minimum-weight-matching algorithms

The information used in the EWD decoder is the weight and number of the most likely error chains in each equivalence class. Standard minimum-weight-matching decoding has two severe shortcomings in this respect. First, bit- and phase-flip errors are treated as separate graph problems. This amounts

to measuring the weight using a model of uncorrelated errors ($\alpha_z = \alpha_x = 1$ and $\alpha_y = 2$) such that the suggested lowest-weight chain(s) may not be the actual lowest-weight chain(s). Second, degeneracy of chains with the same weight is not considered. The latter is crucial for the high performance of the EWD decoder, as without this information, syndromes are always decoded according to the $p \rightarrow 0$ class prediction.

To quantify the effect of the first issue (incorrect weight of chains containing Y), we have studied the failure rate for chains of weight $(d+1)/2$ for depolarizing noise for the XZZX code (with equivalent results for the rotated surface code). These are the lightest error chains that will give failure, since for a small fraction of such chains there is a $(d-1)/2$ -weight chain which is in another equivalence class. These chains will subsequently give the asymptotic failure rate at low p , $P_f \sim p^{(d+1)/2}$ with a prefactor that is given by the number of failing chains. Table I shows the results of this study. The MWPM algorithm fails on a significantly larger fraction of these chains, due to overestimating the weight of chains containing Y errors.

Using Metropolis sampling as in this work eliminates both the incorrect weight of chains and the failure to take into account degeneracy, but is computationally expensive. There are also a number of works which provide improvements to minimum-weight-matching schemes by reweighting graph edges to better reflect the actual error rates and account for degeneracy of chains. We note in particular the method to enhance MWPM presented in Ref. [36], which addresses both of these issues, using multiplicity counts and belief propagation on error rates, resulting in high thresholds for both uncorrelated and depolarizing noise on the surface code.

VI. CONCLUSION

We have implemented a decoding algorithm for topological stabilizer codes, which is based on evaluating the most likely equivalence class of error chains for a given syndrome, based on the effective weight and number of most likely chains in each class. This effective weight and degeneracy decoder is error-rate agnostic in the sense of being invariant under changes that preserve the weight of error chains, corresponding to fixed relative bias of phase- to bit-flip errors. It thus represents a type of decoder which is intermediate between maximum-likelihood decoders, which work at a fixed error rate, and heuristic decoders, e.g., minimum-weight matching, which are insensitive to the error rate.

While related to the Markov-chain Monte Carlo-based decoder of Refs. [63,64], using the Metropolis algorithm to sample error chains, we find that the EWD decoder is significantly more efficient, flexible, and has better convergence properties. The reason for these advantages is that the error-chain sampling can be done independently in the equivalence classes and at a higher error rate than the physical, which facilitates efficient exploration. A drawback of the same freedom is that the results may be sensitive to the sampling error rate. We have found that it is useful to think of the sampling rate as a hyperparameter and tune it for best results. For optimal performance the decoder is limited to moderate code sizes or low error rates, for which the distribution of error chains is dominated by the most likely chains. However, for large

code sizes the Metropolis sampling becomes prohibitively expensive such that the limitations may not be very important for practical purposes.

We have found that the EWD decoder can be used as a relatively fast and flexible maximum-likelihood decoder for exploring novel stabilizer codes [85]. In addition, we argue that the compact data representation, where, for a given error bias, each syndrome is described by a small set of numbers (weight and degeneracy of most likely chains in each class), can be suitable for supervised deep-learning approaches that would allow for fast decoding while having flexibility with respect to the overall error rate.

Future work should also further explore the inclusion of circuit level noise leading to imperfect stabilizer measurements [30,64], which can be incorporated into the algorithm by appropriately including the weight of faulty measurement events on a three-dimensional lattice. Even though the EWD

decoder is slow compared to matching or union-find decoders and not practical for large codes, it may be useful for establishing approximate maximum-likelihood thresholds to provide upper bounds on code performance.

The software for running the EWD and MCMC decoders is available on the repository [83], while the implementation of the MPS decoder used in this work is available from Ref. [77].

ACKNOWLEDGMENTS

Computations were enabled by resources provided by the Swedish National Infrastructure for Computing and Chalmers Centre for Computational Science and Engineering. We acknowledge financial support from the Knut and Alice Wallenberg Foundation through the Wallenberg Centre for Quantum Technology.

-
- [1] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493 (1995).
 - [2] A. M. Steane, Error Correcting Codes in Quantum Theory, *Phys. Rev. Lett.* **77**, 793 (1996).
 - [3] D. Gottesman, *Stabilizer Codes and Quantum Error Correction* (California Institute of Technology, Pasadena, 1997).
 - [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
 - [5] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [6] S. M. Girvin, in Quantum Information Machines, edited by M. Devoret, B. Huard, and I. Pop, *Lecture Notes of the Les Houches Summer School 2019* (Oxford University Press, Oxford, in press).
 - [7] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys. (NY)* **303**, 2 (2003).
 - [8] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, [arXiv:quant-ph/9811052](https://arxiv.org/abs/quant-ph/9811052).
 - [9] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
 - [10] R. Raussendorf and J. Harrington, Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions, *Phys. Rev. Lett.* **98**, 190504 (2007).
 - [11] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [12] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner *et al.*, State preservation by repetitive error detection in a superconducting quantum circuit, *Nature (London)* **519**, 66 (2015).
 - [13] M. Takita, A. W. Cross, A. D. Córcoles, J. M. Chow, and J. M. Gambetta, Experimental Demonstration of Fault-Tolerant State Preparation with Superconducting Qubits, *Phys. Rev. Lett.* **119**, 180501 (2017).
 - [14] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff, Repeated quantum error detection in a surface code, *Nat. Phys.* **16**, 875 (2020).
 - [15] J. F. Marques *et al.*, Logical-qubit operations in an error-detecting surface code, *Nat. Phys.* **18**, 80 (2022).
 - [16] Z. Chen *et al.*, Exponential suppression of bit or phase errors with cyclic error correction, *Nature (London)* **595**, 383 (2021).
 - [17] A. Erhard, H. Poulsen Nautrup, M. Meth, L. Postler, R. Stricker, M. Stadler, V. Negnevitsky, M. Ringbauer, P. Schindler, H. J. Briegel, R. Blatt, N. Friis, and T. Monz, Entangling logical qubits with lattice surgery, *Nature (London)* **589**, 220 (2021).
 - [18] K. J. Satzinger *et al.*, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
 - [19] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, and C. Monroe, Fault-tolerant control of an error-corrected qubit, *Nature (London)* **598**, 281 (2021).
 - [20] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, Realization of real-time fault-tolerant quantum error correction, [arXiv:2107.07505](https://arxiv.org/abs/2107.07505).
 - [21] W. P. Livingston, M. S. Blok, E. Flurin, J. Dressel, A. N. Jordan, and I. Siddiqi, Experimental demonstration of continuous quantum error correction, [arXiv:2107.11398](https://arxiv.org/abs/2107.11398).
 - [22] L. Postler, S. Heußen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, P. Schindler, M. Müller, and T. Monz, Demonstration of fault-tolerant universal quantum gate operations, [arXiv:2111.12654](https://arxiv.org/abs/2111.12654).
 - [23] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, [arXiv:2112.03708](https://arxiv.org/abs/2112.03708).
 - [24] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletic, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, [arXiv:2112.03923](https://arxiv.org/abs/2112.03923).

- [25] C. Wang, J. Harrington, and J. Preskill, Confinement-Higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory, *Ann. Phys. (NY)* **303**, 31 (2003).
- [26] H. Bombin, R. S. Andrist, M. Ohzeki, H. G. Katzgraber, and M. A. Martin-Delgado, Strong Resilience of Topological Codes to Depolarization, *Phys. Rev. X* **2**, 021004 (2012).
- [27] H. G. Katzgraber and R. S. Andrist, Stability of topologically-protected quantum computing proposals as seen through spin glasses, *J. Phys.: Conf. Ser.* **473**, 012019 (2013).
- [28] J. Edmonds, Paths, trees, and flowers, *Can. J. Math.* **17**, 449 (1965).
- [29] D. S. Wang, A. G. Fowler, A. M. Stephens, and L. C. L. Hollenberg, Threshold error rates for the toric and planar codes, *Quantum Inf. Comput.* **10**, 456 (2010).
- [30] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, Surface code quantum computing with error rates over 1%, *Phys. Rev. A* **83**, 020302(R) (2011).
- [31] T. M. Stace, S. D. Barrett, and A. C. Doherty, Thresholds for Topological Codes in the Presence of Loss, *Phys. Rev. Lett.* **102**, 200501 (2009).
- [32] T. M. Stace and S. D. Barrett, Error correction and degeneracy in surface codes suffering loss, *Phys. Rev. A* **81**, 022317 (2010).
- [33] A. G. Fowler, Optimal complexity correction of correlated errors in the surface code, [arXiv:1310.0863](https://arxiv.org/abs/1310.0863).
- [34] N. Delfosse and J.-P. Tillich, *Proceedings of the 2014 IEEE International Symposium on Information Theory, Honolulu* (IEEE, Piscataway, 2014), pp. 1071–1075.
- [35] 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).
- [36] B. Criger and I. Ashraf, Multi-path summation for decoding 2D topological codes, *Quantum* **2**, 102 (2018).
- [37] G. Duclos-Cianci and D. Poulin, Fast Decoders for Topological Quantum Codes, *Phys. Rev. Lett.* **104**, 050504 (2010).
- [38] M. Herold, E. T. Campbell, J. Eisert, and M. J. Kastoryano, Cellular-automaton decoders for topological quantum memories, *npj Quantum Inf.* **1**, 15010 (2015).
- [39] A. Kubica and J. Preskill, Cellular-Automaton Decoders with Provable Thresholds for Topological Codes, *Phys. Rev. Lett.* **123**, 020501 (2019).
- [40] N. Delfosse and N. H. Nickerson, Almost-linear time decoding algorithm for topological codes, *Quantum* **5**, 595 (2021).
- [41] S. Huang, M. Newman, and K. R. Brown, Fault-tolerant weighted union-find decoding on the toric code, *Phys. Rev. A* **102**, 012419 (2020).
- [42] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [43] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning* (MIT Press, Cambridge, 2016).
- [44] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [45] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
- [46] E. P. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nat. Phys.* **13**, 435 (2017).
- [47] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys.: X* **5**, 1797528 (2020).
- [48] S. Ahmed, C. Sánchez Muñoz, F. Nori, and A. F. Kockum, Quantum State Tomography with Conditional Generative Adversarial Networks, *Phys. Rev. Lett.* **127**, 140502 (2021).
- [49] R. Sweke, M. S. Kesselring, E. P. van Nieuwenburg, and J. Eisert, Reinforcement learning decoders for fault-tolerant quantum computation, *Mach. Learn.: Sci. Technol.* **2**, 025005 (2021).
- [50] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, Quantum error correction for the toric code using deep reinforcement learning, *Quantum* **3**, 183 (2019).
- [51] H. P. Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis, Optimizing quantum error correction codes with reinforcement learning, *Quantum* **3**, 215 (2019).
- [52] L. D. Colomer, M. Skotiniotis, and R. Muñoz-Tapia, Reinforcement learning for optimal error correction of toric codes, *Phys. Lett. A* **384**, 126353 (2020).
- [53] D. Fitzek, M. Eliasson, A. F. Kockum, and M. Granath, Deep Q-learning decoder for depolarizing noise on the toric code, *Phys. Rev. Research* **2**, 023230 (2020).
- [54] H. Théveniaut and E. van Nieuwenburg, A NEAT quantum error decoder, *SciPost Phys.* **11**, 005 (2021).
- [55] G. Torlai and R. G. Melko, Neural Decoder for Topological Codes, *Phys. Rev. Lett.* **119**, 030501 (2017).
- [56] S. Krastanov and L. Jiang, Deep neural network probabilistic decoder for stabilizer codes, *Sci. Rep.* **7**, 11003 (2017).
- [57] S. Varsamopoulos, B. Criger, and K. Bertels, Decoding small surface codes with feedforward neural networks, *Quantum Sci. Technol.* **3**, 015004 (2018).
- [58] P. Baireuther, T. E. O'Brien, B. Tarasinski, and C. W. Beenakker, Machine-learning-assisted correction of correlated qubit errors in a topological code, *Quantum* **2**, 48 (2018).
- [59] N. P. Breuckmann and X. Ni, Scalable neural network decoders for higher dimensional quantum codes, *Quantum* **2**, 68 (2018).
- [60] C. Chamberland and P. Ronagh, Deep neural decoders for near term fault-tolerant experiments, *Quantum Sci. Technol.* **3**, 044002 (2018).
- [61] X. Ni, Neural network decoders for large-distance 2D toric codes, *Quantum* **4**, 310 (2020).
- [62] S. Gicev, L. C. Hollenberg, and M. Usman, A scalable and fast artificial neural network syndrome decoder for surface codes, [arXiv:2110.05854](https://arxiv.org/abs/2110.05854).
- [63] J. R. Wootton and D. Loss, High Threshold Error Correction for the Surface Code, *Phys. Rev. Lett.* **109**, 160503 (2012).
- [64] A. Hutter, J. R. Wootton, and D. Loss, Efficient Markov chain Monte Carlo algorithm for the surface code, *Phys. Rev. A* **89**, 022326 (2014).
- [65] S. Bravyi, M. Suchara, and A. Vargo, Efficient algorithms for maximum likelihood decoding in the surface code, *Phys. Rev. A* **90**, 032326 (2014).
- [66] H. Bombin and M. A. Martin-Delgado, Optimal resources for topological two-dimensional stabilizer codes: Comparative study, *Phys. Rev. A* **76**, 012305 (2007).
- [67] D. K. Tuckett, A. S. Darmawan, C. T. Chubb, S. Bravyi, S. D. Bartlett, and S. T. Flammia, Tailoring Surface Codes for Highly Biased Noise, *Phys. Rev. X* **9**, 041031 (2019).
- [68] J. P. B. Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, The XZZX surface code, *Nat. Commun.* **12**, 2172 (2021).
- [69] X.-G. Wen, Quantum Orders in an Exact Soluble Model, *Phys. Rev. Lett.* **90**, 016803 (2003).

- [70] A. Kay, Capabilities of a Perturbed Toric Code as a Quantum Memory, *Phys. Rev. Lett.* **107**, 270502 (2011).
- [71] M. R. Geller and Z. Zhou, Efficient error models for fault-tolerant architectures and the Pauli twirling approximation, *Phys. Rev. A* **88**, 012314 (2013).
- [72] Y. Tomita and K. M. Svore, Low-distance surface codes under realistic quantum noise, *Phys. Rev. A* **90**, 062320 (2014).
- [73] A. M. Stephens, Fault-tolerant thresholds for quantum error correction with the surface code, *Phys. Rev. A* **89**, 022321 (2014).
- [74] A. G. Fowler, Coping with qubit leakage in topological codes, *Phys. Rev. A* **88**, 042308 (2013).
- [75] T. E. O'Brien, B. Tarasinski, and L. DiCarlo, Density-matrix simulation of small surface codes under current and projected experimental noise, *npj Quantum Inf.* **3**, 39 (2017).
- [76] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultrahigh Error Threshold for Surface Codes with Biased Noise, *Phys. Rev. Lett.* **120**, 050505 (2018).
- [77] D. K. Tuckett, Tailoring surface codes: Improvements in quantum error correction with biased noise, Ph.D. thesis, University of Sydney, 2020 (qecsim: <https://github.com/qecsim/qecsim>).
- [78] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, *Phys. Rev. A* **54**, 3824 (1996).
- [79] V. Kolmogorov, Blossom V: A new implementation of a minimum cost perfect matching algorithm, *Math. Program. Comput.* **1**, 43 (2009).
- [80] O. Higgott, PyMatching: A Python package for decoding quantum codes with minimum-weight perfect matching, *arXiv:2105.13082*.
- [81] A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg, Towards Practical Classical Processing for the Surface Code, *Phys. Rev. Lett.* **108**, 180501 (2012).
- [82] D. J. Earl and M. W. Deem, Parallel tempering: Theory, applications, and new perspectives, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- [83] GitHub repository, <https://github.com/QEC-project-2020/EWD-QEC>.
- [84] S. Bravyi and A. Vargo, Simulation of rare events in quantum error correction, *Phys. Rev. A* **88**, 062308 (2013).
- [85] B. Srivastava, A. F. Kockum, and M. Granath, The XYZ² hexagonal stabilizer code, *arXiv:2112.06036*.