



Universiteit
Leiden
The Netherlands

Applications of topology to Weyl semimetals and quantum computing

O'Brien, T.E.

Citation

O'Brien, T. E. (2019, June 20). *Applications of topology to Weyl semimetals and quantum computing*. Retrieved from <https://hdl.handle.net/1887/74471>

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/74471>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The following handle holds various files of this Leiden University dissertation:

<http://hdl.handle.net/1887/74471>

Author: O'Brien, T.E.

Title: Applications of topology to Weyl semimetals and quantum computing

Issue Date: 2019-06-20

4. Density-matrix simulation of small surface codes under current and projected experimental noise

4.1. Introduction

Recent experimental demonstrations of small quantum simulations [164–166] and quantum error correction (QEC) [167–170] position superconducting circuits for targeting quantum supremacy [171] and quantum fault tolerance [172], two outstanding challenges for all quantum information processing platforms. On the theoretical side, much modeling of QEC codes has been made to determine fault-tolerance threshold rates in various models [173–175] with different error decoders [176–178]. However, the need for computational efficiency has constrained many previous studies to oversimplified noise models, such as depolarizing and bit-flip noise channels. This discrepancy between theoretical descriptions and experimental reality compromises the ability to predict the performance of near-term QEC implementations, and offers limited guidance to the experimentalist through the maze of parameter choices and trade-offs. In the planar circuit quantum electrodynamics (cQED) [179] architecture, the major contributions to error are transmon qubit relaxation, dephasing from flux noise and resonator photons leftover from measurement, and leakage from the computational space, none of which are well-approximated by depolarizing or bit-flip channels. Simulations with more complex error models are now essential to accurately pinpoint the leading contributions to the logical error rate in the small-distance surface codes [173, 176, 180] currently pursued by several groups worldwide.

The contents of this chapter have been published in T. E. O’Brien, B. Tarasinski and L. DiCarlo, npj Quant. Inf. **3**, 27 (2017)

4. Density-matrix simulation of small surface codes

In this chapter, we perform a density-matrix simulation of the distance-3 surface code named Surface-17, using the concrete quantum circuit recently proposed in [181] and the measured performance of current experimental multi-transmon cQED platforms [182–185]. For this purpose, we have developed an open-source density-matrix simulation package named `quantumsim` *. We use `quantumsim` to extract the logical error rate per QEC cycle, ϵ_L . This metric allows us to optimize and trade off between QEC cycle parameters, assess the merits of feedback control, predict gains from future improvements in physical qubit performance, and quantify decoder performance. We compare an algorithmic decoder using minimum-weight perfect matching (MWPM) with homemade weight calculation to a simple look-up table (LT) decoder, and weigh both against an upper bound (UB) for decoder performance obtainable from the density-matrix simulation. Finally, we make a low-order approximation to extend our predictions to the distance-5 Surface-49. The combination of results for Surface-17 and -49 allows us to make statements about code scaling and to predict the code size and physical qubit performance required to achieve break-even points for memory and computational performance.

4.2. Results

4.2.1. Error rates for Surface-17 under current experimental conditions

To quantify the performance of the logical qubit, we first define a test experiment to simulate. Inspired by the recent experimental demonstration of distance-3 and -5 repetition codes [167], we first focus on the performance of the logical qubit as a quantum memory. Specifically, we quantify the ability to hold a logical $|0\rangle$ state, by initializing this state, holding it for $k \in \{1, \dots, 20\}$ cycles, performing error correction, and determining a final logical state (see Fig. 4.6 for details). The logical fidelity $\mathcal{F}_L[k]$ is then given by the probability to match the initial state. We observe identical results when using $|1\rangle$ or $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ in place of $|0\rangle$.

We base our error model for the physical qubits on current typical experimental performance for transmons in planar cQED, using parameters from the literature and in-house results (e.g., gate-set tomography measurements). These are summarized in Table 4.1, and further detailed in Table 4.2. We focus on the QEC cycle proposed in [181], which pipelines the execution of X - and Z -type stabilizer measurements. Each

*Please visit <https://github.com/brianzi/quantumsim>




Figure 4.1.: Logical fidelity $\mathcal{F}_L[k]$ of Surface-17 with current experimental parameters (Tables 4.1 and 4.2), simulated with quantumsim as described in Fig. 4.6. The results from a MWPM decoder (green) and an implementation of the LT decoder of [176] (blue) are compared to the decoder upper bound (red). The labeled error rate is obtained from the best fit to Eq. (4.2) (also plotted). A further comparison is given to majority voting (purple, dashed), which ignores the outcome of individual stabilizer measurements, and to the fidelity $\mathcal{F}_{\text{phys}}$ of a single transmon (black) [Eq. (4.1)]. Error bars (2 s.d.) are obtained by bootstrapping.

stabilizer measurement consists of three parts: a coherent step (duration $\tau_c = 2\tau_{g,1Q} + 4\tau_{g,2Q}$), measurement (τ_m), and photon depletion from readout resonators (τ_d), making the QEC cycle time $\tau_{\text{cycle}} = \tau_c + \tau_m + \tau_d$.

Simulating this concrete quantum circuit with the listed parameters using quantumsim, we predict $\mathcal{F}_L[k]$ of Surface-17 (Fig. 4.1). We show $\mathcal{F}_L[k]$ for both a homemade MWPM decoder (green, described in App. 4.F), and an implementation of the LT decoder of [176] (blue, described in App. 4.G). To isolate decoder performance, we can compare the achieved fidelity to an upper bound extractable from the density-matrix simulation (red, described in Sec. 4.4.1). To assess the benefit of QEC, we also compare to a single decohering transmon, whose fidelity is calculated by averaging over the six cardinal points of the Bloch sphere:

$$\mathcal{F}_{\text{phys}}(t) = \frac{1}{6} \left(1 + e^{-t/T_1} \right) + \frac{1}{3} \left(1 + e^{-t(1/2T_1 + 1/T_\phi)} \right). \quad (4.1)$$

4. Density-matrix simulation of small surface codes

The observation of $\mathcal{F}_L[k] > \mathcal{F}_{\text{phys}}(k\tau_{\text{cycle}})$ for large k would constitute a demonstration of QEC beyond the quantum memory break-even point [170]. Equivalently, one can extract a logical error rate ϵ_L from a best fit to $\mathcal{F}_L[k]$ (as derived in Sec. 4.4.1 as the probability of an odd number of errors occurring),

$$\mathcal{F}_L[k] = \frac{1}{2}[1 + (1 - 2\epsilon_L)^{k-k_0}]. \quad (4.2)$$

Here, k_0 and ϵ_L are the parameters to be fit. We compare ϵ_L to the physical error rate

$$\epsilon_{\text{phys}} = -\tau_{\text{cycle}} \left. \frac{d\mathcal{F}_{\text{phys}}(t)}{dt} \right|_{t=0} = \frac{\tau_{\text{cycle}}}{3T_1} + \frac{\tau_{\text{cycle}}}{3T_\phi}. \quad (4.3)$$

We observe $\epsilon_L = 1.44\%_c$ for the LT decoder, $\epsilon_L = 1.07\%_c$ for the MWPM decoder, and $\epsilon_L = 0.68\%_c$ at the decoder upper bound ($\%_c = \%$ per cycle). The latter two fall below $\epsilon_{\text{phys}} = 1.33\%_c$. Defining the decoder efficiency $\eta_d = \epsilon_L^{(\text{UB})}/\epsilon_L$, we find $\eta_d^{(\text{LT})} = 0.47$ and $\eta_d^{(\text{MWPM})} = 0.64$.

We can also compare the multi-cycle error correction to majority voting, in which the state declaration is based solely on the output of the final data qubit measurements (ancilla measurements are ignored). Majority voting corrects any single data qubit error (over the entire experiment), and thus exhibits a quadratic decay for small k *. A decoder should also be able to correct (at least) a single error, and thus should produce the same behavior at low k , delaying the onset of exponential decay in $\mathcal{F}_L[k]$. In fact, a good test for the performance of a MWPM decoder is to ensure it can outperform the majority vote at short timescales, as suboptimal configuration will prevent this (as seen for the look-up table decoder).

With the baseline for current performance established, we next investigate ϵ_L improvements that may be achieved by two means. First, we consider modifications to the QEC cycle at fixed physical performance. Afterwards, we consider the effect of improving physical qubit T_1 and T_ϕ .

4.2.2. Optimization of logical error rates with current experimental conditions

Error sources in current cQED setups derive primarily from transmon decoherence, as opposed to gate and measurement errors produced by control electronics. Thus, a path to reducing ϵ_L may be to decrease

*A distance- d code with majority voting alone should exhibit a $(d+1)/2$ -order decay

τ_{cycle} . Currently, the cycle is dominated by $\tau_{\text{m}} + \tau_{\text{d}}$. At fixed readout power, reducing τ_{m} and τ_{d} will reduce τ_{cycle} at the cost of increased readout infidelity ϵ_{RO} (described in Sec. 4.4.2). We explore this trade-off in Fig. 4.2, using a linear-dispersive readout model [186], keeping $\tau_{\text{m}} = \tau_{\text{d}}$ and assuming no leftover photons. Because of the latter, $\epsilon_{\text{L}}^{(\text{MWPM})}$ reduces from $1.07\%_{\text{c}}$ (Fig. 4.1) to $0.62\%_{\text{c}}$ at $\tau_{\text{m}} = 300$ ns. The minimum $\epsilon_{\text{L}}^{(\text{MWPM})} = 0.55\%_{\text{c}}$ is achieved at around $\tau_{\text{m}} = 260$ ns. This is perhaps counterintuitive, as ϵ_{phys} reduces only $0.13\%_{\text{c}}$ while ϵ_{RO} increases 0.5% . However, it reflects the different sensitivity of the code to different types of errors. Indeed, $\epsilon_{\text{L}}^{(\text{MWPM})}$ is smaller for $\tau_{\text{m}} = 200$ ns than for $\tau_{\text{m}} = 300$ ns, even though ϵ_{RO} increases to 5% . It is interesting to note that the optimal τ_{m} for quantum memory, which minimizes logical error per unit time, rather than per cycle, is $\tau_{\text{m}} = 280$ ns (Fig. 4.2 inset). This shows that different cycle parameters might be optimal for computation and memory applications.




Figure 4.2.: Optimization of the logical error rate (per cycle) of Surface-17 as a function of measurement-and-depletion time [182]. Changes in the underlying physical error rates are shown as well. Decreasing the measurement time causes an increase in the readout infidelity (solid black curve with dots), whilst decreasing the single qubit decay from T_1 and T_2 (black dashed curve) for all qubits. The logical rate with an MWPM decoder (green curve) is minimized when these error rates are appropriately balanced. The logical error rate is calculated from the best fit of Eq. (4.2). Error bars (2 s.d.) are obtained by bootstrapping ($N = 10,000$ runs). Inset: Logical error rate per unit time, instead of per cycle.

4. Density-matrix simulation of small surface codes

Next, we consider the possibility to reduce ϵ_L using feedback control. Since T_1 only affects qubits in the excited state, the error rate of ancillas in Surface-17 is roughly two times higher when in the excited state. The unmodified syndrome extraction circuit flips the ancilla if the corresponding stabilizer value is -1, and since ancillas are not reset between cycles, they will spend significant amounts of time in the excited state. Thus, we consider using feedback to hold each ancilla in the ground state as much as possible. We do not consider feedback on data qubits, as the highly entangled logical states are equally susceptible to T_1 .

The feedback scheme (Inset of Fig. 3) consists of replacing the $R_y(\pi/2)$ gate at the end of the coherent step with a $R_y(-\pi/2)$ gate for some of the ancillas, depending on a classical control bit p for each ancilla. This bit p represents an estimate of the stabilizer value, and the ancilla is held in the ground state whenever this estimate is correct (i.e. in the absence of errors). Figure 4.3 shows the effect of this feedback on the logical fidelity, both for the MWPM decoder and the decoder upper bound. We observe ϵ_L improve only 0.05% in both cases. Future experiments might opt not to pursue these small gains in view of the technical challenges added by feedback control.

4.2.3. Projected improvement with advances in quantum hardware

We now estimate the performance increase that may result from improving the transmon relaxation and dephasing times via materials and filtering improvements. To model this, we return to $\tau_{\text{cycle}} = 800$ ns, and adjust T_1 values with both $T_\phi = 2T_1$ (common in experiment) and $T_\phi = \infty$ (all white-noise dephasing eliminated). We retain the same rates for coherent errors, readout infidelity, and photon-induced dephasing as in Fig. 4.1. Figure 4.4 shows the extracted ϵ_L and ϵ_{phys} over the T_1 range covered. For the MWPM decoder (upper bound) and $T_\phi = 2T_1$, the memory figure of merit $\gamma_m = \epsilon_{\text{phys}}/\epsilon_L$ increases from 1.3 (2) at $T_1 = 30$ μs to 2 (5) at 100 μs . Completely eliminating white-noise dephasing will increase γ_m by 10% with MWPM and 30% at the upper bound.

A key question for any QEC code is how ϵ_L scales with code distance d . Computing power limitations preclude similar density-matrix simulations of the $d = 5$ surface code Surface-49. However, we can approximate the error rate by summing up all lowest-order error chains (as calculated for the MWPM decoder), and deciding individually whether or not these would be corrected by a MWPM decoder (see App. 4.H for details). Figure 4.5 shows the lowest-order approximation to the logical error rates of Surface-17




Figure 4.3.: Logical fidelity of Surface-17 with (solid) and without (dashed) an additional feedback scheme. The performance of a MWPM decoder (green) is compared to the decoder upper bound (red). Curves are fits of Eq. (4.2) to the data, and error bars (2 s.d.) are given by bootstrapping, with each point averaged over 10,000 runs. Inset: Method for implementing the feedback scheme. For each ancilla qubit A_j , we store a parity bit p_j , which decides the sign of the $R_y(\pi/2)$ rotation at the end of each coherent step. The time A_j spends in the ground state is maximized when p_j is updated each cycle t by XORing with the measurement result from cycle $t - 1$, after the rotation of cycle t has been performed.

and -49 over a range of $T_1 = T_\phi/2$. Comparing the Surface-17 lowest-order approximation to the quantumsim result shows good agreement and validates the approximation. We observe a lower ϵ_L for Surface-49 than for -17, indicating quantum fault tolerance over the T_1 range covered. The fault-tolerance figure of merit defined in [172], $\Lambda_t = \epsilon_L^{(17)}/\epsilon_L^{(49)}$, increases from 2 to 4 as T_1 grows from 30 to 100 μs .

As a rough metric of computational performance, we offer to compare ϵ_L (per cycle) to the error accrued by a physical qubit idling over $\tau_{g,1Q}$. We define a metric for computation performance, $\gamma_c = (\epsilon_{\text{phys}}\tau_{g,1Q})/(\epsilon_L\tau_{\text{cycle}})$ and $\gamma_c = 1$ as a computational break-even point. Clearly, using the QEC cycle parameters of Table 4.1 and even with T_1 improvements, neither Surface-17 nor -49 can break-even computationally. However, including the readout acceleration recently demonstrated in [185], which allows $\tau_m = \tau_d = 100$ ns and $\tau_{\text{cycle}} = 400$ ns, Surface-49 can cross $\gamma_c = 1$ by

4. Density-matrix simulation of small surface codes




Figure 4.4.: T_1 dependence of the Surface-17 logical error rate (MWPM and UB) and the physical error rate. We either fix $T_\phi = 2T_1$ (solid) or $T_\phi = \infty$ (dashed). Logical error rates are extracted from a best fit of Eq. (4.2) to $\mathcal{F}_L[k]$ over $k = 1, \dots, 20$ QEC cycles, averaged over $N = 50,000$ runs. Error bars (2 s.d.) are calculated by bootstrapping.

$T_1 = 40 \mu\text{s}$. In view of first reports of T_1 up to $80 \mu\text{s}$ emerging for planar transmons [187, 188], this important milestone may be within grasp.

4.3. Discussion

4.3.1. Computational figure of merit

We note that our metric of computational power is not rigorous, due to the different gate sets available to physical and logical qubits. Logical qubits can execute multiple logical X and Z gates within one QEC cycle, but require a few cycles for two-qubit and Hadamard gates (using the proposals of [175, 180]), and state distillation over many cycles to perform non-Clifford gates. As such, this metric is merely a rough benchmark for computational competitiveness of the QEC code. However, given the amount by which all distance-3 logical fidelities fall above this metric, we find it unlikely that these codes will outperform a physical qubit by any fair comparison in the near future.




Figure 4.5.: Analytic approximation of ϵ_L for Surface-17 (green) and Surface-49 (orange) using a MWPM decoder. Details of the calculation of points and error bars are given in App. 4.H. All plots assume $T_\phi = 2T_1$, and $\tau_{\text{cycle}} = 800$ ns (crosses) or 400 ns (dots). Numerical results for Surface-17 with $\tau_{\text{cycle}} = 800$ ns are also plotted for comparison (green, dashed). The physical-qubit computation metric is given as the error incurred by a single qubit over the resting time of a single-qubit gate (black, dashed).

4.3.2. Decoder performance

A practical question facing quantum error correction is how best to balance the trade-off between decoder complexity and performance. Past proposals for surface-code computation via lattice surgery [180] require the decoder to provide an up-to-date estimate of the Pauli error on physical qubits during each logical T gate. Because tracking Pauli errors through a non-Clifford gate is inefficient, however implemented, equivalent requirements will hold for any QEC code [66]. A decoder is thus required to process ancilla measurements from one cycle within the next (on average). This presents a considerable challenge for transmon-cQED implementations, as $\tau_{\text{cycle}} < 1 \mu\text{s}$. This short time makes the use of computationally intensive decoding schemes difficult, even if they provide lower ϵ_L .

The leading strategy for decoding the surface code is MWPM using the blossom algorithm of Edmonds [173, 177, 189]. Although this algorithm is challenging to implement, it scales linearly in code distance [189]. The algorithm requires a set of weights (representing the probability that two

4. Density-matrix simulation of small surface codes

given error signals are connected by a chain of errors) as input. An important practical question (see App. 4.F) is whether these weights can be calculated on the fly, or must be precalculated and stored. On-the-fly weight calculation is more flexible. For example, it can take into account the difference in error rates between an ancilla measured in the ground and in the excited state. The main weakness of MWPM is the inability to explicitly detect Y errors. In fact, in App. 4.C we see that MWPM is nearly perfect in the absence of Y errors. The decoder efficiency η_d may significantly increase by extending MWPM to account for correlations between detected X and Z errors originating from Y errors [190, 191].

If computational limitations preclude a MWPM decoder from keeping up with τ_{cycle} , the look-up table decoder may provide a straightforward solution for Surface-17. However, at current physical performance, the η_d reduction will make Surface-17 barely miss memory break-even (Fig. 4.1). Furthermore, memory requirements make look-up table decoding already impractical for Surface-49. Evidently, real-time algorithmic decoding by MWPM or improved variants is an important research direction already at low code distance.

4.3.3. Other observations

The simulation results allow some further observations. Although we have focused on superconducting qubits, we surmise that the following statements are fairly general.

We observe that small quasi-static qubit errors are suppressed by the repeated measurement. In our simulations, the $1/f$ flux noise producing 0.01 radians of phase error per flux pulse on a qubit has a diamond norm approximately equal to the T_1 noise, but a trace distance 100 times smaller. As the flux noise increases ϵ_L by only 0.01 %_c, it appears ϵ_L is dependent on the trace distance rather than the diamond norm of the underlying noise components. Quasi-static qubit errors can then be easily suppressed, but will also easily poison an experiment if unchecked.

We further observe that above a certain value, ancilla and measurement errors have a diminished effect on ϵ_L . In our error model, the leading sources of error for a distance d code are chains of $(d-1)/2$ data qubit errors plus either a single ancilla qubit error or readout error, which together present the same syndrome as a chain of $(d+1)/2$ data qubit errors. An optimal decoder decides which of these chains is more likely, at which point the less-likely chain will be wrongly corrected, completing a logical error. This implies that if readout infidelity (ϵ_{RO}) or the ancilla error rate (ϵ_{anc}) is below the data qubit (ϵ_{phys}) error rate, $\epsilon_L \propto (\epsilon_{\text{anc}} + \epsilon_{\text{RO}})\epsilon_{\text{phys}}^{(d-1)/2}$.

However, if $\epsilon_{\text{RO}} (\epsilon_{\text{anc}}) > \epsilon_{\text{phys}}$, ϵ_{L} becomes independent of $\epsilon_{\text{RO}} (\epsilon_{\text{anc}})$, to lowest order. This can be seen in Fig. 4.2, where the error rate is almost constant as ϵ_{RO} exponentially increases. This approximation breaks down with large enough ϵ_{anc} and ϵ_{RO} , but presents a counterintuitive point for experimental design; ϵ_{L} becomes less sensitive to measurement and ancilla errors as these error get worse.

A final, interesting point for future surface-code computation is shown in Fig. 4.2: the optimal cycle parameters for logical error rates per cycle and per unit time are not the same. This implies that logical qubits functioning as a quantum memory should be treated differently to those being used for computation. This idea can be extended further: at any point in time, a large quantum computer performing a computation will have a set S_m of memory qubits which are storing part of a large entangled state, whilst a set S_c of computation qubits containing the rest of the state undergo operations. To minimize the probability of a logical error occurring on qubits within both S_c and S_m , the cycle time of the qubits in S_c can be reduced to minimize the rest time of qubits in S_m . As a simple example, consider a single computational qubit q_c and a single memory qubit q_m sharing entanglement. Operating all qubits at $\tau_{\text{cycle}} = 720$ ns to minimize ϵ_{L} would lead to a 1.09% error rate for the two qubits combined. However, shortening the τ_{cycle} of q_c reduces the time over which q_m decays. If q_c operates at $\tau_{\text{cycle}} = 600$ ns, the average error per computational cycle drops to 1.06%, as q_m completes only 5 cycles for every 6 on q_c . Although this is only a meager improvement, one can imagine that when many more qubits are resting than performing computation, the relative gain will be quite significant.

4.3.4. Effects not taken into account

Although we have attempted to be thorough in the detailing of the circuit, we have neglected certain effects. We have used a simple model for C-Z gate errors as we lack data from experimental tomography (e.g. one obtained from two-qubit gate-set tomography [192]). Most importantly, we have neglected leakage, where a transmon is excited out of the two lowest energy states, i.e., out of the computational subspace. Previous experiments have reduced the leakage probability per C-Z gate to $\sim 0.3\%$ [193], and per single-qubit gate to $\sim 0.001\%$ [194]. Schemes have also been developed to reduce the accumulation of leakage [195]. Extending quantumsim to include and investigate leakage is a next target. However, the representation of the additional quantum state can increase the simulation effort significantly [by a factor of $(9/4)^{10} \approx 3000$]. To still achieve this goal, some further

4. Density-matrix simulation of small surface codes

approximations or modifications to the simulation will be necessary. Future simulations will also investigate the effect of spread in qubit parameters, both in space (i.e., variation of physical error rates between qubits) and time (e.g., T_1 fluctuations), and cross-talk effects such as residual couplings between nearest and next-nearest neighbor transmons, qubit cross-driving, and qubit dephasing by measurement pulses targeting other qubits.

4.4. Methods

4.4.1. Simulated experimental procedure

Surface-17 basics

A QEC code can be defined by listing the data qubits and the stabilizer measurements that are repeatedly performed upon them [68]. In this way, Surface-17 is defined by a 3×3 grid of data qubits $\{D_0, \dots, D_8\}$. In order to stabilize a single logical qubit, $9 - 1 = 8$ commuting measurements are performed. The stabilizers are the weight-two and weight-four X - and Z -type parity operators X_2X_1 , Z_3Z_0 , $X_4X_3X_1X_0$, $Z_5Z_4Z_2Z_1$, $Z_7Z_6Z_4Z_3$, $X_8X_7X_5X_4$, Z_8Z_5 , and X_7X_6 , where X_j (Z_j) denotes the X (Z) Pauli operator acting on data qubit D_j . Their measurement is realized indirectly using nearest-neighbor interactions between data and ancilla qubits arranged in a square lattices, followed by ancilla measurements [Fig. 4.6(a)]. This leads to a total of 17 physical qubits when a separate ancilla is used for each individual measurement. We follow the circuit realization of this code described in [181], for which we give a schematic description in Fig. 4.6(b) (see App. 4.A for a full circuit diagram).

In an experimental realization of this circuit, qubits will regularly accumulate errors. Multiple errors that occur within a short period of time (e.g., one cycle) form error ‘chains’ that spread across the surface. Errors on single qubits, or correlated errors within a small subregion of Surface-17, fail to commute with the stabilizer measurements, creating error signals that allow diagnosis and correction of the error via a decoder. However, errors that spread across more than half the surface in a short enough period of time are misdiagnosed, causing an error on the logical qubit when wrongly corrected [173]. The rate at which these logical errors arise is the main focus of this chapter.

Protocol for measurement of logical error rates

As the performance measure of Surface-17, we study the fidelity of the logical qubit as a quantum memory. We describe our protocol with an example ‘run’ in Fig. 4.6. We initialize all qubits in $|0\rangle$ and perform $k = 1, 2, \dots, 20$ QEC cycles [Fig. 4.6(b)]. Although this initial state is not a stabilizer eigenstate, the first QEC cycle projects the system into one of the 16 overlapping eigenstates within the $+1$ eigenspace for Z stabilizers, which form the logical $|0\rangle$ state [173]. This implies that, in the absence of errors, the first measurement of the Z stabilizers will be $+1$, whilst that of the X stabilizers will be random. In the following cycles, ancilla measurements of each run [Fig. 4.6(c)] are processed using a classical decoding algorithm. The decoder computes a Pauli update after each QEC cycle [Fig. 4.6(d)]. This is a best estimate of the Pauli operators that must be applied to the data qubits to transform the logical qubit back to the logical $|0\rangle$ state. The run ends with a final measurement of all data qubits in the computational basis. From this 9-bit outcome, a logical measurement result is declared [Fig. 4.6(e)]. First, the four Z -type parities are calculated from the 9 data-qubit measurement outcomes and presented to the decoder as a final set of parity measurements. This ensures that the final computed Pauli update will transform the measurement results into a set that measures $+1$ for all Z stabilizers. This results in one of 32 final measurements, from which the value of a logical Z operator can be calculated to give the measurement result (any choice of logical operator gives the same result). The logical fidelity $\mathcal{F}_L[k]$ after k QEC cycles is defined as the probability of this declared result matching the initial $+1$ state.

At long times and with low error rates, Surface codes have a constant logical error rate ϵ_L . The fidelity $\mathcal{F}_L[k]$ is obtained by counting the probability of an odd number of errors having occurred in total (as two σ_x errors cancel)* [183]:

$$\mathcal{F}_L[k] = 1 - \sum_{l \text{ odd}} \binom{k}{l} \epsilon_L^l (1 - \epsilon_L)^{k-l}. \quad (4.4)$$

Here, the combinatorial factor counts the number of combinations of l errors in k rounds, given an ϵ_L chance of error per round. This can be

*We thank Barbara Terhal for providing this derivation.

4. Density-matrix simulation of small surface codes

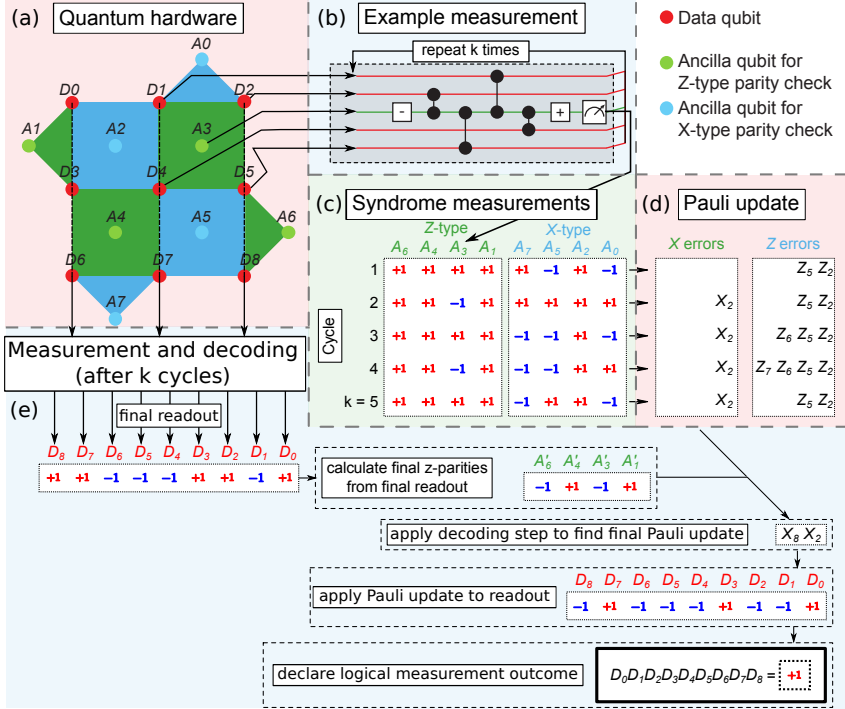


Figure 4.6.: Schematic overview of the simulated experiment. (a) 17 qubits are arranged in a surface code layout (legend top-right). The red data qubits are initialized in the ground state $|0\rangle$, and projected into an eigenstate of the measured X- (blue) and Z- (green) type stabilizer operators. (b) A section of the quantum circuit depicting the four-bit parity measurement implemented by the A_3 ancilla qubit (+/- refer to $R_y(\pm\pi/2)$ single-qubit rotations). The ancilla qubit (green line, middle) is entangled with the four data qubits (red lines) to measure $Z_1 Z_2 Z_4 Z_5$. Ancillas are not reset between cycles. Instead, the implementation relies on the quantum non-demolition nature of measurements. The stabilizer is then the product of the ancilla measurement results of successive cycles. This circuit is performed for all ancillas and repeated k times before a final measurement of all (data and ancilla) qubits. (c) All syndrome measurements of the k cycles are processed by the decoder. (d) After each cycle, the decoder updates its internal state to represent the most likely set of errors that occurred. (e) After the final measurement, the decoder uses the readout from the data qubits, along with previous syndrome measurements, to declare a final logical state. To this end, the decoder processes the Z-stabilizers obtained directly from the data qubits, finalizing its prediction of most likely errors. The logical parity is then determined as the product of all data qubit parities ($\prod_{j=0}^8 D_j$) once the declared errors are corrected. The logical fidelity \mathcal{F}_L is the probability that this declaration is the same as the initial state ($|0\rangle$).

simplified to

$$\begin{aligned}
\mathcal{F}_L[k] &= 1 - \frac{1}{2} \sum_l \binom{k}{l} \epsilon_L^l (1 - \epsilon_L)^{k-l} (1 - (-1)^l) \\
&= 1 - \frac{1}{2} [(1 - \epsilon_L + \epsilon_L)^k - (1 - \epsilon_L - \epsilon_L)^k] \\
&= \frac{1}{2} [1 + (1 - 2\epsilon_L)^k]. \tag{4.5}
\end{aligned}$$

However, at small k , the decay is dominated by the majority vote, for which $\epsilon_L \propto (k\epsilon_{\text{phys}})^{(d+1)/2}$. For example, for all the Surface-17 decay curves, we observe a quadratic error rate at small k , as opposed to the linear slope predicted by Eq. (4.5). In order to correct for this, we shift the above equation in k by a free parameter k_0 , resulting in Eq. (4.2). This function fits well to data with $k \geq 3$ in all plots, and thus allows accurate determination of ϵ_L .

The `quantumsim` simulation package

`Quantumsim` performs calculations on density matrices utilizing a graphics processing unit in a standard desktop computer. Ancillas are measured at the end of each cycle, and thus not entangled with the rest of the system. As such, it is possible to obtain the effect of the QEC cycle on the system without explicitly representing the density matrix of all 17 qubits simultaneously. The simulation is set up as follows: the density matrix of the nine data qubits is allocated in memory with all qubits initialized to $|0\rangle$. One- and two-qubit gates are applied to the density matrix as completely positive, trace preserving maps represented by Pauli transfer matrices. When a gate involving an ancilla qubit must be performed, the density matrix of the system is dynamically enlarged to include that one ancilla.

Qubit measurements are simulated as projective and following the Born rule, with projection probabilities given by the squared overlap of the input state with the measurement basis states. In order to capture empirical measurement errors, we implement a black-box measurement model (Sec. 4.4.2) by sandwiching the measurement between idling processes. The measurement projects the system to a product state of the ancilla and the projected sub-block of the density matrix. We can therefore remove the ancilla from the density matrix and only store its state right after projection, and continue the calculation with the partial density matrix of the other qubits. Making use of the specific arrangement of the

4. Density-matrix simulation of small surface codes

interactions between ancillas and data qubits in Surface-17, it is possible to apply all operations to the density matrix in such an order (shown in App. 4.A) that the total size of the density matrix never exceeds $2^{10} \times 2^{10}$ (nine data qubits plus one ancilla), which allows relatively fast simulation. We emphasize that with the choice of error model in this chapter, this approach gives the same result as a full simulation on a 17-qubit density matrix. Only the introduction of residual entangling interactions between data and ancilla qubits (which we do not consider in this chapter) would make the latter necessary. On our hardware (see App. 4.E), simulating one QEC cycle of Surface-17 with `quantumsim` takes 25 ms.

We highlight an important advantage of doing density-matrix calculations with `quantumsim`. We do not perform projective measurements of the data qubits. Instead, after each cycle, we extract the diagonal of the data-qubit density matrix, which represents the probability distribution if a final measurement were performed. We leave the density matrix undisturbed and continue simulation up to $k = 20$. This is a very useful property of the density-matrix approach, because having a probability distribution of all final readout events greatly reduces sampling noise.

Our measurement model includes a declaration error probability (see Sec. 4.4.2), where the projected state of the ancilla after measurement is not the state reported to the decoder. Before decoding, we thus apply errors to the outcomes of the ancilla projections, and smear the probability distribution of the data qubit measurement. To then determine the fidelity averaged over this probability distribution, we present all 16 possible final Z -type parities to the decoder. This results in 16 different final Pauli updates, allowing us to determine correctness of the decoder for all 512 possible measurement outcomes. These are then averaged over the simulated probability distribution. This produces good results after about $\sim 10^4$ simulated runs.

A second highlight of `quantumsim` is the possibility to quantify the sub-optimality of the decoder. The fidelity of the logical qubit obtained in these numerical simulations is a combination of the error rates of the physical qubits and the approximations made by the decoder. Full density-matrix simulations make it possible to disentangle these two contributions. Namely, the fidelity is obtained by assigning correctness to each of the 512 possible readouts according to 16 outputs of the decoder, and summing the corresponding probabilities accordingly. If the probabilities are known, it is easy to determine the 16 results that a decoder should output in order to maximize fidelity (i.e., the output of the best-possible decoder). This allows placing a decoder upper bound \mathcal{F}_L^{\max} on logical fidelity as limited by the physical qubits independent of the decoder. Conversely, it also allows

quantifying sub-optimality in the decoder used. In fact, we can make the following reverse statement: if our measurement model did not include a declaration error, then we could use the simulation to find the final density matrix of the system conditioned on a syndrome measurement. From this, the simulation could output exactly the 16 results that give \mathcal{F}_L^{\max} , so that `quantumsim` could thus be used as a maximum-likelihood decoder. In this situation, \mathcal{F}_L^{\max} would not only be an upper bound, but indeed the performance of the best-possible decoder. However, as we add the declaration errors after simulation, we can only refer to \mathcal{F}_L^{\max} as the decoder upper bound.

4.4.2. Error models

We now describe the error model used in the simulations. Our motivation for the development of this error model is to provide a limited number of free parameters to study, whilst remaining as close to known experimental data as possible. As such, we have taken well-established theoretical models as a base, and used experimental tomography to provide fixed parameters for observed noise beyond these models. The parameters of the error model are provided in App. 4.B.

Parameter	Symbol	Value	Reference
Qubit relaxation time	T_1	30 μ s	[182]
Qubit dephasing time (white noise)	T_ϕ	60 μ s	[182, 184]
Single-qubit gate time	$\tau_{g,1Q}$	20 ns	[182, 184]
Two-qubit gate time	$\tau_{g,2Q}$	40 ns	[168]
Coherent step time	τ_c	200 ns	[181]
Measurement time	τ_m	300 ns	[182]
Depletion time	τ_d	300 ns	[182]
Fast measurement time	$\tau_m^{(\text{fast})}$	100 ns	[185]
Fast depletion time	$\tau_d^{(\text{fast})}$	100 ns	[185]

Table 4.1.: Standard simulation parameters: Summary of standard times used in all density-matrix simulations, unless otherwise indicated. The two-qubit gate is a conditional phase gate (C-Z). Other error rates and parameters are given in Table 4.2.

4. Density-matrix simulation of small surface codes

Idling qubits

While idling for a time τ , a transmon in $|1\rangle$ can relax to $|0\rangle$. Furthermore, a transmon in superposition can acquire random quantum phase shifts between $|0\rangle$ and $|1\rangle$ due to $1/f$ noise sources (e.g., flux noise) and broadband ones (e.g., photon shot noise [196] and quasiparticle tunneling [197]). These combined effects can be parametrized by probabilities $p_1 = \exp(-\tau/T_1)$ for relaxation, and $p_\phi = \exp(-\tau/T_\phi)$ for pure dephasing. The combined effects of relaxation and pure dephasing lead to decay of the off-diagonal elements of the qubit density matrix. We model dephasing from broadband sources in this way, taking for T_ϕ the value extracted from the decay time T_2 of standard echo experiments:

$$\frac{1}{T_2} = \frac{1}{T_\phi} + \frac{1}{2T_1}. \quad (4.6)$$

We model $1/f$ sources differently, as discussed below.

Dephasing from photon noise

The dominant broadband dephasing source is the shot noise due to photons in the readout resonator. This dephasing is present whenever the coupled qubit is brought into superposition before the readout resonator has returned to the vacuum state following the last measurement. This leads to an additional, time-dependent pure dephasing (rates given in Table 4.2).

One-qubit Y rotations

We model y -axis rotations as instantaneous rotations sandwiched by idling periods of duration $\tau_{g,1Q}/2$. The errors in the instantaneous gates are modeled from process matrices measured by gate-set tomography [192, 198] in a recent experiment [183]. In this experiment, the GST analysis of single-qubit gates also showed that the errors can mostly be attributed to Markovian noise. For simplicity, we thus model these errors as Markovian.

Dephasing of flux-pulsed qubits

During the coherent step, transmons are repeatedly moved in frequency away from their sweetspot using flux pulses, either to implement a C-Z gate or to avoid one. Away from the sweetspot, transmons become first-order sensitive to flux noise, which causes an additional random phase shift. As this noise typically has a $1/f$ power spectrum, the largest contribution

comes from low-frequency components that are essentially static for a single run, but fluctuating between different runs. In our simulation, we approximate the effect of this noise through ensemble averaging, with quasi-static phase error added to a transmon whenever it is flux pulsed. Gaussian phase errors with the variance (calculated in App. 4.B.4) are drawn independently for each qubit and for each run.

C-Z gate error

The C-Z gate is achieved by flux pulsing a transmon into the $|11\rangle \leftrightarrow |02\rangle$ avoided crossing with another, where the 2 denotes the second-excited state of the fluxed transmon. Holding the transmons here for $\tau_{g,2Q}$ causes the probability amplitudes of $|01\rangle$ and $|11\rangle$ to acquire phases [199]. Careful tuning allows the phase ϕ_{01} acquired by $|01\rangle$ (the single-qubit phase ϕ_{1Q}) to be an even multiple of 2π , and the phase ϕ_{11} acquired by $|11\rangle$ to be π extra. This extra phase acquired by $|11\rangle$ is the two-qubit phase ϕ_{2Q} . Single- and two-qubit phases are affected by flux noise because the qubit is first-order sensitive during the gate. Previously, we discussed the single-qubit phase error. In App. 4.B.5, we calculate the corresponding two-qubit phase error $\delta\phi_{2Q}$. Our full (but simplistic) model of the C-Z gate consists of an instantaneous C-Z gate with single-qubit phase error $\delta\phi_{1Q}$ and two-qubit phase error $\delta\phi_{2Q} = \delta\phi_{1Q}/2$, sandwiched by idling intervals of duration $\tau_{g,2Q}/2$.

Measurement

We model qubit measurement with a black-box description using parameters obtained from experiment. This description consists of the eight probabilities for transitions from an input state $|i\rangle \in \{|0\rangle, |1\rangle\}$ into pairs $(m, |o\rangle)$ of measurement outcome $m \in \{+1, -1\}$ and final state $|o\rangle \in \{|0\rangle, |1\rangle\}$. By final state we mean the qubit state following the photon-depletion period. Input superposition states in the computational bases are first projected to $|0\rangle$ and $|1\rangle$ following the Born rule. The probability tree (the butterfly) is then used to obtain an output pair $(m, |o\rangle)$. These experimental parameters can be described by a six-parameter model (described in detail in App. 4.B.6), consisting of periods of enhanced noise before and after a point at which the qubit is perfectly projected, and two probabilities $\epsilon_{RO}^{(i)}$ for wrongly declaring the result of this projective measurement. In App. 4.B.6, a scheme for measuring these butterfly parameters and mapping them to the six-parameter model is described. In experiment, we find that the readout errors $\epsilon_{RO}^{(i)}$ are almost independent of the qubit state $|i\rangle$, and so we

4. Density-matrix simulation of small surface codes

describe them with a single readout error parameter ϵ_{RO} in this chapter.




Figure 4.7.: The quantum circuit for Surface-17 syndrome measurement used in all simulations. (a) Outline of the timing of the standard circuit, including the time shift between X- and Z-type stabilizer measurements described by [181]. Qubit labels correspond to the position in Fig. 6. (b) Full quantum circuit of the QEC cycle. The C-Z gates within each group are slightly offset horizontally for visibility (in reality they are performed simultaneously).

4.A. Full circuit diagram for Surface-17 implementation

The quantum circuit [181] (Fig. 4.7) consists of $R_y(\pi/2)$ (“+”) and $R_y(-\pi/2)$ (“-”) rotations, C-Z gates, and ancilla measurements. The coherent steps of the X and Z ancillas are pipelined (shifted in time with respect to each other) to prevent transmon-transmon avoided crossings. As long as $\tau_m + \tau_d \geq \tau_c$, no time is lost due to this separation.

In a simulation of the given circuit, gates on different qubits commute and may be applied to the density matrix in any order, regardless of the times at which they are performed in an experiment. As described in Sec. IV A 3, by simulating gates in a specific order (Fig. 4.8), one can ensure that only one ancilla is entangled with the data qubits at any point in the simulation. This allows a reduction in the maximum size of the density matrix from $2^{17} \times 2^{17}$ to $2^{10} \times 2^{10}$.

4.B. Parameters of error models

Parameter	Symbol	Value	Reference
In-axis rotation error	p_{axis}	10^{-4}	[183]
In-plane rotation error	p_{plane}	5×10^{-4}	[183]
$1/f$ flux noise.	A	$(1\mu\Phi_0)^2$	[200, 201]
Readout infidelity	ϵ_{RO}	5×10^{-3}	[182]
Photon relaxation time	$1/\kappa$	250 ns	[182]
Dispersive shift	χ/π	-2.6 MHz	[182]
photon # post-measurement	n_0	0.8 photons	[182]

Table 4.2.: Standard parameters of error models used in `quantumsim`, unless indicated otherwise.

This appendix provides mathematical details of the sources of error described in the main text. Standard values for the parameters used throughout the text are given in Table 4.2.

In the `quantumsim` module, all gates are applied in the Pauli transfer matrix representation [324]. These are given in the form

$$(R_\Lambda)_{ij} = \frac{1}{2} \text{Tr}(\sigma_i \Lambda \sigma_j), \quad (4.7)$$

where matrices σ_i are the Pauli operators: $\sigma_0 = I$, $\sigma_1 = X$, $\sigma_2 = Y$ and $\sigma_3 = Z$.

4.B.1. Qubit idling

Idling qubits are described by the amplitude-phase damping model [23], corresponding to the transfer matrices

$$R_{\Lambda_{T_1}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-p_1} & 0 & 0 \\ 0 & 0 & \sqrt{1-p_1} & 0 \\ p_1 & 0 & 0 & 1-p_1 \end{pmatrix} \quad (4.8)$$

$$R_{\Lambda_{T_\phi}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-p_\phi} & 0 & 0 \\ 0 & 0 & \sqrt{1-p_\phi} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (4.9)$$

Idling for a duration t is thus described by

$$R_{AP(t)} = R_{\Lambda_{T_1}} R_{\Lambda_{T_\phi}} \quad (4.10)$$

with $p_1 = 1 - e^{-t/T_1}$ and $p_\phi = 1 - e^{-t/T_\phi}$.

4.B.2. Photon decay

In the presence of photons in a readout resonator, the coupled qubit is affected according to the effective stochastic master equation [186]:

$$\frac{d\rho}{dt} = -i\frac{B}{2}[\sigma_z, \rho] + \frac{\Gamma_d}{2}\mathcal{D}[\sigma_z]\rho.$$

Here, ρ is the qubit density matrix, $\mathcal{D}[X]$ is the Lindblad operator $\mathcal{D}[X]\rho = X\rho X^\dagger - \frac{1}{2}X^\dagger X\rho - \frac{1}{2}\rho X^\dagger X$, $B = 2\chi\text{Re}(\alpha_g\alpha_e^*)$ is the measurement-induced detuning (Stark shift), and $\Gamma_d = 2\chi\text{Im}(\alpha_g\alpha_e^*)$ is the measurement-induced dephasing, with α_i the qubit-state-dependent photon field in the resonator and 2χ the qubit frequency shift per photon. At time $t - t_g$ after the qubit superposition is created,

$$\alpha_g\alpha_e^* = \alpha(t_m)\exp(-\kappa(t - t_m))\exp(2i\chi(t - t_g)),$$

with $t - t_m$ the time since the end of measurement excitation pulse. Integrating over the interval $[t_1, t_2]$ gives a dephasing term with coefficient

$$\begin{aligned} p_{\phi, \text{photon}} &= \exp\left(-\int_{t_1}^{t_2} \Gamma_d(t) dt\right) \\ &= \exp\left(2\chi\alpha(0) \exp(\kappa(t_m - t_g))\right. \\ &\quad \left.\times \left[\frac{e^{-\kappa t}}{4\chi^2 + \kappa^2} [-\kappa \sin(2\chi t) - 2\chi \cos(2\chi t)]\right]_{t_1 - t_g}^{t_2 - t_g}\right). \end{aligned}$$

This dephasing is then implemented via the same Pauli transfer matrix as (4.9).

4.B.3. Single-qubit $R_y(\pi/2)$ rotations

Single-qubit rotations are modeled by sandwiching an instantaneous Pauli transfer matrix, representing the rotation, with periods of duration $\tau_{g,1Q/2}$ of amplitude and phase damping. This allows to model the gate for different T_1 and T_ϕ . However, comparison of this model with Pauli transfer matrices obtained from gate-set tomography experiments shows that actual gates are more accurately described when adding a phenomenological depolarizing noise to the instantaneous part. In the Bloch sphere, this decay corresponds to shrinking toward the origin, with factor $1 - p_{\text{axis}}$ along the y axis and $1 - p_{\text{plane}}$ along the x - and z -axes. We thus model

$$R_{R_y(\pi/2)} = R_{AP(\tau_{g,1Q/2})} R'_{R_y(\pi/2)} R_{\text{dep}} R_{AP(\tau_{g,1Q/2})}, \quad (4.11)$$

where

$$R_{\text{dep}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 - p_{\text{plane}} & 0 & 0 \\ 0 & 0 & 1 - p_{\text{axis}} & 0 \\ 0 & 0 & 0 & 1 - p_{\text{plane}} \end{pmatrix},$$

and $R'_{R_y(\pi/2)}$ is the Pauli transfer matrix describing a perfect $\pi/2$ rotation around the y axis.

4.B.4. Flux noise

Shifting the transmon from its sweetspot $f_{q,\text{max}}$ to a lower frequency

$$f_q(t) = (f_{q,\text{max}} + E_C) \sqrt{|\cos(\pi\Phi(t)/\Phi_0)|} - E_C$$

4. Density-matrix simulation of small surface codes

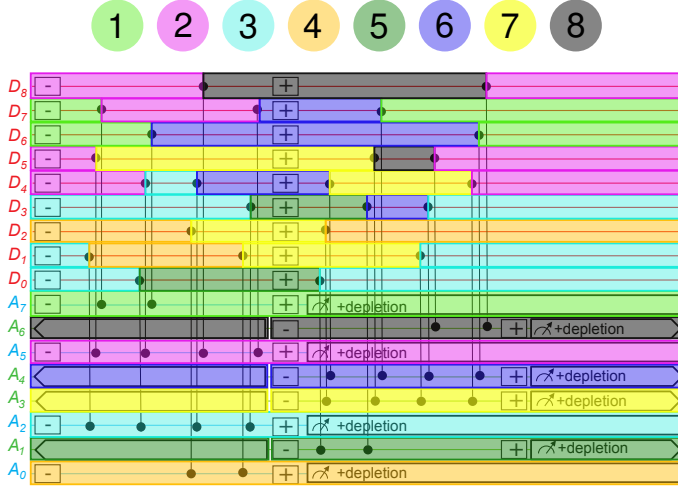


Figure 4.8.: Isolation of ancilla interactions in the Surface-17 circuit given in Fig. 4.7. Throughout a simulation, `quantsim` stores the density matrix of all data qubits. Each error correction cycle is split up into 8 steps as labeled. In each step, a single ancilla qubit is added to the density matrix, the correspondingly colored pieces of the circuit are executed, and the ancilla is read out and removed from the density matrix. This scheme is only possible because on each data qubit all gates are executed in order. Note that steps after the final C-Z gate on a data qubit are executed during the next cycle.

makes it first-order sensitive to flux noise, with sensitivity

$$\frac{\partial f_q}{\partial \Phi} = \frac{-\pi}{2\Phi_0} (f_q + E_C) \tan\left(\frac{\pi\Phi}{\Phi_0}\right).$$

Here, Φ is the flux bias and $\Phi_0 = h/2e$ is the flux quantum. For a deviation of $\delta\Phi$, the pulsed transmon incurs a phase error

$$\delta\phi = -2\pi\tau_{g,2Q} \frac{\partial f_q}{\partial \Phi} \delta\Phi.$$

Flux noise has a characteristic (single-sided) spectral density

$$S_{\Phi}(f) \approx A/f,$$

where $A \approx (1 \mu\Phi_0)^2$ with f in Hz. We model this noise as quasi-static over the duration ($1/f_{\min} \sim 20 \mu\text{s}$, or 20 QEC cycles) of individual runs,

but fluctuating between subsequent runs ($1/f_{\max} \sim 20$ sec, or 10^5 runs at $200 \mu\text{s}$ intervals). The root-mean-square (rms) fluctuations of flux are therefore

$$\begin{aligned}\delta\Phi_{\text{rms}} &= \left(\int_{f_{\min}}^{f_{\max}} S_{\Phi}(f) df \right)^{1/2} \\ &= A(\ln(f_{\max}/f_{\min}))^{1/2} \\ &\approx 4 \mu\Phi_0.\end{aligned}$$

For our quantum circuit based on [181], we estimate the corresponding rms phase error induced in a pulsed transmon to be

$$\delta\phi_{\text{rms}} \approx 0.01 \text{ rad.}$$

4.B.5. C-Z gates

We now focus on the two-qubit phase error. For an adiabatic gate,

$$\phi_{2\text{Q}} = \phi_{11} - \phi_{01} = -2\pi \int_{t_1}^{t_2} \zeta(t) dt,$$

with t_1 and $t_2 = t_1 + \tau_{\text{g},2\text{Q}}$ the start and end of the gate and ζ the time-dependent frequency deviation of the lower branch of the $|11\rangle \leftrightarrow |02\rangle$ avoided crossing from the sum of frequencies for $|01\rangle$ and $|10\rangle$. Near the flux center Φ_c of the $|11\rangle - |02\rangle$ avoided crossing,

$$\zeta \approx \beta(\Phi - \Phi_c) - \sqrt{\beta^2(\Phi - \Phi_c)^2 + (2J/2\pi)^2},$$

where $2J/2\pi \sim 50$ MHz is the minimum splitting between $|11\rangle$ and $|02\rangle$, and

$$\beta = \frac{1}{2} \frac{\partial f_{\text{q}}}{\partial \Phi} \Big|_{\Phi=\Phi_c}.$$

Differentiating with respect to Φ at Φ_c gives

$$\frac{\partial \zeta}{\partial \Phi} \Big|_{\Phi=\Phi_c} = \beta.$$

To estimate the $\delta\phi_{2\text{Q}}$ error, we make the following simplification: we replace the exact trajectory created by the flux pulse by a shift to $\Phi = \Phi_c + \delta\Phi$ with duration $\tau_{\text{g},2\text{Q}}$. For a deviation of $\delta\Phi$,

$$\delta\phi_{2\text{Q}} \approx -2\pi\tau_{\text{g},2\text{Q}} \frac{\partial \zeta}{\partial \Phi} \Big|_{\Phi=\Phi_c} \delta\Phi.$$

4. Density-matrix simulation of small surface codes

Probability	Value	Probability	Value
$\epsilon_0^{+1,0}$	0.9985	$\epsilon_1^{+1,0}$	0.0050
$\epsilon_0^{+1,1}$	0.0000	$\epsilon_1^{+1,1}$	0.0015
$\epsilon_0^{-1,0}$	0.0015	$\epsilon_1^{-1,0}$	0.0149
$\epsilon_0^{-1,1}$	0.000	$\epsilon_1^{-1,1}$	0.9786

Table 4.3.: Measurement butterfly matching a recent characteristic experiment [182] using a Josephson parametric amplifier [203] in phase-preserving mode as the front end of the readout amplification chain.

Note that this two-qubit phase error is correlated with the single-qubit phase error on the fluxed transmon. The former is smaller by a factor ≈ 2 .

4.B.6. Measurement

The probabilities $\epsilon_i^{m_a, m_b}$ are calibrated using the statistics of outcomes in back-to-back measurements (a followed by b) with the qubit initialized in $|i\rangle$.

$$\begin{aligned}
 P(m_a = +1)_i &= \epsilon_i^{+1,0} + \epsilon_i^{+1,1}, \\
 P(m_a = -1)_i &= \epsilon_i^{-1,0} + \epsilon_i^{-1,1}, \\
 P(m_b = m_a = +1)_i &= \left(\epsilon_0^{+1,0} + \epsilon_0^{+1,1} \right) \epsilon_i^{+1,0} \\
 &\quad + \left(\epsilon_1^{+1,0} + \epsilon_1^{+1,1} \right) \epsilon_i^{+1,1}, \\
 P(m_b = -m_a = +1)_i &= \left(\epsilon_0^{+1,0} + \epsilon_0^{+1,1} \right) \epsilon_i^{-1,0} \\
 &\quad + \left(\epsilon_1^{+1,0} + \epsilon_1^{+1,1} \right) \epsilon_i^{-1,1}, \\
 P(-m_b = m_a = +1)_i &= \left(\epsilon_0^{-1,0} + \epsilon_0^{-1,1} \right) \epsilon_i^{+1,0} \\
 &\quad + \left(\epsilon_1^{-1,0} + \epsilon_1^{-1,1} \right) \epsilon_i^{+1,1}, \\
 P(-m_b = -m_a = +1)_i &= \left(\epsilon_0^{-1,0} + \epsilon_0^{-1,1} \right) \epsilon_i^{-1,0} \\
 &\quad + \left(\epsilon_1^{-1,0} + \epsilon_1^{-1,1} \right) \epsilon_i^{-1,1}.
 \end{aligned}$$

We obtain the six free parameters of the black-box description from these 12 equations, using experimental values on the left-hand side [202]. Table 4.3 shows the values used, achieved in a recent experiment [182]. For the simulation, we reproduce this behaviour of the measurement process by a model with several steps. The qubit undergoes dephasing, followed by periods of decay or excitation between which the measurement result is




Figure 4.9.: The model for measurements consists of a dephasing of the qubit followed by a period of decay and excitation with probability $p_{\downarrow/\uparrow}^{(1)}$. At this point, the qubit state is sampled. The sampling result is subject to a declaration error ϵ_{RO} , and the qubit state is subject to further decay or excitation with probabilities $p_{\downarrow/\uparrow}^{(2)}$ before the end of the measurement block.

sampled. This measurement result is further subject to a state-dependent declaration error ϵ_{RO} before reported to the decoder (see Fig.4.9). The six parameters of this model are in a one-to-one correspondence with the butterfly parameters described above, and can be mapped by solving the corresponding system of equations. The experimental results in Tab.4.3 are very well explained by assuming unmodified amplitude-phase damping (with the zero excitation probabilities) during the measurement period, and an outcome-independent declaration error of $\epsilon_{\text{RO}} = \epsilon_{\text{RO}}^1 = \epsilon_{\text{RO}}^0 = 0.15\%$. We use this result to extrapolate measurement performance to different values of T_1 .

Reduction of measurement time is expected to reduce assignment fidelity. For the results presented in Fig. 2, we do not rely on experimental results, but assume a simplified model for measurement, following Ref. 186. A constant drive pulse of amplitude ϵ and tuned to the bare resonator frequency, $\Delta_r = 0$, excites the readout resonator for time τ_m . The dynamics of the resonator is dependent on the transmon state (we approximate linear behavior), and the transmitted signal is amplified and detected in a homodyne measurement as a noisy transient. This transient is processed by a linear classifier, which declares the measurement outcome. For resonator depletion, we use a two-step clearing pulse with amplitude ϵ_{c1} and ϵ_{c2} , each active for $\tau_d/2$ and chosen (by numerical minimization) so that, at the end of the depletion pulse, the transients for both transmon states return to zero. While the resonator dynamics is easily found if the transmon is in the ground state, amplitude damping of the transmon in the excited

4. Density-matrix simulation of small surface codes

state leads to non-deterministic behavior. We thus numerically obtain an ensemble of noisy transients for each input qubit state, and optimize the decision boundary of the linear classifier for this ensemble. Generating a second verification ensemble, the “butterfly” of the measurement setup is estimated.

The dynamics of the resonator is determined by the resonator linewidth κ as well as the dispersive shift χ . We chose the parameters of the setup used in [182], $1/\kappa = 250$ ns and $\chi/\pi = -2.6$ MHz. The signal-to-noise ratio of the detected transient is reduced by the quantum efficiency $\eta = 12.5\%$. The driving strength ϵ is chosen to approximate the “butterfly” used in most of the main text, and corresponds to a steady-state average photon population of about $\bar{n} = 15$. We then keep ϵ constant while changing the measurement time, keeping $\tau_m = \tau_d$, to obtain the butterflies used in the density matrix simulation. We ignore effects leading to measurement-induced mixing and non-linearity of the readout resonator. Finally, since these simulations do not allow to make a realistic prediction about residual photon numbers achievable in experiments, we ignore this effect when using these results.

4.C. Effect of over-rotations and two-qubit phase noise on logical error rate

In this section we provide additional numerical data showing the effect of some common noise sources on the logical error rate. In Fig. 4.10 we show the effect of a coherent over-rotation, whereby the $R'_Y(\pi/2)$ operator in Eq. 4.11 is replaced by $R'_Y(\pi/2 + \delta\phi)$. This can be caused by inaccurate calibration of the flux pulse used to perform the gate. In Fig. 4.11 we show the effect of an increase in the two-qubit flux noise $\delta\phi_{\text{rms}}$ as described in Sec. 4.B.4.

4.D. Calculation of decoder upper bound

We provide a detailed description how the decoder upper bound is obtained from the simulation results. As described in the main text, after each cycle of simulation, the diagonal of the reduced density matrix of the data qubits in the Z basis is stored. It contains the probability distribution for the $2^9 = 512$ different possible measurement outcomes of the data qubits. In the quantum memory experiment described in the main text, each of

4.D. Calculation of decoder upper bound




Figure 4.10.: Logical error rate for Surface-17 as a function of single-qubit over-rotation, using the MWPM decoder. Other parameters are as given in the main text.

these outcomes are passed to the decoder, which then declares a logical measurement outcome.

It is evident that any decoder must declare opposite logical outcomes if two of the 512 possible measurements m and m' are related by the application of a logical X operator. Thus, any decoder can give the correct result only for half of the measurement outcomes. Subject to this constraint, we can find the set of 256 declarations which maximize the probability that the declaration is correct. It immediately follows that no decoder can achieve a declaration fidelity larger than this maximal probability. We thus refer to it as the decoder upper bound.

In practice, the upper bound is found according to the following approach. Since declarations are opposite if two outcomes differ by a logical X operator, they must be equal if they differ by the application of one or more X stabilizers (applying two different logical X operators amounts to the application of a product of X stabilizers). We thus group the outcomes in 32 cosets which are related by the application of X -stabilizers. (There are 4 X -stabilizers in Surface-17, so there are $512/2^4 = 32$ cosets). For outcomes from the same coset, the declaration from a decoder must be the same. We obtain the probability of a final measurement falling within each coset by summing the probabilities from the density matrix diagonal. We further group the 32 cosets to 16 pairs, which differ by the application

4. Density-matrix simulation of small surface codes

of a logical operator. The upper bound is then obtained by selecting the more probable coset from each pair and summing the corresponding probabilities. This upper bound can also be interpreted as the internal decoherence of the logical qubit: it represents the maximal overlap of the final state with the initial state, under any possible correction of errors.

We finally emphasize that this upper bound can be found only because we have access to the complete probability distribution of outcomes (for a given result of syndrome measurements), a major advantage of the density matrix simulation. However, we do not expect that any decoder can actually achieve this upper bound: This is because we add syndrome measurement events independently after the situation, which will decrease the logical error rate further.




Figure 4.11.: Logical error rate for Surface-17 as a function of two-qubit phase error, using the MWPM decoder. Other parameters are as given in the main text.

4.E. Hardware requirements of simulation

The simulations are performed using the `quantumsim` package *, which were developed by the authors for this chapter. The package is accelerated by performing the density matrix manipulations on a GPU (graphics

*The `quantumsim` package can be found at <http://github.com/brianzi/quantumsim>

4.F. Homemade MWPM decoder with asymmetric weight calculation

card). The simulations for this chapter were performed on a NVidia Tesla K40 GPU, on which we observed runtimes of about 0.5 seconds for the simulation of a run of $k=20$ cycles (25 ms per QEC cycle). We also had the opportunity to test the software on a more modern GPU (NVidia Tesla P100), observing about 15 ms per cycle, and on a consumer-grade GPU (NVidia Quadro M2000), observing about 40 ms per cycle. By comparison, the CPU is mostly idle during the simulation, except for handling of input and output. The memory requirements are modest for both CPU and GPU RAM. They are dominated by the storage of the density matrices and amount to a few ten megabytes.

4.F. Homemade MWPM decoder with asymmetric weight calculation

Every QEC code requires a decoder to track the most likely errors consistent with a given set of stabilizer measurements. The MWPM decoder has gained popularity since it was shown to have threshold values above 1% [177]. The motivation behind MWPM is that single X or Z errors on data qubits in the bulk of a surface-code fabric cause changes of two stabilizers in the code. These signals can then be considered vertices on a graph, with the error the edge connecting them. Errors in measurement, or errors on a single ancilla qubit, behave as changes in the stabilizer that are separated in time. Multiple errors that would join the same vertices create longer paths in the graph, of which an experiment only records the endpoints. Thus, the problem becomes that of finding the most likely set of generating errors given the error signals that mark their ends. This is made slightly simpler, as in the surface code any chain of errors that forms a closed loop does not change the logical state. This implies that all paths that connect two points are equivalent, and can be considered together. The problem then is to join error signals, either in pairs, or to a ‘boundary’ vertex. The latter corresponds to errors on data qubits at the boundary, which belong to only one X or Z stabilizer. This pairing P should be chosen as the most likely combination of single-qubit errors that could generate the measured error signals. This has then been reduced to the problem of minimum-weight perfect matching on a graph, which can be solved in polynomial time by the blossom algorithm [173, 209].

The MWPM decoder we use differs from previous methods by its weight calculation. As part of the decoding process, it is required to calculate to some degree of accuracy [204] the probability p_{e_1, e_2} of two measured

4. Density-matrix simulation of small surface codes

error signals e_1 and e_2 being connected by a chain of individual logical errors. This is then converted to a weight $w_{e_1, e_2} = -\log(p_{e_1, e_2})$, which form the input to the blossom algorithm of Edmonds to find the most likely matching of error signals [173, 209]. An exact calculation of p_{e_1, e_2} requires a sum over all such chains between e_1 and e_2 that do not cross the boundary (these are equivalent modulo stabilizer operators that do not change the logical state). In this appendix we detail a method of computing this sum, and approximations to make it viable within the runtime of the experiment.

Let us define the ancilla graph $\mathcal{G}_A = (V_A, E_A)$ containing a vertex $v \in V_A$ for every ancilla measurement, and an edge $e \in E_A$ connecting $v, u \in V_A$ if a single component (gate, single-qubit rest period, or faulty measurement) in the simulation can cause the u and v measurements to return an error. We include a special ‘boundary’ vertex v_B , to which we connect another vertex v if single components can cause errors on v alone. Then, to each edge e we associate a probability p_e , being the sum of the probabilities of each component causing this error signal. These error rates can be obtained directly from `quantumsim`, by cutting the circuit at each C-Z gate and measuring the decay of single qubits between. Then, for a given experiment with given syndrome measurements, let us define the syndrome graph $\mathcal{G}_S = (V_S, E_S)$ containing a vertex $v \in V_S$ for each syndrome measurement that records an error, and an edge $\lambda_{u, v} \in E_S$ connecting $u, v \in V_S$ if u and v are either both X ancilla qubits or both Z ancilla qubits. To each edge $\lambda_{u, v}$ we associate a probability $p_{u, v}$ given by the sum of the probabilities of a chain of errors causing error signals solely on u and v .

If we assume that single-qubit errors are uncorrelated, we have to lowest order

$$p_{u, v} \approx \sum_{\text{paths } (e_1, e_2, \dots, e_n) \text{ between } u \text{ and } v} \prod_{j=1}^n p_{e_j}, \quad (4.12)$$

Let A_A be the adjacency matrix on \mathcal{G}_A weighted by the probabilities p_e (i.e., $(A_A)_{u, v} = p_e$ with e connecting u and v), and A_S the same for \mathcal{G}_S . Then, the above becomes

$$A_S = A_A + A_A^2 + A_A^3 + \dots = \frac{\mathbf{1}}{\mathbf{1} - A_A} - \mathbf{1}, \quad (4.13)$$

noting that A_S contains a subset of the indices that are used to construct A_A .

The boundary must be treated specially in the above calculation. For

4.F. Homemade MWPM decoder with asymmetric weight calculation

the purposes of the surface code, the boundary can be described as a single vertex which has no limit on the number of other vertices it may pair to [173]. For the purposes of weight calculation, any path that passes through the boundary is already counted by pairing both end vertices to the boundary. This can be treated by making \mathcal{G}_A directed, and breaking the symmetry $A_A^T = A_A$. In particular, either $(A_A)_{v_B,u} = 0$ for all u or $(A_A)_{u,v_B} = 0$ for all u .

The above calculation requires inversion of a $N_{\text{mat}} \times N_{\text{mat}}$ matrix, with N_{mat} the total number of ancilla measurements per experiment. Furthermore, as ancilla error rates depend upon the previous ancilla state, elements in A_A are not completely known until the previous cycle. This implies that in an actual computation with runtime decoding, this inversion would need to be completed within a few microseconds (with a transmon-cQED architecture), which is practically unfeasible. We suggest two approximations that can be made to shorten the decoding time. The first is to average all errors over the ancilla population, ignoring any asymmetry in the system. The adjacency matrix is now the same for any experiment, and can be precalculated and stored as a look-up table for the run-time decoder. We call this the decoder with symmetrized weights. The size of such a look-up table scales poorly with the number of qubits and the number of cycles. However, $(A_S)_{u,v}$ is approximately invariant under simultaneous translation of u and v (excluding boundary effects). This implies that a precalculated A_S can be vastly compressed, making this method feasible.

The second approximation to the full A_S calculation is to perform it iteratively. We divide our graph \mathcal{G}_A (\mathcal{G}_S) by time steps; let \mathcal{G}_A^t (\mathcal{G}_S^t) be the subgraph of \mathcal{G}_A (\mathcal{G}_S) containing only ancillas measured before time step t , and let $\partial\mathcal{G}_A^t$ ($\partial\mathcal{G}_S^t$) be the subgraph of \mathcal{G}_A (\mathcal{G}_S) containing only ancillas measured during time step t . Then, if we assume we have an approximation to the matrix A_S^t (being the adjacency matrix of \mathcal{G}_S^t), we can approximate

$$A_S^{t+1} \approx \begin{pmatrix} A_S^t & C_S^{t+1} \\ (C_S^{t+1})^T & (\mathbf{1} - \partial A_A^{t+1})^{-1} \end{pmatrix} \quad (4.14)$$

to lowest order in physical errors. Here, ∂A_A^{t+1} is the weighted adjacency matrix on $\partial\mathcal{G}_A^{t+1}$, and the coupling matrix C_S^{t+1} is approximated by

$$C_S^{t+1} = A_S^t C_A^{t+1} (\mathbf{1} - \partial A_A^{t+1})^{-1}, \quad (4.15)$$

4. Density-matrix simulation of small surface codes

with C_A^{t+1} the adjacency matrix containing only edges between $\partial\mathcal{G}_A^{t+1}$ and \mathcal{G}_A^t . This procedure corresponds to a sum over all paths that are made by moving within $\partial\mathcal{G}_A^{t+1}$, shifting back in time to \mathcal{G}_A^t , and then taking any precalculated path in \mathcal{G}_A^t . C_A^{t+1} and $(\mathbf{1} - \partial A_A^{t+1})^{-1}$ can be precalculated, and so the runtime computation requirement is reduced to the product in Eq. (4.15). This in turn can be sparsified, as C_A^{t+1} only contains connections to vertices in \mathcal{G}_A^t close to the time boundary, and we can delete all terms in A_S^t that do not connect from these vertices to errors.

We have used the second method for our MWPM decoder, as we expect the error from neglecting higher-order combinations of errors to be small. In order to check this assumption, in Fig. 4.12 we repeat our simulation protocol with a modified physical error model that excludes all Y and measurement errors. We see that in the absence of these errors, the MWPM decoder performs within the error margin of the decoder upper bound. Note that a small deviation is expected from the discrepancy between a MWPM decoder and a maximum-likelihood decoder [178]. With the parameters used in this chapter, we do not observe any loss of fidelity when we stop accounting for the difference in error rates between ancilla states. We account this to the large error contribution from photon noise and gate infidelity on the ancilla qubits, which do not have this asymmetry. We further note that we operate in a regime of large ancilla error; as described in the text this makes the system counter-intuitively less sensitive to ancilla noise. In systems where this is not the case, it could be that accounting for ancilla asymmetry provides a useful computational method to improve ϵ_L .

4.G. Implementation of a look-up table decoder

In [176], the authors describe a decoding scheme specific to Surface-17, which is optimized to be implementable with limited computational resources in a short cycle time. This decoding scheme works by using a short decision tree to connect errors to each other in a style similar to blossom. Indeed, this scheme is equivalent to a blossom decoder with all horizontal, vertical and diagonal weights equal [176]. As such, we have implemented the new weights in the blossom decoder rather than utilizing the exact method given.

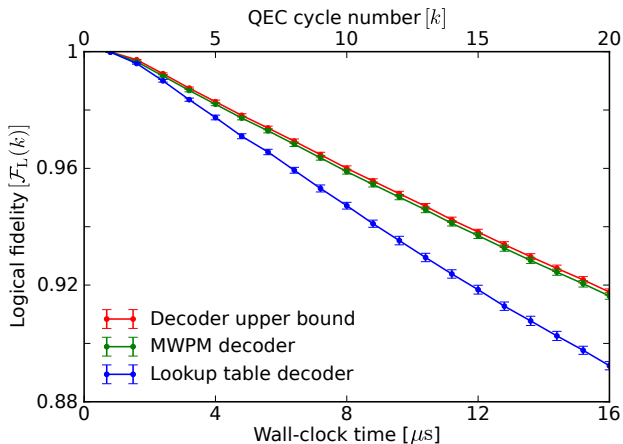


Figure 4.12.: Simulation of the experimental protocol used throughout the chapter, but using an error model that has Y errors and readout infidelity removed. With these errors absent, the MWPM decoder achieves the decoder upper bound within simulation error. The look-up table approach (blue) retains some inaccuracy beyond this.

4.H. Details of lowest-order approximation

We detail the approximation made to study Surface-49 in Sec. II C. Note that this calculation is only for X errors, which are measured by the Z ancillas. This implies that our approximation should attempt to realize the result of blossom, rather than the decoder upper bound.

We begin with the \mathcal{G}_A graph defined in App. 4.F. In the absence of correlated errors that cause more than two error signals, any experiment can be approximately described by choosing a set $S \subset E_A$ of edges on the graph and assuming the errors that correspond to these edges have occurred. Each ancilla measurement corresponds to a vertex in \mathcal{G}_A , which records an error if an odd number of edges in S point to the vertex. Each combination M_a of ancilla measurements can be generated by multiple error sets S .

Formally, let us write \mathcal{M} for the set of all combinations of ancilla measurements and \mathcal{S} for the set of all combinations of errors (so $\mathcal{S} = 2^{E_A}$). We then define a function $\phi : \mathcal{S} \rightarrow \mathcal{M}$ that takes a combination of errors to the resultant measurement outcomes. Let us fix a logical Z operator Z_L on the surface-code fabric. Then to each $S \in \mathcal{S}$ we can assign a parity

4. Density-matrix simulation of small surface codes

$p(S) = \pm 1$ depending on whether the product of all errors in S commute with Z or not. A decoding then consists of a choice of parity $p_d(M)$ for each $M \in \mathcal{M}$. Such a decoding correctly decodes $S \in \mathcal{S}$ if $p_d(\phi(S)) = p(S)$, and creates a logical error otherwise. The source of logical errors in a perfect decoder is then precisely the fact that we can have two error combinations $S_1, S_2 \in \mathcal{S}$ such that $\phi(S_1) = \phi(S_2)$ but $p(S_1) \neq p(S_2)$.

The above suggests a method by which a perfect decoder can be constructed. As defined, $\phi^{-1}(M) \subset \mathcal{S}$ is the set of error combinations S that return a measurement $M \in \mathcal{M}$. For each error combination S , we can calculate the probability of this occurring:

$$r(S) = \prod_{e \in S} p_e \prod_{e \notin S} (1 - p_e). \quad (4.16)$$

The optimal choice of $p_d(M_a)$ is the one maximizing

$$\sum_{S \in \phi^{-1}(M), p(S)=p_d(M)} r(S), \quad (4.17)$$

and the fidelity of such a decoder (over the entire experiment) can be calculated as

$$\mathcal{F}_L = 1 - \sum_{M \in \mathcal{M}} \min \left(\sum_{S \in \phi^{-1}(M)} \delta_{p(S), +1} r(S), \sum_{S \in \phi^{-1}(M)} \delta_{p(S), -1} r(S) \right). \quad (4.18)$$

At this point the only approximation that has been made is to neglect the T_1 asymmetry in the system, which we have shown previously in this chapter to be negligible. Unfortunately, the above function cannot be evaluated exactly; the number of error combinations S is approximately 2^{200} for 4 cycles of Surface-49. Our goal instead is to approximate this to the lowest order in the physical qubit error rate.

Let us make the approximation that our error combinations S can be split into small, well-separated pieces of errors containing separate correctable and non-correctable parts, $S = \cup_i S^i$. To each S^i we can assign a time step $t(S^i)$, being the earliest time of the first error measurement observed (in $\phi(S^i)$). The error rate per round, ϵ_L , can be determined by summing Eq. 4.18 over all pieces S^i of all combinations S such that $t(S^i) = T$ (with

4.H. Details of lowest-order approximation

arbitrary T), as the effect of repeated errors from $S^i, S^j \subset S$ is taken into account during the derivation of the logical fidelity equation (Eq. 2 in the main text).

$$\epsilon_L = \sum_{M \in \mathcal{M}} \min \left(\sum_{S \in \phi^{-1}(M)} \sum_{t(S^i)=T} \delta_{p(S^i),+1} r(S), \right. \\ \left. \sum_{S \in \phi^{-1}(M)} \sum_{t(S^i)=T} \delta_{p(S^i),-1} r(S) \right). \quad (4.19)$$

Let us also extend the above division of S to a division of M into separate pieces M^a , and rewrite our sum slightly,

$$\epsilon_L = \sum_{M^a} \min \left(\sum_{S^i \in \phi^{-1}(M^a), t(S^i)=T} \delta_{p(S^i),+1} \bar{r}(S^i), \right. \\ \left. \sum_{S^i \in \phi^{-1}(M^a), t(S^i)=T} \delta_{p(S^i),-1} \bar{r}(S^i) \right), \quad (4.20)$$

Where here we have brought the sum over the global combinations of syndromes and measurements inside a new function \bar{r}

$$\bar{r}(S^i) = \prod_{e \in S^i} p_e \sum_{M \supset M^a} \sum_{(S \supset S^i, S \in \phi^{-1}(M))} \\ \prod_{f \in S/S^i} p_f \prod_{g \notin S^i} (1 - p_g) \\ = \prod_{e \in S^i} p_e \sum_{S \supset S^i} \prod_{f \in S/S^i} p_f \prod_{g \notin S^i} (1 - p_g) \quad (4.21)$$

If we took this approximation literally and considered the sum over every possible combination S containing S^i , the final sum in Eq. 4.21 would reduce to

$$\bar{r}^{(u)}(S^i) = \prod_{e \in S^i} p_e. \quad (4.22)$$

However, this includes error combinations S that cannot be easily separated into S^i and ‘something else’, i.e. they contain other errors e that cannot be separated from S^i . Eq. 4.22 is then equivalent to assuming that if S^i is an uncorrectable logical error, no nearby combination of physical errors

4. Density-matrix simulation of small surface codes

S' can be combined such that $S^i \cup S'$ is correctable unless S' itself is an uncorrectable logical error. Such combinations would serve to reduce the calculated ϵ_L , and so $\bar{r}^{(u)}$ gives an upper bound for ϵ_L in Eq. 4.20. For a lower bound, we approximate that for any uncorrectable error combination S^i , approximately one rounds-worth of single errors would undo the logical error, leading to the approximation

$$\bar{r}^{(l)}(S^i) = \prod_{e \in S^i} p_e \prod_{t(\{e\})=T} (1 - p_e). \quad (4.23)$$

We now make one further approximation, and sum Eq. 4.20 only over the shortest S^i that can be expected to contribute to the final error rate. That is, we sum over those S^i with $|S^i| \leq (d+1)/2$, and that spread directly across the chain. The error incurred from this approximation is roughly proportional to the largest single error, which is no more than 5% throughout our study. We use $\bar{r}^{(u)}$ and $\bar{r}^{(l)}$ to give the error bars shown in Fig. 5. Points in the plot are taken as a log average of the upper and lower bounds, and thus have no particular relevance themselves. We see that the numerical calculation falls within the corresponding error bars for almost the entire dataset, giving verification for our method, save a slight deviation at one point where it falls below. Moreover, as the simulated Surface-17 error rate lies above the upper bound found for the Surface-49 error rate (with the standard set of parameters from the main text), our claim that Surface-17 will operate below the fault-tolerant threshold is quite strong.