



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Calibrated Decoders for Experimental Quantum Error Correction

Edward H. Chen, Theodore J. Yoder, Youngseok Kim, Neereja Sundaresan, Srikanth Srinivasan, Muyuan Li, Antonio D. Córcoles, Andrew W. Cross, and Maika Takita

Phys. Rev. Lett. **128**, 110504 — Published 17 March 2022

DOI: [10.1103/PhysRevLett.128.110504](https://doi.org/10.1103/PhysRevLett.128.110504)

Calibrated decoders for experimental quantum error correction

Edward H. Chen,^{1, a} Theodore J. Yoder,^{2, b} Youngseok Kim,² Neereja Sundaresan,² Srikanth Srinivasan,² Muiyuan Li,² Antonio D. Córcoles,² Andrew W. Cross,² and Maika Takita²

(IBM Quantum)

¹*Almaden Research Center, San Jose, CA 95120, USA*

²*T.J. Watson Research Center, Yorktown Heights, NY 10598, USA*

(Dated: February 11, 2022)

Arbitrarily long quantum computations require quantum memories that can be repeatedly measured without being corrupted. Here, we preserve the state of a quantum memory, notably with the additional use of flagged error events. All error events were extracted using fast, mid-circuit measurements and resets of the physical qubits. Among the error decoders we considered, we introduce a perfect matching decoder that was calibrated from measurements containing up to size-4 correlated events. To compare the decoders, we used a *partial* post-selection scheme shown to retain ten times more data than *full* post-selection. We observed logical errors per round of $2.2 \pm 0.1 \times 10^{-2}$ (decoded *without* post-selection) and $5.1 \pm 0.7 \times 10^{-4}$ (*full* post-selection), which was less than the physical measurement error of 7×10^{-3} and therefore surpasses a pseudo-threshold for repeated logical measurements.

Keywords: Quantum Information Processing, Quantum Memories, Quantum Correlations, Quantum Error Correction, Quantum Protocols

Introduction.— Preparing and preserving logical quantum states is necessary for performing long quantum computations [1]. Because noise inevitably corrupts the underlying physical qubits, quantum error correction (QEC) codes have been designed to detect and recover from errors [2–6]. Significant efforts are currently focused on demonstrating capabilities that will be necessary for implementing practical QEC. An optimal choice of a code varies depending on the device and its noise properties [7]. Notable experimental implementations include NMR [8, 9], ion traps [10–13], donors [14–16], quantum dots [17, 18], and superconducting qubits [19–23]. Recent developments of high-fidelity mid-circuit measurements and resets of superconducting qubits have enabled the preparation and repeated stabilization of logical states [24–26]; demonstrations of such quantum memories with enhanced lifetimes have been limited by, among other reasons, a combination of gate and measurement cross-talk.

One way to mitigate cross-talk [27] is to reduce the lattice connectivity [28, 29]. Consequently, fault-tolerant operations require intermediary qubits; such qubits can be used to *flag* high-weight errors originating from low-weight errors [30, 31]. In certain QEC codes and lattice geometries, flag qubits supply the information needed to extend the effective distance of a QEC code up to its intended distance, and thus enable maximal efficiency at detecting and correcting errors [32].

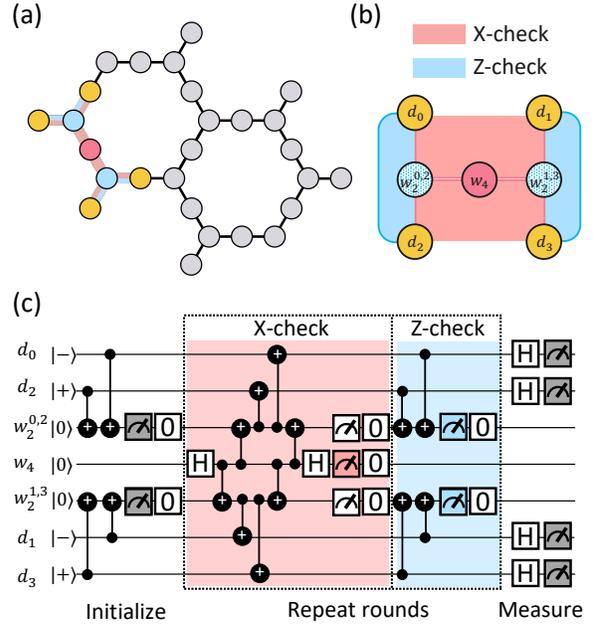


FIG. 1. (Color) (a) Experiments were performed on *ibmq_kolkata*, which had 27 qubits connected in a heavy hexagon (HH) topology. The 7 qubits used for the $[[4, 1, 2]]$ code are colored yellow, blue and pink. (b) The code layout indicates a single weight-4, X stabilizer (pink), and two weight-2, Z stabilizers (blue) on the four data qubits (yellow) labelled d_i for integers ‘i’ from 0 to 3. For the weight-2 stabilizers, superscripts ‘0,2’ (‘1,3’) indicates the data qubits on which they operate. The reduced connectivity of the graph is addressed by flag qubits (blue) alternating between (i) being used as weight-2 stabilizers, and (ii) as intermediary qubits used to detect errors on the center, syndrome qubit (pink). (c) Circuit diagram for the code layout in (b) applied to an initial $|-\rangle_L$ logical state with repeated X- (pink), flag (white) and Z-check (blue) stabilizer measurements, together comprising a round, with mid-circuit reset operations (‘0’) applied between rounds. In this illustration, the final measurement measures the four data qubits in the X-basis (gray).

^a Equally contributing spokesperson ; ehchen@ibm.com

^b Equally contributing spokesperson ; ted.yoder@ibm.com

We demonstrated repeated error detection and correction of a $[[4, 1, 2]]$ error-detecting topological stabilizer code on a heavy-hexagonal (HH) device designed to mitigate the limiting effects of cross-talk using flag qubits. The combination of fast readout with reduced qubit connectivity improved, after post-selecting on instances in which no errors were detected, logical errors per round when compared to the physical measurement error rate. A thorough analysis of this code led us to introduce a *partial* post-selection scheme allowing us to discard ten times less data for comparing matching decoding algorithms. Compared against previously known decoding strategies on the entire data set, we found that a decoder performed best with experimentally-calibrated edge weights that account for the correlations between syndromes; we show that the computational cost of calibrating such a decoder scales linearly with the size of the syndrome extraction circuit for topological codes due to the local nature of the parity checks, and thus can be extended to larger distance codes even without any post-selection. Furthermore, we showed that correlations between five or more syndromes can be eliminated by the application of a “deflagging” procedure. The minimal impact of deflagging on logical errors is an encouraging sign that this technique, and its extension to general flag-based codes, is a viable way to process flag outcomes in practice.

Theory of correlation analysis, decoding, deflagging, and post-selection.— Active error-correction involves decoding, using syndrome measurements, the errors that occurred in the circuit so that the proper corrections can be applied. We define error-sensitive events to be linear combinations of syndrome measurement bits that, in an ideal circuit, would be zero. Thus, a non-zero error-sensitive event indicates some error has occurred. For the HH code, there are two types of error-sensitive events defined as: (1) the difference of two subsequent measurements of the same stabilizer, and (2) flag qubit measurements.

Error-sensitive events are depicted as nodes in a decoding graph with edges representing errors that are detected by both events at their end points (Figure 2(a) and [33]). If the probability an edge occurs is P , then the edge is given weight $\log((1 - P)/P)$. The decoding graph may also have a boundary node, so that an error detected by just one error-sensitive event can be represented as an edge from that event to the boundary node. In practice, there are also errors detected by more than two error-sensitive events that could be represented as *hyperedges* in a more general decoding *hypergraph*.

Given a set of non-zero error-sensitive events, minimum-weight perfect-matching (MWPM) finds paths of edges connecting pairs of those events with minimum total weight, and is a simple and effective decoding algorithm for a topological stabilizer code that only operates on a decoding *graph* [34], as opposed to a decoding *hypergraph*. While MWPM is computationally efficient, the analogous matching algorithm on a hypergraph is not, which limits the practicality of a decoding hypergraph.

The effectiveness of MWPM depends crucially on edge weights in the decoding graph. We explored three strategies for setting these edge weights: (1) In the *uniform* approach, all edge weights were identical. (2) In the *analytical* approach, edge weights were individually calculated in terms of Pauli error rate parameters ρ_j , where the index j indicates one of the six errors being considered: CNOT gates, single-qubit gates, idle locations,

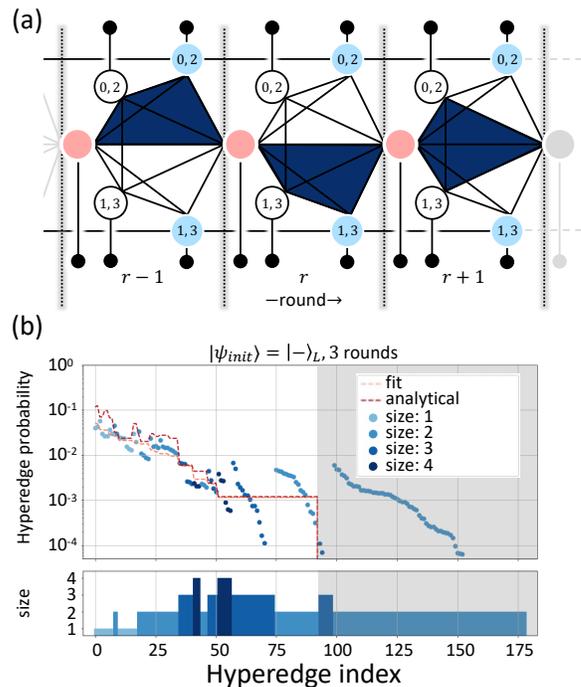


FIG. 2. (Color) (a) Decoder graph for the code layout depicted in Figure 1(b). Syndrome measurements from weight-4 (2) stabilizers are mapped to the pink (blue) nodes, and the weight-2 flag measurements are mapped onto the white nodes. Identical to Figure 1, ‘0,2’ (‘1,3’) denotes the left (right) hand side of the code layout, with the colors of the circuits and syndrome measurement matching in both figures. For initial $| - / + \rangle_L$ states stabilized by the circuit in Figure 1(c), there are three different possible size-4 hyperedges within each round, each highlighted in dark blue across three consecutive rounds. The boundary nodes in black have, by definition, edges with weight ‘0’ connecting them; rendering all boundary nodes to be effectively a single node for the purposes of the decoding process. (b) Applying the technique introduced in the main text to experiment in Figure 1(c), we estimated the correlation probabilities for all of the hyperedges shown in (a). The probabilities are sorted from largest to smallest based on the results from a least-squares fit using a six-parameter noise model. Points with darker colors represent hyperedges of greater sizes, as shown in the lower half of the plot. Hyperedges with indices greater than 93 (shaded gray) had no analytical expression, but were still experimentally adjusted to quantify the impact of computational leakage. The result of fitting the six-parameter noise model (fit, pink dash) agreed well with the analytical (red dash) curve generated using noise terms from simultaneous randomized benchmarking).

initialization, resets, and measurements. The numerical values of the parameters ρ_j can be chosen in several ways as discussed in the next section on *logical error measurements and decoding*. (3) In the *correlation* approach, we analyzed experimental data to determine a set of edge probabilities that are likely to have produced it. This approach involved first calculating the probabilities for all hyperedges in the decoding *hypergraph* before determining the edge probabilities used in the decoder *graph*. Importantly, this calibrated decoder graph, informed by the experimentally estimated hyperedge probabilities, can be used for any fault-tolerant protocol, and is not limited to the HH code or when post-selection is employed.

The key idea of the estimation protocol is as follows: a hyperedge in the decoding hypergraph represents any of a number of Pauli faults in the circuit that are indistinguishable from one another because they each lead to the same set of non-zero error-sensitive events. If several faults occur together, the symmetric difference of their hyperedges is denoted S , the syndrome, or, in other words, the set of non-zero error-sensitive events that is observed. The probability that we observe a particular S is the probability that hyperedges occur in combination to produce S . Since this is related to the probability α_h of an individual hyperedge h occurring, we can learn α_h from many observations of S .

Realistically, the possible hyperedges are limited in size $|h|$ by locality of the circuits. In the $[[4, 1, 2]]$ code, we found that hyperedges are limited to sizes four or less. Finding α_h in practical time begins by considering local clusters and then adjusting local estimates recursively from size-4 hyperedges down to size-1 and -2 (Figure 2 and [35]). Only size-1 and -2 edges are required for MWPM, but ignoring larger hyperedges can result in nonphysical, negative size-1 correlations. Another way we explored decoding strategies was to consider analyzing only a subset of all data. By Pauli tracing, we classified edges in the decoding graph into three categories depending on whether its inclusion in the minimum-weight matching necessitated: (1) flipping the logical measurement, (2) not flipping the logical measurement, or (3) is ambiguous [36]. The ambiguous case occurs specifically for error-detecting codes, like the $[[4, 1, 2]]$ code presented here, because some errors result in the decoder having to choose between two equally probable corrections.

Using these classifications for edges in the decoder graph, we explored three degrees of post-selection. The most conservative approach, using *full* post-selection, involved discarding all results showing any non-zero error-sensitive event; this approach was the only one in which further decoding cannot be done. In the opposite regime, *without* post-selection, all results were kept and any ambiguous edges in the MWPM were treated without flipping the logical measurement; here, logical error rates could have been improved by decoding but was not strictly needed. Finally, the intermediate regime involved a *partial* post-selection scheme whereby results were only discarded if the MWPM algorithm highlighted an am-

biguous edge; here, decoding had to be done so that results with ambiguous edges that were highlighted could be discarded.

Logical error measurements and decoding.— Fitting the adjusted hyperedge probabilities to analytical expressions produces approximate estimates for the six-noise parameters in the error correcting experiments (Figure 2(b)). These noise estimates were found to be in good agreement with benchmarks based on simultaneous randomized benchmarking. Experiments were performed on four logical states ($| - / + \rangle_L$ and $| 0 / 1 \rangle_L$) each of which was stabilized up to 10 rounds to extract a logical error per round of stabilizers (Figure 3). This logical error varied depending on the analysis method.

For the *full* post-selection scheme, the logical error for some rounds fell below the best and average physical initialization and measurement errors - a hallmark of being below a so-called pseudo-threshold for fault-tolerant quantum computing. Fitting the decay curves resulted in inferred logical errors per round of $6.4 \pm 1.3 \times 10^{-4}$ for $| - / + \rangle_L$, and $11 \pm 1 \times 10^{-4}$ for $| 0 / 1 \rangle_L$.

If *none* of the instances of the experiment were discarded, then the logical error remained consistently above the pseudo-threshold. In this analysis without any post-selection and without decoding, we inferred logical errors per round of $40.4 \pm 0.2 \times 10^{-3}$ for $| - / + \rangle_L$, and $102 \pm 2 \times 10^{-3}$ for $| 0 / 1 \rangle_L$.

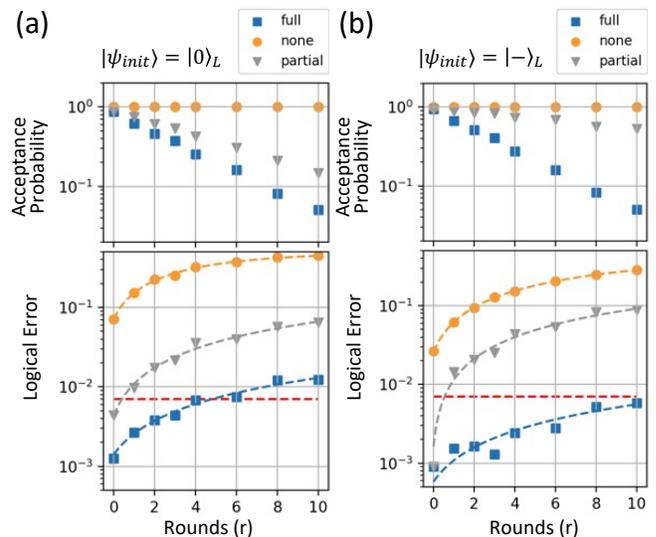


FIG. 3. (Color) (top) Fraction of total results used for the logical states ((a) $|0\rangle_L$, (b) $|- \rangle_L$) as the number of stabilizer rounds were repeated from 0 to 10 times when *full* (blue squares), *none* (yellow circles), or *partial* (gray triangles) post-selection analysis was used. For *partial* post-selection, the analytical decoder was used to exclude ambiguous shots. (bottom) The corresponding logical errors versus number of rounds. The dashed red lines indicate the pseudo-threshold as determined by the best (average) physical measurement errors of 7×10^{-3} (7.7×10^{-3}).

Recalling that the $[[4, 1, 2]]$ is an error detecting code, we used the syndrome outcomes from each stabilizer round to perform a *post-facto* logical correction in software. Discarding instances where ambiguous edges were highlighted by the decoder allowed us to apply the *partial*, in contrast with the *full*, post-selection scheme. With this scheme, significantly more instances of the

experimental runs remained, resulting in inferred logical errors per round of $10.7 \pm 0.7 \times 10^{-3}$ for $| - / + \rangle_L$, and $6.2 \pm 0.3 \times 10^{-3}$ for $| 0 / 1 \rangle_L$.

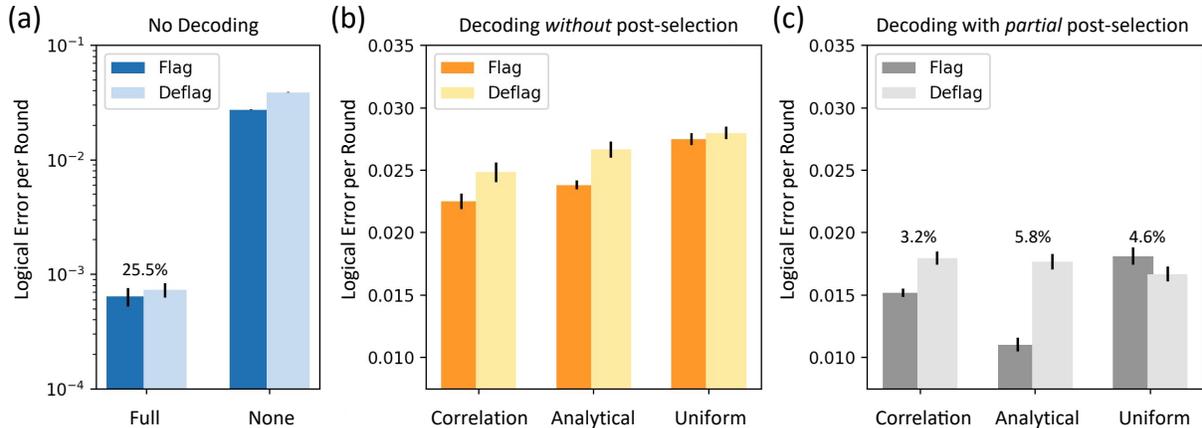


FIG. 4. (Color) Logical errors per round initially in $| - / + \rangle_L$ states under various analysis methods with acceptance probability per round labelled above. Results varied depending on whether the flag events were directly used for decoding (*flag*) or indirectly used for decoding using a *deflagging* procedure. (a) Logical error per round for *full* and *none* post-selection methods. 25.5% of the counts were rejected with each round for the *full* post-selection scheme. (b) Comparison between errors using three decoder graphs on data *without* post-selection. (c) Comparison between errors using three decoder graphs on data with *partial* post-selection. The approximate percentage of counts rejected for each stabilizer round are indicated above each bar.

Within the *none* and *partial* post-selection schemes, we were able to compare the performance of three different instances of decoders (Figure 4). The most generic decoder assumes there was no known noise model for the underlying physical system. Such a *uniform* decoder graph, in which every edge of the decoder graph was given equal weights, was expected to perform better than no decoding at all; but, was expected to be worse than any other graph whose edges were informed by some knowledge of the underlying noise. For instance, by selecting a simple, Pauli noise model, analytical expressions for the edge weights were calculated and led to improved logical error rates. Alternatively, if no assumptions were made about the noise, then edge weights were populated by the experimentally calibrated, correlation probabilities described earlier. We found that, as expected, such a correlation decoder graph indeed corrected for logical errors more effectively than the uniform decoding strategy and compared well with the analytical method (Figure 4(b)). However, when the *partial* post-selection scheme was used, this trend no longer held since an analytical decoder with noise parameters from simultaneous randomized benchmarking outperformed the correlation analysis (Figure 4(c)).

While the correlation analysis should, in principle, contain complete information about all of the noise in our experiments, its implementation is expected to become

more computationally costly when applied to codes at larger distances. We simplified the decoder graph and thus the number and size of hyperedges needed in the correlation analysis, by feeding-forward information from each round of flag measurements. This procedure, known as “deflagging” [37], allowed us to eliminate all 30 of the size-4 hyperedges in an experiment with 10 rounds of stabilizer measurements without a significant increase in the logical error per round (Figure 4). Furthermore, the logical errors were mostly preserved compared to results without the deflagging procedure.

Naïvely extending the HH code to distance-3 would result in size-5 hyperedges arising in the decoder hypergraph. However, when deflagging is applied, we found that there were no longer any size-5 hyperedges, and the number of size-4 hyperedges reduced from $148r - 12$ to $60r - 12$, where $r \geq 1$ is the number of rounds. Since the computational resources scale exponentially with the largest weight hyperedge in a graph, we expect that the deflagging procedure will provide a dramatic reduction in the computational resources needed to carry out the correlation analysis for codes beyond distance-3.

Conclusions.— Experimentally preparing and repeatedly stabilizing a logical quantum state, with error rates nearly ten times smaller than the lowest physical measurement error rate, is an important step towards executing larger, fault-tolerant circuits. The hexagonal lattice

on which we demonstrated our findings can be extended to operate larger distance versions of the fault-tolerant HH code used here, or for related codes [38–40] or even other codes using additional, bosonic degrees of freedom [41, 42]. Although the distance-2 version was implemented on a subset of qubits within a hexagonal lattice, other topologies are also expected to benefit; for example, a heavy-square topology akin to the rotated surface code with added flag qubits [32] or for distance-3 demonstrations of the widely studied surface code [43, 44]. Nevertheless, our probabilistic error correction methods and higher order error correlation analysis represents an approach for improving decoders for codes with or without flags within any device topology. We also demonstrated an effective use of flags to limit the extent of the correlations needed for efficient decoding. Our approach for extracting quantitative noise figures from the experiments creates a path to diagnose and reduce the logical errors per round of codes at larger distances.

As quantum computing devices become larger and less noisy, approaches such as ours may form the basis for efficiently decoding experimentally relevant errors. Other decoding strategies such as maximum-likelihood algorithms are known to scale unfavorably with code distances but may also benefit from our approach [45–47]. Eventually, decoders will need to be trained in real-time [48], whereby logical operations could be interleaved with calibration circuits to periodically update the de-

coder graph’s prior information with calibrated correlation probabilities. Previously studied bootstrapping techniques [26] coupled with the periodic re-calibration of the correlation edges may eventually approach near-optimal decoding efficiencies, although the existence of an optimal strategy remains an open question.

Acknowledgements.— We wish to acknowledge Ben W. Reichardt (USC) for proposing the deflagging procedure, Joe Latone for discussions on optimizing the numerical analysis, Ted Thorbeck for initial suggestions on read-out tune-up, David Lokken-Toyli and Oliver Dial for discussions on characterizing measurement impact on qubit states, and Isaac Lauer, Andrew Eddins, David McKay and Sarah Sheldon for valuable discussions. The comments from the anonymous reviewers also helped to improve this manuscript. These results were enabled by the work of the IBM Quantum software and hardware teams. We acknowledge support by IARPA under contract W911NF-16-1-0114 for the theoretical and experimental work (including partial device bring-up, characterization, and gate and measurement calibration) presented in this manuscript [49]. The device was designed and fabricated internally at IBM.

E.H.C. and T.J.Y. contributed equally to this work. M.L. and A.W.C. performed simulations and analysis. Y.K., N.S., S.S., A.D.C. and M.T. designed and conducted the experiments. All authors contributed to writing the manuscript.

-
- [1] É. Gouzien and N. Sangouard, Factoring 2048-bit RSA Integers in 177 Days with 13 436 Qubits and a Multimode Memory, *Physical Review Letters* **127**, 140503 (2021).
- [2] D. Aharonov and M. Ben-Or, Fault-Tolerant Quantum Computation with Constant Error Rate, *SIAM Journal on Computing* **38**, 1207 (2008).
- [3] E. Knill, R. Laflamme, and W. Zurek, Threshold Accuracy for Quantum Computation, arXiv:quant-ph/9610011 (1996), arXiv:quant-ph/9610011.
- [4] D. Gottesman, *Stabilizer Codes and Quantum Error Correction* (California Institute of Technology, 1997).
- [5] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Physical review A* **52**, R2493 (1995).
- [6] E. Knill and R. Laflamme, Theory of quantum error-correcting codes, *Physical Review A* **55**, 900 (1997).
- [7] P. Iyer and D. Poulin, A small quantum computer is needed to optimize fault-tolerant protocols, *Quantum Science and Technology* **3**, 030504 (2018).
- [8] O. Moussa, J. Baugh, C. A. Ryan, and R. Laflamme, Demonstration of Sufficient Control for Two Rounds of Quantum Error Correction in a Solid State Ensemble Quantum Information Processor, *Physical Review Letters* **107**, 160501 (2011).
- [9] J. Zhang, D. Gangloff, O. Moussa, and R. Laflamme, Experimental quantum error correction with high fidelity, *Physical Review A* **84**, 034303 (2011).
- [10] D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, Quantum computations on a topologically encoded qubit, *Science* **345**, 302 (2014).
- [11] N. M. Linke, M. Gutierrez, K. A. Landsman, C. Figgatt, S. Debnath, K. R. Brown, and C. Monroe, Fault-tolerant quantum error detection, *Science advances* **3**, e1701074 (2017).
- [12] 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* , 1 (2021).
- [13] 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 (2021), arXiv:2107.07505.
- [14] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. F. Du, P. Neumann, and J. Wrachtrup, Quantum error correction in a solid-state hybrid spin register, *Nature* **10.1038/nature12919** (2014).
- [15] M. H. Abobeih, Y. Wang, J. Randall, S. J. H. Loenen, C. E. Bradley, M. Markham, D. J. Twitchen, B. M. Terhal, and T. H. Tamminiau, Fault-tolerant operation of a logical qubit in a diamond quantum processor, arXiv:2108.01646 (2021), arXiv:2108.01646.
- [16] C. D. Hill, M. Usman, and L. C. L. Hollenberg, An exchange-based surface-code quantum computer architecture in silicon, arXiv:2107.11981 (2021),

- arXiv:2107.11981.
- [17] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, Computing with spin qubits at the surface code error threshold, arXiv:2107.00628 [cond-mat, physics:quant-ph] (2021), arXiv:2107.00628 [cond-mat, physics:quant-ph].
- [18] R. W. Andrews, C. Jones, M. D. Reed, A. M. Jones, S. D. Ha, M. P. Jura, J. Kerckhoff, M. Levendorf, S. Meenehan, S. T. Merkel, A. Smith, B. Sun, A. J. Weinstein, M. T. Rakher, T. D. Ladd, and M. G. Borselli, Quantifying error and leakage in an encoded Si/SiGe triple-dot qubit, *Nature Nanotechnology* **14**, 747 (2019).
- [19] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Realization of three-qubit quantum error correction with superconducting circuits, *Nature* **482**, 382 (2012).
- [20] 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, A. N. Cleland, and J. M. Martinis, State preservation by repetitive error detection in a superconducting quantum circuit, *Nature* **519**, 66 (2015).
- [21] A. D. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, Demonstration of a quantum error detection code using a square lattice of four superconducting qubits, *Nature Communications* **6**, 6979 (2015).
- [22] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo, Detecting bit-flip errors in a logical qubit using stabilizer measurements, *Nature Communications* **6**, 6983 (2015).
- [23] 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, *Physical Review Letters* **119**, 180501 (2017).
- [24] 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, *Nature Physics* **16**, 875 (2020).
- [25] J. F. Marques, B. M. Varbanov, M. S. Moreira, H. Ali, N. Muthusubramanian, C. Zachariadis, F. Battistel, M. Beekman, N. Haider, W. Vlothuizen, A. Bruno, B. M. Terhal, and L. DiCarlo, Logical-qubit operations in an error-detecting surface code, arXiv:2102.13071 (2021), arXiv:2102.13071.
- [26] Google Quantum AI, Exponential suppression of bit or phase errors with cyclic error correction, *Nature* **595**, 383 (2021).
- [27] M. Takita, A. D. Córcoles, E. Magesan, B. Abdo, M. Brink, A. Cross, J. M. Chow, and J. M. Gambetta, Demonstration of Weight-Four Parity Measurements in the Surface Code Architecture, *Physical Review Letters* **117**, 210505 (2016).
- [28] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, E. Magesan, J. A. Smolin, J.-B. Yau, V. P. Adiga, M. Sandberg, M. Brink, J. M. Chow, and J. S. Orcutt, Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors, *npj Quantum Information* **7**, 1 (2021).
- [29] IBM Quantum, High-fidelity superconducting quantum processors via laser-annealing of transmon qubits, arXiv:2012.08475 (2020), arXiv:2012.08475.
- [30] R. Chao and B. W. Reichardt, Quantum Error Correction with Only Two Extra Qubits, *Physical Review Letters* **121**, 050502 (2018).
- [31] C. Chamberland and M. E. Beverland, Flag fault-tolerant error correction with arbitrary distance codes, *Quantum* **2**, 53 (2018).
- [32] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and subsystem codes on low-degree graphs with flag qubits, *Physical Review X* **10**, 011022 (2020), arXiv:1907.09528.
- [33] See Supplemental Material at [URL will be inserted by publisher] for discussion of the Pauli tracer in Supp. D which does not include computational leakage, and results from simultaneous randomized benchmarking in Table S6.
- [34] O. Higgott, PyMatching: A Python package for decoding quantum codes with minimum-weight perfect matching, arXiv:2105.13082 (2021), arXiv:2105.13082.
- [35] See Supplementary Material at [URL will be inserted by publisher] for detailed derivations of the adjustment procedure in Supp. G.
- [36] See Supplementary Material at [URL will be inserted by publisher] for decoder graphs with classified edges labelled in color in Figure S1.
- [37] See Supplemental Material at [URL will be inserted by publisher] for a detailed explanation of the deflagging procedure in Supp. F.
- [38] M. B. Hastings and J. Haah, Dynamically Generated Logical Qubits, arXiv:2107.02194 (2021), arXiv:2107.02194.
- [39] C. Gidney, M. Newman, A. Fowler, and M. Broughton, A Fault-Tolerant Honeycomb Memory, arXiv:2108.10457 (2021), arXiv:2108.10457.
- [40] J. R. Wootton, Hexagonal matching codes with 2-body measurements, arXiv:2109.13308 (2021), arXiv:2109.13308.
- [41] S. Puri, A. Grimm, P. Campagne-Ibarcq, A. Eickbusch, K. Noh, G. Roberts, L. Jiang, M. Mirrahimi, M. H. Devoret, and S. M. Girvin, Stabilized cat in a driven nonlinear cavity: a fault-tolerant error syndrome detector, *Physical Review X* **9**, 041009 (2019).
- [42] B. M. Terhal, J. Conrad, and C. Vuillot, Towards scalable bosonic quantum error correction, arXiv: Quantum Physics (2020).
- [43] S. Krinner, N. Lacroix, A. Remm, A. D. 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 (2021), arXiv:2112.03708 [quant-ph].
- [44] Y. Zhao, Y. Ye, H.-L. Huang, Y. Zhang, D. Wu, H. Guan, Q. Zhu, Z. Wei, T. He, S. Cao, F. Chen, T.-H. Chung, H. Deng, D. Fan, M. Gong, C. Guo, S. Guo, L. Han, N. Li, S. Li, Y. Li, F. Liang, J. Lin, H. Qian, H. Rong, H. Su, L. Sun, S. Wang, Y. Wu, Y. Xu, C. Ying, J. Yu, C. Zha, K. Zhang, Y.-H. Huo, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, Realizing an error-correcting surface code with superconducting qubits (2021), arXiv:2112.13505 [quant-ph].
- [45] S. Bravyi, M. Suchara, and A. Vargo, Efficient Algorithms for Maximum Likelihood Decoding in the Surface Code, *Physical Review A* **90**, 032326 (2014).
- [46] S. Bravyi and A. Cross, Doubled Color Codes, arXiv:1509.03239 (2015), arXiv:1509.03239.

- [47] B. Heim, K. M. Svore, and M. B. Hastings, Optimal Circuit-Level Decoding for Surface Codes, arXiv:1609.06373 (2016), [arXiv:1609.06373](#).
- [48] P. Das, A. Locharla, and C. Jones, LILLIPUT: A Lightweight Low-Latency Lookup-Table Based Decoder for Near-term Quantum Error Correction, arXiv:2108.06569 (2021), [arXiv:2108.06569](#).
- [49] See Supplemental Material at [URL will be inserted by publisher] for further discussion of references [23-26, 28, 32, 50-57].
- [50] S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, and J. M. Gambetta, Mitigating measurement errors in multiqubit experiments, *Physical Review A* **103**, 042605 (2021).
- [51] IBM Quantum and Community, [Qiskit: An open-source framework for quantum computing](#) (2021).
- [52] S. Garion, N. Kanazawa, H. Landa, D. C. McKay, S. Sheldon, A. W. Cross, and C. J. Wood, Experimental implementation of non-clifford interleaved randomized benchmarking with a controlled-s gate, *Physical Review Research* **3**, 013204 (2021).
- [53] C. Gidney, [Decorrelated depolarization](#) (2020).
- [54] R. Chao, M. E. Beverland, N. Delfosse, and J. Haah, Optimization of the surface code design for Majorana-based qubits, *Quantum* **4**, 352 (2020).
- [55] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, *Physical Review A* **69**, 062320 (2004).
- [56] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Mineev, J. M. Chow, and J. M. Gambetta, Exploiting Dynamic Quantum Circuits in a Quantum Algorithm with Superconducting Qubits, *Physical Review Letters* **127**, 100501 (2021).
- [57] N. Sundaresan, I. Lauer, E. Pritchett, E. Magesan, P. Jurcevic, and J. M. Gambetta, Reducing Unitary and Spectator Errors in Cross Resonance with Optimized Rotary Echoes, *PRX Quantum* **1**, 020318 (2020).