





Resilience of the surface code to error burstsShi Jie Samuel Tan ^{1,2}, Christopher A. Pattison ³, Matt McEwen ⁴, and John Preskill ^{3,5}¹*Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, Maryland 20742, USA*²*Haverford College, Haverford, Pennsylvania 19041, USA*³*Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, California 91125, USA*⁴*Google Quantum AI, Santa Barbara, California 93111, USA*⁵*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

(Received 16 October 2025; accepted 5 January 2026; published 26 February 2026)

Quantum error correction works effectively only if the error rate of gate operations is sufficiently low. However, some rare physical mechanisms can cause a temporary increase in the error rate that affects many qubits; examples include ionizing radiation in superconducting hardware and large deviations in the global control of atomic systems. We refer to such rare transient spikes in the gate error rate as error bursts. In this work, we investigate the resilience of the rotated surface code to generic error bursts. We assume that, after appropriate mitigation strategies, the spike in the error rate lasts for only a single syndrome-extraction cycle; we also assume that the enhanced error rate is uniform across the code block. Under these assumptions, and for a circuit-level depolarizing noise model, we perform Monte Carlo simulations to determine the regime in burst error rate and background error rate for which the memory time becomes arbitrarily long as the code block size grows. Our results indicate that suitable hardware mitigation methods combined with standard decoding methods may suffice to protect against transient error bursts in the rotated surface code.

DOI: [10.1103/vwtl-hwfx](https://doi.org/10.1103/vwtl-hwfx)**I. INTRODUCTION**

Because quantum hardware is intrinsically prone to error, a large-scale quantum computer will need to be robust against noise. This can be accomplished by executing an encoded version of a quantum circuit, where every qubit is replaced by a logical qubit protected by a quantum error-correcting code (QECC).

If the error rate does not exceed a critical value, the *accuracy threshold*, the output of an ideal quantum circuit can be sampled with arbitrary accuracy by a noisy fault-tolerant circuit if the code block is sufficiently large. The most studied candidate QECC for achieving fault tolerance, the surface code [1,2], has a high threshold and may be implemented using nearest-neighbor connectivity on a two-dimensional lattice. These features make the surface code well suited for implementation in a solid-state platform such as a superconducting circuit with transmon qubits [3,4].

Recent experimental studies of superconducting quantum processors [5–7] suggest that ionizing radiation may pose serious limitations on the size of quantum circuits that can be executed fault tolerantly using these devices. Such ionizing radiation, for example a cosmic ray muon, can create many quasiparticles, significantly reducing the qubit T_1 time, and hence greatly enhancing the error rate in a large spatial region of the device. Because many qubits in a single code block may be affected, this event can cause an uncorrectable logical

error. To avoid this catastrophe, one might place the quantum computer deep underground to suppress the muon flux, or use a hierarchical error-correction scheme which protects against events that damage many qubits in a subblock of a larger code block [8,9]. Here we consider a third option: We envision that hardware mitigation methods such as quasiparticle traps [10] limit the temporal extent of the ionization event, and investigate whether a transient spike in the error rate can be tolerated by the surface code.

Such error bursts may occur in other platforms aside from superconducting circuits, and for other reasons than ionizing radiation. For example, brief episodes of elevated gate error rates might arise from rare large deviations from average behavior in global control systems. Here too suitable mitigation strategies can alleviate such global control noise. Our results provide guidance concerning how much mitigation is required to execute long fault-tolerant quantum computations with high success probability.

Sustainable threshold and burst threshold

To achieve fault tolerance using surface codes, logical operations are interleaved with so-called *recovery operations*. In each recovery operation, the stabilizer generators [11] that define the code subspace are projectively measured using a *syndrome extraction circuit*. Although syndrome extraction is noisy and therefore might introduce additional errors, error correction succeeds with high probability if the error rate is below threshold and the code block is sufficiently large. We refer to a single execution of a syndrome extraction circuit producing one measurement outcome for each stabilizer generator as a *round* of syndrome extraction. For the surface code, in the presence of measurement errors, it is necessary

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

to repeat the syndrome extraction circuit a number of times roughly equal to the code distance d [2]. We refer to this $O(d)$ -times repetition of the syndrome extraction circuit as one *logical cycle*.

In this work, we quantitatively study the resilience of the rotated surface code decoded with the minimum-weight perfect-matching (MWPM) decoder under the assumption that during ionizing radiation events the device is much noisier than the background error rate for a brief period on the order of the gate time. For the sake of brevity, we refer to rotated surface codes as surface codes. We refer to the ionizing radiation events as *error bursts* and assume that (1) each error burst occurs within the duration of a single round of syndrome extraction and (2) error bursts are rare. Because error bursts are rare, we may regard them as isolated events; that is, typically many rounds of syndrome measurement occur in-between successive error bursts.

We contrast the threshold error rate during error bursts with the threshold for the background error rate. We refer to these two values as the *burst threshold* and the *sustainable threshold*, respectively. In the absence of measurement and gate errors (the code capacity model), all errors are corrected within a single syndrome-measurement round and, therefore, in that case the sustainable threshold and burst threshold are equal. However, when there are measurement errors, the burst threshold can be higher than the sustainable threshold; while an excessive number of errors occurring during a burst might be intolerable if followed immediately by further rounds with a similarly large error rate, these errors might become correctable if the error rate following the burst quickly settles down to its background value. In other words, if the background error rate is sufficiently low, we might expect the rotated surface code to be resilient to error bursts with high error rates up to the burst threshold. In this work, we investigate the separation between the burst threshold and the sustainable threshold in a simple noise model. Our main result is a relation between the values of these thresholds, depicted in Fig. 4.

II. BACKGROUND

We will analyze a noise model which, although highly idealized, provides a useful caricature of actual errors in realistic devices. Specifically, we consider a depolarizing-channel noise model in which the error rate is enhanced during a burst that lasts for a single syndrome-extraction cycle. Here we briefly sketch two possible sources of such error bursts, namely, high-energy impact events in superconducting hardware, and global control noise in atomic, molecular, and optical (AMO) and other hardware platforms. See Appendix A for further details.

Ionizing radiation impacting solid-state superconducting devices has been shown to produce highly correlated error bursts [7, 12–14]. The spatial region affected by such an impact event can be larger than the projected size of surface-code logical qubits; hence, a spatially uniform error burst is a reasonable representation of the errors induced by these events. While the error bursts found in current devices typically take hundreds to thousands of error correction cycles to abate, mitigation strategies have been proposed that would substantially reduce this timescale [10, 15]. Therefore, studies of our

idealized error burst model can illuminate the usefulness of these mitigation strategies.

Global control is a desirable architectural feature, where a single control system influences many qubits simultaneously. For instance, in AMO systems a single laser might facilitate many gates in parallel, providing very favorable scaling of the control hardware as the number of qubits grows. However, noise in a global control system presents a potential single point of failure, as an aberration in the control signal output can produce errors on all qubits the signal is applied to. If substantial deviations in the global control signal are rare and independent over time, they can be accurately described as error bursts. Here, too, analysis of our error burst model is helpful for assessing architectural tradeoffs.

A. Noise model

If an error burst elevates the gate error rate over many consecutive syndrome extraction cycles, then a logical error is likely to occur unless the error rate during the burst is below the surface code sustainable threshold. However, if error bursts are brief and rare, then additional resilience may be possible; therefore, we study the case where the error burst has a limited temporal extent. The rapid return to the background error rate after the burst might occur because of effective quasiparticle mitigation strategies such as gap engineering or because classical control noise has appreciable high-frequency components. Recent experimental results in a gap-engineered transmon qubit array [16] have noise characteristics that are qualitatively similar to the model that we consider here. After mitigations, Ref. [16] observed large error bursts lasting several syndrome extraction cycles in duration. Future mitigations may result in noise characteristics approaching the model we consider here.

In a realistic setting one might expect errors in successive rounds to be correlated; for simplicity we assume there are no such correlations. Another feature of error bursts is their large spatial extent; we pessimistically assume that the elevated error rate during the burst applies uniformly across the entire surface code block. It would be useful to extend our analysis to the case where errors in distinct rounds are correlated and error bursts have limited spatial extent, but we have not attempted that more general analysis. A recent work by Chadwick *et al.* [17] studies optimization of magic state distillation under a related model of ionizing radiation where the region of elevated noise has limited spatial extent but very long temporal extent.

The one-qubit depolarizing noise channel with depolarizing rate p applies the identity with probability $1 - p$ and one of X , Y , or Z each with probability $p/3$. Likewise, the two-qubit depolarizing noise channel with depolarizing rate p applies identity with probability $1 - p$ and one of the 15 non-trivial two-qubit Pauli operators each with probability $p/15$.

We consider the surface code under two noise models, phenomenological noise and circuit-level depolarizing noise. In both cases, we assign a noise parameter p to each round of syndrome extraction.

In the phenomenological noise model, bit-flip noise (Pauli X applied with probability p) is applied to all data qubits with

rate p followed by a round of perfect syndrome extraction. Finally, all measurement outcomes are flipped independently with probability p .

In the circuit-level depolarizing noise model, a standard syndrome extraction circuit is executed [18] where for each one- (two-) qubit gate in the circuit, the gate is assumed to execute perfectly followed by one- (two-) qubit depolarizing noise with rate p . In addition, every data qubit that is idle in any particular time step undergoes depolarizing noise with rate p . All measurement outcomes are flipped with probability p . In addition, to simulate state-preparation error, we apply a bit-flip error with probability p to each data qubit.

During a round of syndrome extraction where an error burst occurs, we replace the noise parameter with the elevated error rate p_B . In particular, for the circuit-level depolarizing noise model, this implies that all gate operations fail with probability p_B within the syndrome extraction round where the error burst occurs. We refer to the difference between the background error rate p and the elevated error rate p_B as the burst error rate p_B' where $p_B' = p_B - p \geq 0$.

The performance of a stabilizer code under such a noise model can be efficiently simulated using standard open-source tools such as stim [19].

B. Decoding surface codes

Surface codes can be efficiently decoded both in theory and in practice [2,20,21]. The efficient decoding of surface codes usually relies on a construction known as the *decoding graph* containing a vertex for every measured stabilizer generator. For every independent error e , the decoding graph contains an edge incident to the vertices corresponding to the flipped measurement outcomes due to the presence of e . If the error occurs with probability p_e , the edge is assigned a weight $\ln \frac{1-p_e}{p_e}$. To ensure that the decoding graph is a well-defined graph, Pauli errors are decomposed into a product of Pauli X and Pauli Z operators which are assumed to be independent.

In all numerics, we assume knowledge of the time at which the error burst occurs and update the edge weights appropriately. This decoding graph for a distance-4 surface code under the phenomenological noise model is illustrated in Fig. 1. However, we find that it is easy to infer the presence of an error burst event by considering the Hamming weight of the change in syndrome in each round: An error burst with elevated error rate much higher than background error rate can be easily distinguished from the background error rate. Appendix B shows an example maximum-likelihood estimate of identifying whether the noise in a syndrome extraction round is drawn from a distribution with one of two error rates. In a real system, the distribution of the global noise parameter will have support on many values. Estimation of the noise is still possible, but it may be the case that a Bayesian framework with carefully selected prior would be more suitable.

In this work, we use an open source implementation of the minimum-weight perfect-matching (MWPM) decoder [22] on the weighted decoding graph derived from the noise model [19]. While the error-correction performance of surface codes in the presence of error bursts depends on our choice of the decoder, we restrict our study to the MWPM decoder for

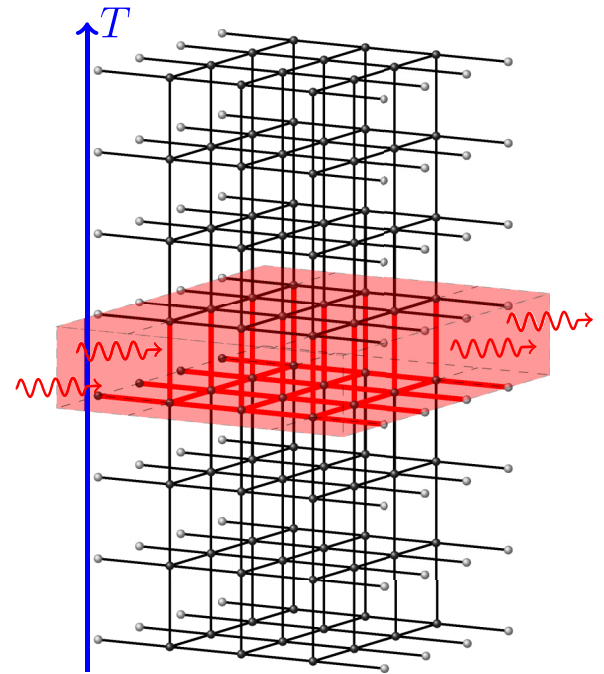


FIG. 1. Decoding graph for a distance-4 surface code. Black edges correspond to independent error events. The vertices correspond to the syndrome bits at each spatial position and time. A syndrome bit v is 1 if the error set contains an odd number of edges incident to v . The black edges parallel to the x - y plane correspond to errors on the data qubits of the surface code. The black edges parallel to the z axis correspond to syndrome measurement errors and are in the error set if the associated measurement is flipped relative to its ideal value. The decoding graph also corresponds to the noise model of the code: To every edge, we can associate an independent random indicator variable that is 1 if the edge is contained in the error set and 0 otherwise. The surface code undergoes T consecutive rounds of syndrome extraction with an error burst occurring at round $T/2$. The edges in that cycle are marked in red, indicating that the indicator variables are 1 with enhanced probability during that cycle.

simplicity and leave the comparison with other decoders for future work.

C. Threshold experiment

We consider the rotated surface code [23] subjected to a memory experiment of duration $T = 2d$ syndrome extraction cycles where a single error burst occurs at time $T/2$ and d is the surface code distance. We choose T to be $2d$ to make sure that the error burst is d time steps away from the boundaries to prevent boundary effects from affecting the performance of our fitting *Ansatz*. Using Monte Carlo sampling, we estimate the threshold of the surface code under phenomenological and circuit-level depolarizing noise with rate p including an error burst at time $T/2$ with rate p_B . The memory experiment consists of sampling the noise via Monte Carlo sampling with a final round of perfect readout in the Z basis. A failure is declared if the eigenvalue of the logical Z operator has been flipped by the error and correction.

Below the threshold, increasing the code distance exponentially suppresses logical errors whereas above the threshold

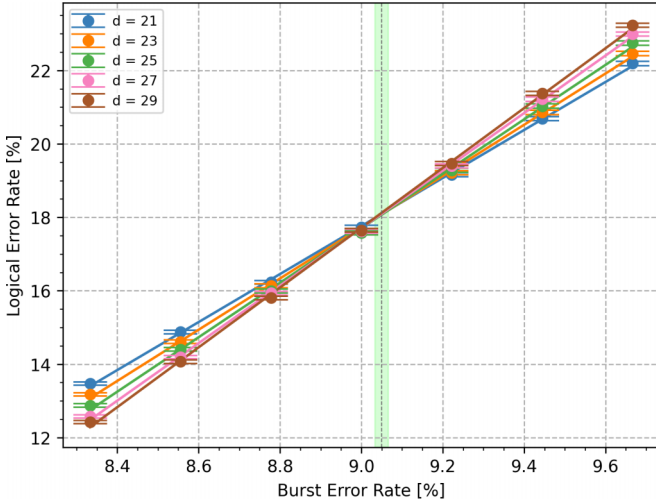


FIG. 2. Sweep of p_B indicating a threshold for phenomenological noise with background error rate $p = 2.0\%$. The error burst threshold $p_B^* = 7.050(16)\%$ is indicated by the gray vertical line with green error bars.

increasing code distance increases the failure rate of the memory. For each code distance, we run a sweep of the memory experiment with different values of the background or burst error rate to observe the characteristic crossing of the logical failure rates corresponding to the threshold (Fig. 2). To reduce finite-size effects, in the different limits, low measurement failure rate and low burst error rate, we sweep the burst error rate and background error rates, respectively. The threshold estimation method is described in more detail in Appendix C.

A single error burst is modeled in order to have a well-defined threshold. However, the conclusions should hold as long as the surface code distance is much smaller than the average time between error bursts: In other words, the “correlation length” of the decoding problem under independent noise is on the order of the surface code distance [2]. In typical devices, the average time between error bursts is more than six orders of magnitude greater than the gate time [5], so we are in this regime for all practical purposes. We observe similar results in the setting of depolarizing circuit-level noise.

Our main results are shown in Figs. 3 and 4. We find that the memory is below threshold even for relatively large values of both the background error rate and burst error rate. This is the regime of practical interest for near-term devices. We note that achieving good resilience to both noise sources simultaneously is nontrivial since the presence of one noise source weakens resilience to the other noise source. Our results for the circuit level depolarizing noise model suggest that, based on estimates in [10], hardware mitigation and standard fault tolerance techniques alone may suffice to protect against ionizing radiation and global control noise.

III. RESULTS

A. Thresholds

We begin by studying the threshold behavior of surface codes in a memory experiment of duration $T = 2d$

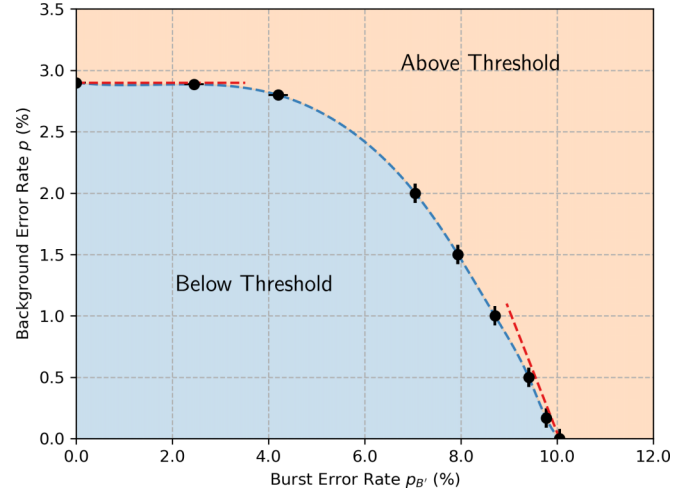


FIG. 3. Phase diagram of the surface code under phenomenological noise. The x axis and y axis indicate burst and background error rates, respectively. The light tan area indicates a noise rate that is above threshold while the dark blue area indicates a noise rate below threshold. Data points are connected by a blue dashed line as a guide to the eye. Previously known limiting behavior in the $p \rightarrow 0$ and $p_B \rightarrow 0$ limits is drawn as a red dashed line, indicating upper bounds on the threshold error rates.

where a single error burst occurs at time $T/2$ and d is the surface code distance. We verified that similar threshold estimates are obtained when the duration is longer than $T = 2d$.

The statistical mechanics mapping for such an error model [2,24] corresponds to a three-dimensional (3D) random-bond Ising model with increased disorder along a two-dimensional

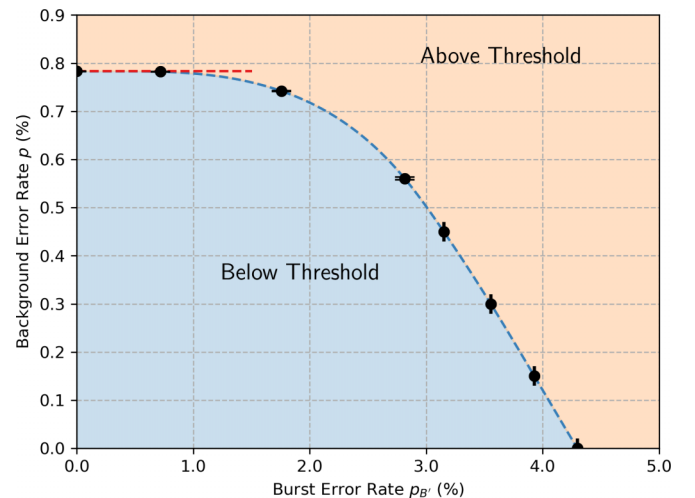


FIG. 4. Phase diagram of the surface code under circuit-level depolarizing noise. The x axis and y axis indicate burst and background error rates, respectively. The light tan area indicates a noise rate that is above threshold while the dark blue area indicates a noise rate below threshold. Data points are connected by a blue dashed line as a guide to the eye. Previously known limiting behavior in the $p_B \rightarrow 0$ limit is drawn as a red dashed line indicating an upper bound on the threshold error rate.

(2D) surface immersed in the bulk. Related models have been considered previously in the context of noisy lattice surgery [25], where the elevated error rate occurs at the interface between surface code blocks, and quantum communication [26], where the elevated error rate characterizes noisy Bell pairs shared by nodes in a quantum network. We present the phase diagram under phenomenological noise in Fig. 3. Notably, we find that the surface code is stable to even large error bursts at finite background error rates which *a priori* is not a given. That is, the presence of additional errors accumulating at a small rate does not significantly harm the ability to recover from a large amount of preexisting errors. Quantitatively, under phenomenological noise, the presence of an error burst with burst error rate $p_B = 7\%$ only depresses the sustainable threshold from $\approx 3\%$ to $\approx 2.7\%$.

A question of practical relevance is the setting where error bursts occur at a fixed rate Γ . After each error burst, the residual population of qubit errors settles back down to the background value after about d syndrome measurement cycles. Thus, for $d \ll 1/\Gamma$, the *rare event regime*, the actual resilience to error bursts is captured well by our idealized setting in which only a single error burst is considered. But in the opposite regime where $d \gg 1/\Gamma$, a second error burst is likely to occur before the effects of the previous error burst have fully relaxed. Indeed we expect that in the thermodynamic limit $d \rightarrow \infty$, $\Gamma = \text{const}$, there may not be a stable phase with $p_B > p$. However, d is only required to increase polylogarithmically with the size of the computation, so a polynomial increase in $1/\Gamma$ leads to an exponential increase in the size of circuit that can be executed while remaining in the rare event regime.

By performing a similar study for circuit-level depolarizing noise, we obtain the phase diagram in Fig. 4; here too we find that the threshold background error rate is relatively stable when the burst error rate is not too high. Notably, though, the threshold burst error rate has a stronger dependence on the background error rate than for the phenomenological noise case. A reasonable hypothesis is that this distinct behavior exhibited by the two phase diagrams occurs because in the circuit-based noise model a larger portion of the errors introduced during the burst remain extant in the following rounds. Thus, the background error rate must be reduced accordingly to prevent the error population from surpassing what the decoder can handle.

To validate this explanation, we evaluate a proxy for the density of errors by computing the marginal expected value of syndrome measurements averaged over a spatial slice, as shown in Fig. 5. The effect of the burst on the error population is clearly confined to a single time slice in the phenomenological noise model, but the error density remains elevated for two consecutive time slices in the circuit-level depolarizing noise model. This effect arises because some of the errors produced in the burst are not detected until the following syndrome extraction cycle. These surviving errors from the burst are augmented by additional errors which occur at the background error rate in the following cycle, explaining the stronger combined effects of burst errors and background errors revealed by the phase diagram.




FIG. 5. Marginal expected value of syndrome bits averaged over each time slice in a surface-code memory experiment with an error burst. The marginal syndrome bit value is a proxy for the density of errors in a syndrome extraction cycle. Here the code distance is $d = 3$, and 20 rounds of syndrome extraction are shown, with an error burst in the 11th round. The background error rate is $p = 0.1\%$ and the elevated error rate is $p_B = 1\%$. The effect of the error burst is confined to a single syndrome-extraction cycle for the phenomenological noise model; however, it extends into the following syndrome extraction cycle for circuit-level noise because some errors occurring during the burst are not detected until the following round.

B. Teraquop footprint

In order to quantitatively estimate the effect of error bursts on the amount of resources required to execute large circuits under circuit-level depolarizing noise, we estimate the teraquop footprint [27]. Roughly, the teraquop footprint is the space overhead required to execute a circuit of size 10^{12} such that the probability of success is roughly $1/e$. This is achieved by a logical failure rate of 10^{-12} .

1. Logical error rate

To separately extract the background logical failure rate and logical failure rate due to error bursts, we run two memory experiments *A* and *B*. Experiment *A* is a standard D logical-cycle memory experiment using a distance- d surface code with background error rate p and background logical error rate $q_d(p)$. In the regime $D \gg 1$ and assuming the failure for each logical cycle is independent and identically distributed distributed [28], the failure probability $\tilde{p}_{d,A}(p, D)$ of the memory experiment is related to the failure probability per logical cycle by the following fitting *Ansatz* in Ref. [29]:

$$2\tilde{p}_{d,A}(p, D) \approx 1 - [1 - q_d(p)]^D. \quad (1)$$

The error rate approaches $\frac{1}{2}$ as $D \rightarrow \infty$ because a correction selected uniformly at random has probability $\frac{1}{2}$ of being in the same equivalence class as the error.

In memory experiment *B*, we use the same noise model except we inject an error burst with error rate p_B in the $(\frac{D}{2})^{\text{th}}$ logical cycle. If the logical cycle containing the error burst fails with probability $q_{d,B}(p_B, p)$, then we can extend the *Ansatz* in Eq. (1) to express the failure probability

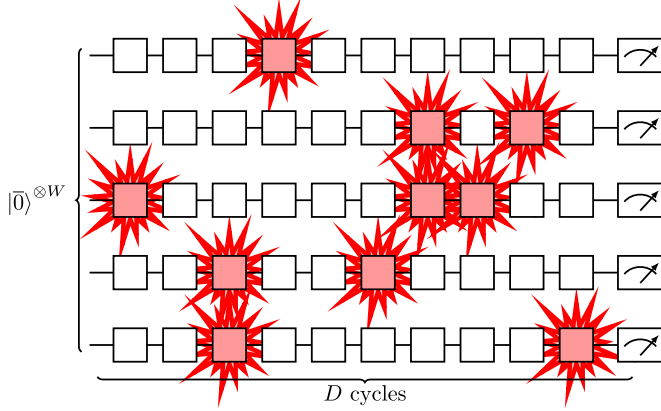


FIG. 6. Logical quantum circuit for memory experiment B with W logical qubits and D logical cycles such that $WD = 10^{12}$. Uncolored blocks are logical cycles that only suffer from the background error rate and fail with probability $\tilde{q}_d(p)$. Red blocks are logical cycles in which an error burst occurs; these fail with probability $\tilde{q}_{d,B}(p_B, p)$.

$\tilde{p}_{d,B}(p_B, p, D)$ in memory experiment B as

$$2\tilde{p}_{d,B}(p, D) \approx 1 - [1 - q_{d,B}(p_B, p)][1 - q_d(p)]^{D-1}. \quad (2)$$

In other words, we attribute all excess logical failures in experiment B to the presence of the error burst, allowing separate determinations of the background and burst logical error rates. Using Eqs. (1) and (2), we estimate $q_d(p)$ and $q_{d,B}(p_B, p)$, finding that these quantities approach constant values independent of D , denoted $\tilde{q}_d(p)$ and $\tilde{q}_{d,B}(p_B, p)$, when D is large. A more extended discussion of finite-size effects from finite duration memory experiments can be found in [29], Appendix D. We extrapolate $\tilde{q}_d(p)$ and $\tilde{q}_{d,B}(p_B, p)$ to large d , and correspondingly small logical failure rates, in order to estimate the teraquop footprint. A memory experiment with a moderate error burst rate is illustrated in in Fig. 6. The low error rate extrapolation is explained in greater detail in Appendix D.

2. Failure model

If the average number of logical cycles between error bursts is $\tau \geq 1$, the probability $P_L(p_B, p)$ of a logical failure per logical cycle qubit is approximately

$$P_L(p_B, p) \approx \frac{1}{\tau} \tilde{q}_{d,B}(p_B, p) + \left(1 - \frac{1}{\tau}\right) \tilde{q}_d(p). \quad (3)$$

TABLE I. Fit parameters for circuit-level depolarizing noise with elevated error rate $p_B = 2.5\%$ and five different values of the background error rate p . Here $\log_{10} \tilde{q}_d = c[\tilde{q}_d] + dm[\tilde{q}_d]$, and $\log_{10} \tilde{q}_{d,B} = c[\tilde{q}_{d,B}] + dm[\tilde{q}_{d,B}]$. All confidence intervals are two standard deviations.

Background error rate (%)	$c[\tilde{q}_d]$	$m[\tilde{q}_d]$	$c[\tilde{q}_{d,B}]$	$m[\tilde{q}_{d,B}]$
0.08	-1.77(3)	-0.421(7)	-0.902(5)	-0.0459(7)
0.12	-1.655(13)	-0.343(4)	-0.872(5)	-0.0459(6)
0.18	-1.531(7)	-0.267(2)	-0.847(5)	-0.0420(6)
0.27	-1.419(4)	-0.1856(7)	-0.827(6)	-0.0342(6)
0.40	-1.311(2)	-0.1069(3)	-0.833(8)	-0.0188(7)




FIG. 7. Teraquop footprint for background error rate $p = 0.1\%$ as a function of the burst error rate p_B for various values of the mean time τ between error bursts. The burst error threshold in the limit $p \rightarrow 0$ is indicated as a vertical dashed line.

We assume a surface code of distance d requires a total of $2d^2$ qubits including ancilla qubits [30]. We define the teraquop footprint as the total number of physical qubits (including ancilla qubits) per logical qubit of the minimum distance surface code that achieves a logical error rate of 10^{-12} .

If a logical cycle takes about $10 \mu\text{s}$, then an error burst occurring once per minute corresponds roughly to $\tau \approx 10^7$. In Fig. 7, we show the teraquop footprint as a function of the mean time between error bursts τ and the error burst rate when the background error rate is $p = 10^{-3}$. We provide fit parameters used for the teraquop footprint estimate in Table I in the Appendix.

Notably, an error burst rate an order of magnitude above the background error rate does not appreciably affect the teraquop footprint. However, performance quickly degrades beyond this point. At low background error rates, the teraquop footprint is also relatively insensitive to further reduction in the background error rate.

IV. DISCUSSION AND CONCLUSION

We have shown that rotated surface codes are resilient to rare “error bursts” in which the physical error rate is briefly enhanced. Such error bursts have many potential sources and become relevant once the frequency of burst

events becomes comparable to the inverse of the duration of the longest computation wall time we wish to consider (e.g., days to weeks). Because such events are rare they may escape notice in experiments with insufficient runtime. Superconducting qubit devices stand apart in that error bursts originating from ionizing radiation have been directly detected. But in a variety of quantum platforms engineering constraints favor control signals shared by many qubits. This engineering convenience carries the risk of potential error burst failure modes which should be understood and mitigated.

With appropriate mitigation methods, the burst event can be brief; furthermore, the burst error rate, though higher than the background error rate, may not be so high as to render quantum error correction ineffective. We find that if the mean time between error bursts is much larger than the duration of a full logical cycle, and if the burst error rate is below the burst threshold, then the overhead cost of error correction is rather insensitive to the burst error rate. For reasonable parameters (10 μ s logical cycle, 1 min^{-1} mean time between bursts, background error rate $p = 0.1\%$), we find that if the elevated error rate is 15 times larger than the background error rate, then the teraquop footprint increases by less than a factor of 2 compared to the case without any burst errors.

In our computations, we assumed that the decoder knows the time step in which a burst error occurs. This assumption is reasonable because the elevated error rate during the burst is easily detected in the syndrome measurement history. Furthermore, we expect the decoder to perform nearly as well even when prior information about the time of the burst is not provided.

Our studies have been limited to the surface code, but we anticipate that qualitatively similar results apply when using other families of quantum low-density parity-check codes. Our conclusions bolster the hope that, given appropriate mitigation strategies that diminish the temporal extent and strength of error bursts, deep quantum circuits with valuable practical applications can be executed reliably on fault-tolerant quantum computers.

ACKNOWLEDGMENTS

We thank S. Flammia, H. Zhou, D. Bluvstein, P. S. Rodriguez, and A. Shaw for valuable conversations. S.J.S.T. acknowledges funding from Caltech's Summer Undergraduate Research Fellowship (SURF) and QuICS's Lanzaos Graduate Fellowship. C.A.P. acknowledges funding from the Air Force Office of Scientific Research (AFOSR) Grant No. FA9550-19-1-0360 and U. S. Department of Energy Office of Science, Grant No. DE-SC0020290. J.P. acknowledges support from the U. S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (Grants No. DE-NA0003525 and No. DE-SC0020290), the U. S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator, and the National Science Foundation (Grant No. PHY-1733907). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

APPENDIX A: SOURCES OF ERROR BURSTS

Here, we discuss in more detail two motivating cases of error sources in specific hardware implementations that are modeled well as error bursts, namely, quasiparticle bursts in superconducting qubit platforms, and global control noise in AMO and other platforms.

1. Quasiparticle bursts from high-energy impact events

When ionizing radiation (for example, a cosmic ray muon or gamma radiation) interacts with the chip substrate, large amounts of energy are deposited (100 keV–1 MeV), which quickly radiates outwards in the form of high-energy phonons [7,12]. When these phonons impinge on the superconducting qubit material, they break Cooper pairs and produce Bogoliubov quasiparticles. The quasiparticles can diffuse through the qubit and tunnel through the Josephson junction, where they can absorb the qubit energy. This tunneling process severely limits the T_1 relaxation time of the qubits, leading to a substantial increase in the error rate. Due to the rapid ballistic transport of high-energy phonons through the substrate, these elevated error rates are experienced over a large area, comparable to the size of current superconducting qubit arrays.

The error rate returns to normal when the quasiparticle population density relaxes to its background value. Quasiparticles disappear as they recombine to form Cooper pairs, releasing phonons. These phonons escape slowly off the device, producing long recovery timescales between 100's of microseconds [12] and 10's of milliseconds [7,13,14]. Recombination requires two quasiparticles to interact, and such interactions become less frequent as the population thins out; therefore, the recombination rate declines as the density is reduced, and as a result the quasiparticle density decreases only polynomially with time [31].

Various mitigation strategies have been proposed [10] and implemented. Phonon trapping [15] aims to down-convert phonons to a frequency lower than the qubit superconducting gap, preventing quasiparticle generation in the qubits. Quasiparticle trapping [32] aims to concentrate the quasiparticle populations away from the qubit Josephson junction, limiting the errors induced by the elevated quasiparticle population and aiding recombination. These mitigation strategies can enhance the speed of recovery to normal performance to the extent that only a single round of error correction is affected [32]. While the spatial extent of the error burst is not infinite, and may be significantly improved by these mitigation strategies, a uniform error burst over all qubits represents a worst case for QEC performance.

2. Global control noise

Implementing gates in quantum computation experiments requires precise time variation in Hamiltonian parameters. To ease engineering challenges, it is frequently the case that a single device is utilized in some step of the control electronics for the control of many qubits (e.g., shared microwave sources, arbitrary waveform generators, lasers). Such sharing of resources makes the system *formally* not fault tolerant, but it may be reliable enough for practical purposes. Since the




FIG. 8. (Top) Illustration of hypothetical global control parameter g with periodic pulses (e.g., envelope of laser amplitude modulation). The ideal value is indicated in blue while orange indicates a trace with white noise added [independent in time, marginals distributed as $\mathcal{N}(0, \sigma^2 = 9 \times 10^{-4})$]. This is the “best” case scenario as such noise is well behaved and has low tail weight. (Bottom) Histogram of the deviation Δ of the integrated signal within a pulse from the ideal value, taken over 10^5 pulses. Note the logarithmic Y axis. In the discussion of the behavior of Δ , the region $|\Delta| \geq \Delta_*$ is illustrated as a yellow shaded region, while a “never exceed” value of Δ is illustrated as a red shaded region. Due to the unbounded nature of the noise, arbitrarily large deviations can be encountered.

time-varying Hamiltonian control parameter is a real number (analog), the deviation in the control system from the ideal output is also characterized by a real number which is time varying and in principle unbounded.

This raises the question of the correct figure of merit for reliability. A failure corresponds to a large deviation in a control signal such that the executed gate is very far from the ideal gate. If such a failure occurs in a component used for global control, the resulting error burst may be fatal for a fault-tolerant quantum memory [33].

Let g denote a global control parameter that controls the execution of many gates in parallel in an error-corrected quantum memory, and let $\Delta \in \mathbb{R}$ denote the deviation of g from its ideal value. Naively, to ensure good performance, one might require that $|\Delta| \leq \Delta_*$ be satisfied with high probability in each syndrome extraction step for some suitably small Δ_* . However, if the memory is sufficiently resilient to brief error bursts, then occasional large excursions of g with $|\Delta|$ significantly larger than Δ_* (but not too much larger) can also be tolerated. This weaker control requirement is easier to achieve in practice. We illustrate an example trace of g in Fig. 8 and the distribution of integrated deviation for each control pulse.

For illustrative purposes, we proceed to describe how rare error bursts might arise from global control in an optical tweezer array, and to note that other quantum computing architectures have similar vulnerabilities.

In platforms based on atomic, molecular, and optical (AMO) physics, global control is often very natural due to the resource-intensive nature of coherent light sources. Here, we describe the control method of a particular neutral atom quantum computing experiment reported on in [34], but one can expect similar error mechanisms to be present in other experiments.

Evered *et al.* [34] report on a high-fidelity ($\approx 99.5\%$) entangling gate in a quantum computing experiment based on neutral rubidium-87 atoms trapped in optical tweezers. The qubit computational basis is formed by the two hyperfine states $|0\rangle = |F=0, m_F=0\rangle$ and $|1\rangle = |F=0, m_F=2\rangle$. Entanglement of two nearby atoms can be realized by transferring the $|1\rangle$ state to a state with high principal quantum number known as a Rydberg state. These excited states have strong interactions such that, after returning to the qubit subspace, the state $|11\rangle$ accumulates a phase while $|00\rangle$, $|10\rangle$, and $|01\rangle$ do not.

In [34], the physical entangling gate is applied to all atoms in a predetermined region by shining 420-nm and 1013-nm laser light modulated by an acousto-optic modulator over the course of approximately 200 ns. This corresponds to a two-photon transition to an $n=53$ Rydberg state via an intermediate state from which the laser light is detuned. Such a gate scheme naturally lends itself to scaling up because the control degrees of freedom are mostly independent of the

number of atoms in the gate region. However, variation in the pulse shape may cause a large global error on all atoms in the gate zone. For example, intensity noise can change the two-photon Rabi frequency, leading to over- and under-rotation and a final state that is not completely in the qubit subspace (i.e., a coherent leakage error). For moderate values of the rotation error, this deviation from ideal control results in an error burst affecting many qubits.

Because such error bursts are rare events, it is difficult to detect them directly except by running an error-correction experiment or some other specialized experiment that gathers a large amount of data. We are not aware of any dedicated search for error bursts in AMO platforms performed to date, but imminent experiments running full-scale quantum error correction are likely to provide instructive quantitative information about the nature and rate of such correlated errors.

A fundamental question is that of mitigation strategies after identifying a global noise source. As global noise sources are highly setup specific, the mitigation strategy will also be setup specific. Our numerics provide a guideline for how much mitigation is necessary. Using the previous example of laser intensity noise on a global gate, closed loop control may suffice for low-frequency (slowly varying) noise. For high-frequency noise, a time-multiplexing strategy where gates are only done on subsets of atoms at a time may suffice. One could also take advantage of the long coherence time of atomic platforms and “pause” the computation until the noise returns to an acceptable range.

Correlated gate errors resulting from global control are not unique to AMO platforms; in particular, spin-qubit architectures frequently rely on global control. For instance, in nitrogen-vacancy center qubits in diamonds, architectures commonly feature global optical and microwave control signals [35–37]. Similarly, in electron spin qubits, global microwave or DC electrical control is a common architectural feature [38–40]. In all of these cases, classical control noise can induce fidelity correlations among gates performed in parallel, and rare large deviations of control signals from their ideal values could cause error bursts.

APPENDIX B: ERROR BURST DETECTION

Here, we show that *a priori* knowledge of the time at which an error burst occurs is not necessary. The presence of an error burst can be inferred directly from the syndrome. We will prove bounds on the probability of a false positive or false negative in a toy model using a maximum likelihood estimator. For simplicity, consider a subset of vertices V in a decoding graph contained in a single time slice such that each vertex has w fault locations incident to it and no two vertices in V are adjacent. Then, suppose that each fault location contains an error with the background error probability p while during an error burst it contains an error with probability $p_B > p$ less than $\frac{1}{2}$.

The assumption of uniform noise and degree can be straightforwardly relaxed, and the subset of vertices can be found by exploiting structure in the decoding graph. If the decoding graph is sparse, one can also find a greedy approximate solution to the maximum independent set problem [41].

We wish to determine whether, at an arbitrary time t , there was an error burst. Define $p_1 = \sum_{i=1}^{\lceil w/2 \rceil} \binom{w}{2i-1} p^{2i-1} (1-p)^{w-2i+1}$, $p_2 = \sum_{i=1}^{\lceil w/2 \rceil} \binom{w}{2i-1} p_B^{2i-1} (1-p_B)^{w-2i+1}$ to be the marginal probabilities that an odd number of faulty edges are incident to a fixed vertex of V . Let q taking values in $\{p_1, p_2\}$ be the random variable corresponding to the marginal flip probability of syndrome vertices in V at time t . Let x be the syndrome bit string. Since the vertices of V were picked such that they are independent, the Hamming weight $|x|$ is a sufficient statistic. The maximum likelihood estimate of q is

$$\begin{aligned} \hat{q} &= \operatorname{argmax}_{\tilde{q} \in \{p_1, p_2\}} \ln \Pr(|x| | q = \tilde{q}) \\ &= \operatorname{argmax}_{\tilde{q} \in \{p_1, p_2\}} [|x| \ln \tilde{q} + (n - |x|) \ln(1 - \tilde{q})]. \end{aligned}$$

The maximization problem can be conveniently rewritten in terms of the log-likelihood ratio $\ell(|x|)$:

$$\ell(|x|) = \left[|x| \ln \frac{p_1}{p_2} + (n - |x|) \ln \frac{1 - p_1}{1 - p_2} \right].$$

The value of \hat{q} is then

$$\hat{q} = \begin{cases} p_1, & \ell(|x|) \geq 0 \\ p_2, & \ell(|x|) < 0. \end{cases}$$

Let $D(a||b) = a \ln \frac{a}{b} + (1-a) \ln \frac{1-a}{1-b}$ be the Kullback-Leibler divergence between two Bernoulli distributions with parameters a and b , respectively. Employing a Chernoff bound on $|x|$, we have an upper bound on the probability of a false positive

$$\Pr(\hat{q} = p_2 | q = p_1) = \Pr(\ell < 0 | q = p_1) \quad (\text{B1})$$

$$= \Pr\left(\frac{|x|}{n} > \alpha | q = p_1\right) \quad (\text{B2})$$

$$\leq e^{-nD(\alpha||p_1)}, \quad (\text{B3})$$

where $\alpha = \frac{\ln \frac{1-p_2}{1-p_1}}{\ln \frac{p_1(1-p_2)}{p_2(1-p_1)}} \in (p_1, p_2)$. Likewise, we can upper bound the probability of a false negative:

$$\Pr(\hat{q} = p_1 | q = p_2) = \Pr(\ell \geq 0 | q = p_2) \quad (\text{B4})$$

$$= \Pr\left(\frac{|x|}{n} \leq \alpha | q = p_2\right) \quad (\text{B5})$$

$$\leq e^{-nD(\alpha||p_2)}. \quad (\text{B6})$$

To pick some representative parameters, for $p = 0.01$, $p_B = 0.1$, $n = 100$, and $w = 10$, the probability of false negatives and false positives is at most 10^{-4} .

APPENDIX C: ESTIMATION OF THRESHOLD VALUES FOR THRESHOLD EXPERIMENTS

To obtain the phase diagrams in Sec. III A, we perform memory experiments (2×10^6 samples) with surface-code distances between $d = 13$ and 27 with $T = 2d$ syndrome extraction cycles and an error burst after the d th syndrome extraction cycle.

For the high p_B portions of the phase diagram, we adopt the same fitting strategy as Wang *et al.* [42] and perform a fit of the logical error rates as function of p_B (p fixed) to a quadratic

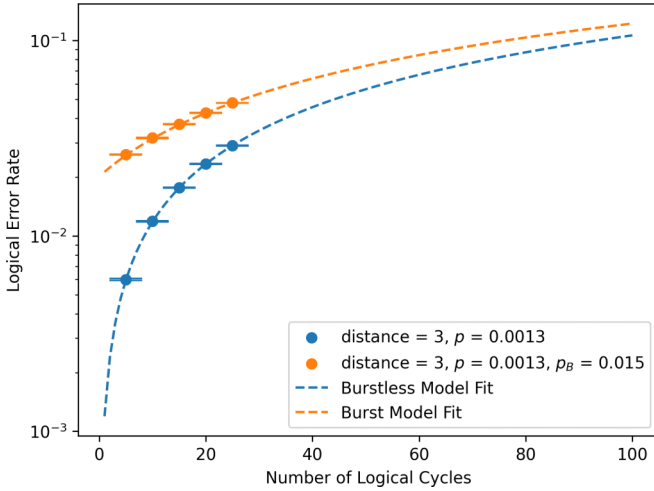


FIG. 9. Plot of overall logical error rate against number of logical cycles for background error rate $p = 0.0013$, elevated error rates $p_B = 0$ and 0.015 , and distance $d = 3$. The dashed lines are obtained with our *Ansatz* stated in Eqs. (1) and (2) using q_d and $q_{d,B}$ values for $D = 25$ logical cycles.

function in a scaled variable:

$$P_L(x) = A + Bx + Cx^2, \quad (\text{C1})$$

where $x = (p_B - p_B^*)d^{1/\nu_0}$ for code distance d , constants A, B, C , and some critical exponent ν_0 .

We determine the burst threshold p_B^* by fitting the values to (C1) and then subtracting the background error rate from it.

When the background error rate p is high, we adopt an alternate approach in order to mitigate large finite-size effects. The procedure is nearly identical to the standard procedure, but p_B is held fixed and p is swept. The fitting *Ansatz* (C1) is performed with p_B replaced by p .




FIG. 10. Logical error rates for various values of the background error rate p , the elevated error rate p_B , and the code distance d .

APPENDIX D: LOW ERROR RATE EXTRAPOLATION

To study the subthreshold behavior, we first determine the logical error rates obtained from memory experiments with various values of p and p_B and number of logical cycles D between $D = 5$ and 25 , obtaining the results plotted in Fig. 9. We derive the values for q_d and $q_{d,B}$ for different values of d , p , and p_B by using Eqs. (1) and (2). We observe that q_d and $q_{d,B}$ approach constant values as $D \rightarrow 25$ and set $\tilde{q}_d = q_d$ and $\tilde{q}_{d,B} = q_{d,B}$ for $D = 25$. We fit the plotted points corresponding to different \tilde{q}_d and $\tilde{q}_{d,B}$ in Fig. 10 to the expression $\log_{10} \tilde{q}_d = c_{\tilde{q}_d} + d m_{\tilde{q}_d}$, and similarly for $\log_{10} \tilde{q}_{d,B}$. The parameters determined by the fits shown in Table I were used for the teraquop footprint estimate.

- [1] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, [arXiv:quant-ph/9811052](#).
- [2] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
- [3] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [4] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [5] Google Quantum AI, Suppressing quantum errors by scaling a surface code logical qubit, *Nature (London)* **614**, 676 (2023).
- [6] L. Cardani, I. Colantoni, A. Cruciani, F. De Dominicis, G. D’Imperio, M. Laubenstein, A. Mariani, L. Pagnanini, S. Pirro, C. Tomei, *et al.*, Disentangling the sources of ionizing radiation in superconducting qubits, *Eur. Phys. J. C* **83**, 94 (2023).
- [7] M. McEwen, L. Faoro, K. Arya, A. Dunsworth, T. Huang, S. Kim, B. Burkett, A. Fowler, F. Arute, J. C. Bardin, *et al.*, Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits, *Nat. Phys.* **18**, 107 (2022).
- [8] Q. Xu, A. Seif, H. Yan, N. Manno, B. O. Sane, R. Van Meter, A. N. Cleland, and L. Jiang, Distributed quantum error correction for chip-level catastrophic errors, *Phys. Rev. Lett.* **129**, 240502 (2022).
- [9] C. A. Pattison, A. Krishna, and J. Preskill, Hierarchical memories: Simulating quantum LDPC codes with local gates, *Quantum* **9**, 1728 (2025).
- [10] J. M. Martinis, Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays, *npj Quantum Inf.* **7**, 90 (2021).
- [11] D. Gottesman, Stabilizer codes and quantum error correction, [arXiv:quant-ph/9705052](#).
- [12] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D’Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, Correlated charge noise and relaxation errors in superconducting qubits, *Nature (London)* **594**, 369 (2021).
- [13] M. McEwen, K. C. Miao, J. Atalaya, A. Bilmes, A. Crook, J. Bovaird, J. M. Kreikebaum, N. Zobrist, E. Jeffrey, B. Ying, A. Bengtsson, H.-S. Chang, A. Dunsworth, J. Kelly, Y. Zhang, E. Forati, R. Acharya, J. Iveland, W. Liu, S. Kim, *et al.*, Resisting

- high-energy impact events through gap engineering in superconducting qubit arrays, *Phys. Rev. Lett.* **133**, 240601 (2024).
- [14] P. M. Harrington, M. Li, M. Hays, W. V. D. Pontseele, D. Mayer, H. D. Pinckney, F. Contipelli, M. Gingras, B. M. Niedzielski, H. Stickler, J. L. Yoder, M. E. Schwartz, J. A. Grover, K. Serniak, W. D. Oliver, and J. A. Formaggio, Synchronous detection of cosmic rays and correlated errors in superconducting qubit arrays, *Nat. Commun.* **16**, 6428 (2025).
- [15] V. Iaia, J. Ku, A. Ballard, C. Larson, E. Yelton, C. Liu, S. Patel, R. McDermott, and B. Plourde, Phonon downconversion to suppress correlated errors in superconducting qubits, *Nat. Commun.* **13**, 6425 (2022).
- [16] V. D. Kurilovich, G. Roberts, L. S. Martin, M. McEwen, A. Eickbusch, L. Faoro, L. B. Ioffe, J. Atalaya, A. Bilmes, J. M. Kreikebaum, A. Bengtsson, P. Klimov, M. Neeley, W. Mruzckiewicz, K. Miao, I. L. Aleiner, J. Kelly, Y. Chen, K. Satzinger, and A. Opremcak, Correlated error bursts in a gap-engineered superconducting qubit array, [arXiv:2506.18228](https://arxiv.org/abs/2506.18228).
- [17] J. D. Chadwick, C. Kang, J. Vizslai, S. F. Lin, and F. T. Chong, Averting multi-qubit burst errors in surface code magic state factories, in *IEEE International Conference on Quantum Computing and Engineering (QCE), Montreal, QC, Canada* (IEEE, Piscataway, NJ, 2024), pp. 1089–1101.
- [18] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
- [19] C. Gidney, Stim: a fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [20] N. Delfosse and N. H. Nickerson, Almost-linear time decoding algorithm for topological codes, *Quantum* **5**, 595 (2021).
- [21] O. Higgott and C. Gidney, Sparse blossom: Correcting a million errors per core second with minimum-weight matching, *Quantum* **9**, 1600 (2025).
- [22] O. Higgott and C. Gidney, Pymatching v2, <https://github.com/oscarhiggott/PyMatching>.
- [23] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
- [24] C. T. Chubb and S. T. Flammia, Statistical mechanical models for quantum codes with correlated noise, *Ann. Inst. Henri Poincaré D* **8**, 269 (2021).
- [25] J. Ramette, J. Sinclair, N. P. Breuckmann, and V. Vuletić, Fault-tolerant connection of error-corrected qubits with noisy links, *npj Quantum Inf.* **10**, 58 (2024).
- [26] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg, Surface code quantum communication, *Phys. Rev. Lett.* **104**, 180503 (2010).
- [27] C. Gidney, M. Newman, A. Fowler, and M. Broughton, A fault-tolerant honeycomb memory, *Quantum* **5**, 605 (2021).
- [28] Google Quantum AI, Exponential suppression of bit or phase errors with cyclic error correction, *Nature (London)* **595**, 383 (2021).
- [29] C. A. Pattison, M. E. Beverland, M. P. da Silva, and N. Delfosse, Improved quantum error correction using soft information, [arXiv:2107.13589](https://arxiv.org/abs/2107.13589).
- [30] A rotated surface code only requires $2d^2 - 1$, however, a regular tiling of surface codes will require an extra qubit.
- [31] M. Lenander, H. Wang, R. C. Bialczak, E. Lucero, M. Mariantoni, M. Neeley, A. O’Connell, D. Sank, M. Weides, J. Wenner, *et al.*, Measurement of energy decay in superconducting qubits from nonequilibrium quasiparticles, *Phys. Rev. B* **84**, 024501 (2011).
- [32] T. Thorbeck, A. Eddins, I. Lauer, D. T. McClure, and M. Carroll, Two-level-system dynamics in a superconducting qubit due to background ionizing radiation, *PRX Quantum* **4**, 020356 (2023).
- [33] An extremely large global coherent error deviates from our definition of an error burst as elevated independent noise, but moderately large coherent errors are approximately captured by our definition.
- [34] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, *et al.*, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature (London)* **622**, 268 (2023).
- [35] P. Neumann, R. Kolesov, B. Naydenov, J. Beck, F. Rempp, M. Steiner, V. Jacques, G. Balasubramanian, M. Markham, D. Twitchen, *et al.*, Quantum register based on coupled electron spins in a room-temperature solid, *Nat. Phys.* **6**, 249 (2010).
- [36] C. E. Bradley, J. Randall, M. H. Abobeih, R. C. Berrevoets, M. J. Degen, M. A. Bakker, M. Markham, D. J. Twitchen, and T. H. Taminiau, A ten-qubit solid-state spin register with quantum memory up to one minute, *Phys. Rev. X* **9**, 031045 (2019).
- [37] M. H. Abobeih, Y. Wang, J. Randall, S. Loenen, C. E. Bradley, M. Markham, D. J. Twitchen, B. M. Terhal, and T. H. Taminiau, Fault-tolerant operation of a logical qubit in a diamond quantum processor, *Nature (London)* **606**, 884 (2022).
- [38] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. L. Hollenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson, Silicon quantum electronics, *Rev. Mod. Phys.* **85**, 961 (2013).
- [39] C. D. Hill, E. Peretz, S. J. Hile, M. G. House, M. Fuechsle, S. Rogge, M. Y. Simmons, and L. C. Hollenberg, A surface code quantum computer in silicon, *Sci. Adv.* **1**, e1500707 (2015).
- [40] M. Veldhorst, H. Eenink, C.-H. Yang, and A. S. Dzurak, Silicon CMOS architecture for a spin-based quantum computer, *Nat. Commun.* **8**, 1766 (2017).
- [41] M. Halldórsson and J. Radhakrishnan, Greed is good: Approximating independent sets in sparse and bounded-degree graphs, in *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 1994), pp. 439–448.
- [42] C. Wang, J. Harrington, and J. Preskill, Confinement-Higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory, *Ann. Phys.* **303**, 31 (2003).