

# Practical Quantum Error Correction with the XZZX Code and Kerr-Cat Qubits

Andrew S. Darmawan,<sup>1,2,\*</sup> Benjamin J. Brown<sup>3</sup>, Arne L. Grimsmo<sup>3</sup>, David K. Tuckett<sup>3</sup> and Shrutu Puri<sup>4,5,†</sup>

<sup>1</sup>Yukawa Institute for Theoretical Physics (YITP), Kyoto University, Kitashirakawa Oiwakecho, Sakyo-ku, Kyoto 606-8502, Japan

<sup>2</sup>Japan Science and Technology Agency (JST), Precursory Research for Embryonic Science and Technology (PRESTO), 4-1-8 Honcho, Kawaguchi, Saitama 332-0012, Japan

<sup>3</sup>Centre for Engineered Quantum Systems, School of Physics, University of Sydney, Sydney, New South Wales 2006, Australia

<sup>4</sup>Department of Applied Physics, Yale University, New Haven, Connecticut 06511, USA

<sup>5</sup>Yale Quantum Institute, Yale University, New Haven, Connecticut 06511, USA



(Received 11 May 2021; accepted 9 August 2021; published 16 September 2021)

The development of robust architectures capable of large-scale fault-tolerant quantum computation should consider both their quantum error-correcting codes and the underlying physical qubits upon which they are built, in tandem. Following this design principle, we demonstrate remarkable error-correction performance by concatenating the XZZX surface code with Kerr-cat qubits. We contrast several variants of fault-tolerant systems undergoing different circuit-noise models that reflect the physics of Kerr-cat qubits. Our simulations show that our system is scalable below a threshold gate infidelity of  $p_{CX} \sim 6.5\%$  within a physically reasonable parameter regime, where  $p_{CX}$  is the infidelity of the noisiest gate of our system, the controlled-NOT gate. This threshold can be reached in a superconducting-circuit architecture with a Kerr nonlinearity of 10MHz, an approximately 6.25-photon cat qubit, single-photon lifetime of  $\gtrsim 64 \mu\text{s}$ , and a thermal photon population  $\lesssim 8\%$ . Such parameters are routinely achieved in superconducting circuits.

DOI: [10.1103/PRXQuantum.2.030345](https://doi.org/10.1103/PRXQuantum.2.030345)

## I. INTRODUCTION

A scalable quantum computer will require a large number of almost perfect qubits. To this end, impressive progress has been made in engineering high-quality quantum systems in the laboratory that can perform some requisite set of operations that are needed for quantum computation [1–14]. While we cannot expect to reduce qubit noise arbitrarily with laboratory engineering alone, the theory of quantum error correction [15–20] has shown that we can produce logical qubits that experience next to no noise by using a redundancy of high-quality physical qubits, provided that the noise strength on the physical qubits as they perform computational operations is below some threshold value [21]. Nevertheless, it still remains a significant

experimental challenge to produce scalable quantum error-correcting codes using available technology. In order to overcome this challenge, we should co-design the code in unison with the underlying qubit architecture. In this way, we can build better error-correcting systems that specifically target the dominant sources of errors.

A promising approach to realize a high-quality physical qubit is to encode it in a bosonic mode [7–13,22–34]. Broadly speaking, encoding a qubit in the larger Hilbert space of an oscillator gives more room to protect the encoded information from noise, either through active error correction or autonomous stabilization of the code space. One of the simplest incarnations of this idea is the Kerr-cat qubit, where orthogonal superpositions of two coherent states  $|\pm\alpha\rangle$  define a physical qubit [23,35–37]. The Kerr-cat qubit is realized by parametric pumping of a Kerr-nonlinear oscillator, as recently demonstrated in a superconducting-circuit platform using a superconducting nonlinear asymmetric inductive element (SNAIL) [11]. This qubit is stabilized against bit flips. Specifically, bit flips are exponentially suppressed with the photon number  $|\alpha|^2$ , leading to a highly biased noise channel for large  $\alpha$  [37]. Surprisingly, it has recently been proposed that the controlled-NOT gate (CX) can be executed with

\*andrew.darmawan@yukawa.kyoto-u.ac.jp

†shruti.puri@yale.edu

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

Kerr-cat qubits without changing this noise profile [37]. This enables us to go beyond schemes where it is assumed that only diagonal entangling gates can be used without introducing bit-flip errors at a significant rate [38,39].

A number of error-correction protocols have been proposed recently that are adapted specifically to biased noise [31,38–49]. In particular, the XZZX surface code [50,51], a surface-code variant [16,17,52], has recently been shown to have a very high threshold under biased noise. This is because it mimics the behavior of the classical repetition code in the limit that a particular type of Pauli error is dominant. As such, minimum-weight perfect-matching decoders [17,53,54] have been developed by exploiting the symmetries of this code [16,17,45,55]. These practical decoders demonstrate exceptionally high fault-tolerant thresholds when the dominant source of noise is dephasing [51].

The error-correction capabilities of the XZZX code complement the biased noise of the Kerr-cat qubits. This motivates the present work, where we demonstrate a very robust architecture for scalable fault-tolerant quantum computing by combining Kerr-cat qubits and the XZZX code. To this end, we perform elaborate simulations for a concatenated scheme undergoing a circuit error model that reflects noise sources that have been identified in recent experimental demonstrations of Kerr-cat qubits [11]. The simulation of these details enables us to express the threshold error rates that we obtain in terms of real laboratory parameters to interrogate the experimental viability of our proposal directly.

Figure 1 shows what a Kerr-cat XZZX chip may look like if implemented with SNAILs. The dominant physical error channels for Kerr-cat qubits are photon loss and gain, quantified by a single-photon loss rate  $\kappa$  and the thermal population  $n_{\text{th}}$  of the oscillator [11]. The noise model of the Kerr-cat qubit is a function of  $\kappa$  and  $n_{\text{th}}$ , as well as the “cat size”  $|\alpha|^2$  and the Kerr nonlinearity of the oscillator  $K$ . With  $|\alpha|^2 = 6.25$  and a thermal population of 8%, we find a threshold of  $\kappa/K = 2.5 \times 10^{-4}$ . This corresponds to a total infidelity of approximately 6.5% for the noisiest logical operation in the error-correction circuit, CX, for which pure Pauli-Z errors occur approximately 351 times more frequently than other Pauli errors during the gate. This noise threshold can be reached, for example, using a superconducting Kerr-cat realization with  $K/(2\pi) = 10$  MHz, a parametric pump strength of 62.5 MHz, and a single-photon lifetime  $\kappa^{-1} = 63.6 \mu\text{s}$ . We remark that the experimental demonstration of the Kerr-cat qubit in Ref. [11] has shown that frequency fluctuations are heavily suppressed when the SNAIL is pumped into the cat subspace, such that the requirements on flux noise (characterized by the  $T_2$  time of the unpumped SNAIL) are greatly relaxed compared to frequency-tunable transmons. These parameters are routinely achieved in superconducting circuits [56,57], suggesting that Kerr cats




FIG. 1. The schematic for a possible realization of the XZZX surface code, using a two-dimensional lattice of nearest-neighbor capacitively coupled SNAIL-based Kerr-cat qubits. Data qubits are represented in magenta and ancilla qubits in green. A dc flux line is present to bias the SNAIL at the point of the desired nonlinearities. The lattice arrangement is such that neighboring SNAILs resonate at different frequencies to minimize crosstalk. Microwave drives applied at appropriate frequencies through the out-of-plane interconnects parametrically activate gate operations [11,37]. Importantly, no extra nonlinear control and/or coupling elements are required, unlike other biased-noise bosonic qubits. For direct readout of ancilla and data qubits, an additional readout resonator is available at each site. This readout resonator also provides the ability to parametrically introduce two-photon dissipation in the SNAILs, which could be used for autonomous leakage correction.

well below threshold may be realized in the near future, making the proposed architecture a promising approach to large-scale fault-tolerant quantum computing.

While our protocol achieves exceptional performance using Kerr-cat qubits, we also show that similar improvements may be attainable in other biased-noise architectures. Using the XZZX code with a generic biased-circuit-noise model, we observe a roughly factor-of-2 improvement in the threshold error rate over the Calderbank, Shor, Steane (CSS) version of the surface code. Surprisingly, we still find that the XZZX code has a 50% higher threshold error rate than the CSS surface code when using standard CX gates rather than bias-preserving CX gates.

We also perform exact simulations on small codes that represent near-term experiments. These simulations allow us to study non-Pauli noise models with optimal decoders

and reveal that the high performance of our scheme at both high and low biases is relatively insensitive to suboptimal implementation of gates and decoding compared to other schemes.

This paper is structured as follows. In Sec. II, we review the definition of the XZZX surface code and describe how it can be implemented in a Kerr-cat architecture. In Sec. III, we review Kerr-cat qubits and describe the bias-preserving gates available to them. In Sec. III B, we describe the important noise processes affecting Kerr-cat qubits and the methods we use to simulate noisy gate execution. In Sec. IV, we present the results of numerical simulations, demonstrating that very high thresholds can be achieved using XZZX codes and Kerr-cat qubits. The performance of low-distance codes under biased noise is presented in Sec. V. A discussion of the results along with potential future research directions are provided in Sec. VI. Details of the surface-code simulation methods and some advantages of Kerr-cat qubits over purely-dissipative-cat qubits are included in Appendices A and B, respectively.

## II. THE XZZX SURFACE CODE

We describe the XZZX code using the stabilizer formalism, i.e., a list of commuting Pauli operators the common  $+1$  eigenspace of which defines the code space of the code. We lay out the code with a single qubit on each vertex of the square lattice. Up to its boundaries, all of its stabilizers are the product of four Pauli terms;  $S_f = X \otimes Z \otimes Z \otimes X$  acting on the four qubits at the corners of each of the faces  $f$  of the square lattice [see Fig. 2(a)].

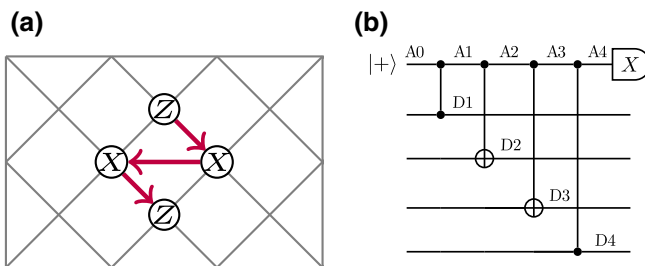


FIG. 2. (a) An XZZX surface-code check as a product of two  $X$  and two  $Z$  Pauli operators on qubits arranged on the vertices of a square lattice. Every face corresponds to a measured check operator and all checks (except for the three-qubit checks on the boundary) have the same form. The order in which data qubits of a check are coupled to the ancilla (at the center of each face) is specified by the arrows. We use the same order for every check. (b) The check-measurement circuit for the XZZX surface code. The four data qubits are coupled to an ancilla qubit prepared in the  $|+\rangle$  state before measuring the ancilla in the  $X$  basis. Distinct error locations on the ancilla qubits are labeled A0–4 and on data qubits they are labeled D1–4. The error probabilities at these locations determine the weights of the syndrome graph for decoding.

The XZZX code is locally equivalent [58,59] to the conventional CSS surface code [16,17]. As such, it inherits its code properties. We choose a lattice of size  $d_x \times d_z$  using  $n = O(d_x d_z)$  qubits, where  $d_x$  and  $d_z$  correspond to the lengths of the smallest  $X$  and  $Z$  logical operators, which have a stringlike support and traverse homologically non-trivial paths between opposite boundaries of the code. With open boundary conditions, the XZZX code encodes one logical qubit.

### A. Stabilizer measurements

We measure stabilizer operators to identify the locations of errors that occur on the qubits of the code. As the code space is defined as the  $+1$  eigenspace of the stabilizers, measuring a stabilizer in the  $-1$  state indicates that an error has occurred that takes the system out of its code space. Our stabilizer readout circuit is based on that given in Ref. [60], up to the Hadamard gates that rotate the standard surface code onto the XZZX code. Here, we briefly summarize the circuit that we use.

We measure each stabilizer generator using the circuit shown in Fig. 2(b). All of the stabilizer generators can be measured in parallel. We place a single ancilla in the  $|+\rangle$  state on each face. We subsequently apply two-qubit entangling gates between the ancilla and its four nearest-neighbor data qubits such that the ancilla collects the parity information for the stabilizer  $S_f$  on its respective face  $f$ . We finally measure the ancilla in the Pauli- $X$  basis to learn the value of the check.

In Fig. 2(a), we indicate the order in which qubits should be coupled to the ancilla. With the exception of qubits on the boundary, every data qubit is coupled with an ancilla qubit at each time step. The same ordering is used for checks on the boundary, except when there is no neighboring ancilla or data qubit to couple with, at which time the qubit is left idle. This ordering enables parallel readout of all the stabilizer operators [60]. The exception to this ordering is for the small  $n = 9$  codes in Sec. V, where we use the rotated surface-code layout. Here, on every second face, the order of application of the two CX gates is swapped. For that particular layout, this order has been shown in Ref. [61] to mitigate the effect of hook errors, in which a single error on an ancilla can propagate to two data qubits.

In a realistic device, errors may occur during gate execution, state preparation, measurements, and on idle qubits. Repeated measurement of stabilizer generators provides data, called the syndrome, that allows us to correct for all of these types of errors, provided that we adopt an appropriate decoding strategy.

### B. Decoding

A decoder is a classical algorithm that uses the error syndrome as an input to estimate the best choice of

correction that can be applied to recover the encoded state. We use a minimum-weight perfect-matching (MWPM) decoding algorithm to correct for all types of errors [17]. Note that correctly identifying the error, up to stabilizer equivalence, is equivalent to finding a correction for that error. We format the error syndrome as a list of defects, where we say that a defect appears at face  $f$  and at time  $t$  if  $S_f(t-1)S_f(t) = -1$ , where  $S_f(t)$  denotes the outcome of the stabilizer measurement  $S_f$  made at time  $t$ . Describing the error syndrome as a list of defects means that errors can be interpreted as strings where defects lie at their endpoints [17]. This interpretation enables us to employ MWPM to estimate the error that has occurred.

The MWPM algorithm takes a graph with weighted edges and returns a matching such that the sum of the weights of the matching is minimal. Given the error model and the syndrome, the graph input to MWPM is constructed such that each node corresponds to a defect and the weight of an edge connecting two defects is  $-\log p$ , where  $p$  is the probability of the most likely error that can give rise to that pair of defects. This choice of weights implies that the edges returned by the MWPM algorithm correspond globally to the most likely error consistent with the observed syndrome.

This practical decoder is particularly well suited for decoding biased noise using the XZZX code. For the XZZX code, a  $Z$  error on any data qubit will produce a pair of defects displaced in the same direction, independent of the qubit, and this is perpendicular to the displacement produced by an  $X$  error. Thus, when the noise is strongly biased toward  $Z$  errors, the edge weights in the syndrome graph have much higher weight in one direction. This allows errors to be identified less ambiguously using MWPM, compared to when edges in both directions are more equally weighted, and results in a high threshold under biased noise. This has been demonstrated using a phenomenological noise model in Ref. [51].

In realistic circuit-level noise, single error events can give rise to pairs of defects separated “diagonally,” i.e., in both space and time, as well as hook errors, which occur when the ancilla qubit experiences a  $X$  error halfway through the stabilizer readout circuit [an  $X$  error at A2 in Fig. 2(b)]. This error is copied onto two of the data qubits by the entangling gates. However, as the hook error is caused by a low rate event, the effect of these errors is relatively benign for biased noise. Our decoder takes all of these types of errors into account by appropriately assigning edge weights in the construction of the syndrome graph.

### III. CAT QUBITS AND BIAS-PRESERVING GATES

Here, we describe the Kerr cat [11,36], which, in our proposal, forms the elementary qubit from which the XZZX code is built. The Kerr-cat qubit is realized in a

two-photon driven Kerr nonlinear oscillator. The Hamiltonian of such an oscillator in a frame rotating at the oscillator frequency  $\omega_l$  is given by

$$H_{\text{cat}} = -Ka^{\dagger 2}a^2 + K\alpha^2(a^{\dagger 2} + a^2). \quad (1)$$

Here,  $K > 0$  is the strength of the nonlinearity while  $K\alpha^2$  is the amplitude of the two-photon drive. The orthogonal even- and odd-parity cat states  $|\mathcal{C}_{\alpha}^{\pm}\rangle = \mathcal{N}_{\pm}(|\alpha\rangle \pm |-\alpha\rangle)$  are degenerate eigenstates of this Hamiltonian. As illustrated in Fig. 3, the degenerate cat subspace  $\mathcal{C} = \text{span}\{|\mathcal{C}_{\alpha}^{\pm}\rangle\}$  is separated from the rest of the Hilbert space  $\mathcal{C}_{\perp}$  by an energy gap, which in the limit of large  $\alpha$  is well approximated by  $|\omega_{\text{gap}}| \sim 4K\alpha^2$ . We work in the basis where the cat states  $|\mathcal{C}_{\alpha}^{+}\rangle \equiv |+\rangle$  and  $|\mathcal{C}_{\alpha}^{-}\rangle \equiv |-\rangle$  are aligned along the  $X$  axis of the Bloch sphere and their superpositions  $(|\mathcal{C}_{\alpha}^{+}\rangle + |\mathcal{C}_{\alpha}^{-}\rangle)/\sqrt{2} \equiv |0\rangle \simeq |\alpha\rangle$  and  $(|\mathcal{C}_{\alpha}^{+}\rangle - |\mathcal{C}_{\alpha}^{-}\rangle)/\sqrt{2} \equiv |1\rangle \simeq |-\alpha\rangle$  are aligned along the  $Z$  axis (where the approximations hold for large  $|\alpha|$ ).

The Kerr-cat qubit has recently been realized experimentally with a capacitively shunted SNAIL [11]. The SNAIL is a flux-biased nonlinear device comprising of four Josephson junctions (Fig. 1). Both third- and fourth-order (Kerr) nonlinearities exist in the SNAIL at an appropriate value of the external magnetic flux. Due to the third-order nonlinearity, a single microwave drive applied at twice the resonance frequency generates the two-photon term in the Hamiltonian in Eq. (1).

It has been shown that the error channel of the Kerr-cat qubit is dominated by phase-flip errors, while  $X$  and  $Y$  errors are exponentially suppressed with the average photon number  $\langle a^{\dagger}a \rangle \sim |\alpha|^2$  [37]. Some noise sources, such as thermal excitation, frequency fluctuations, etc., may introduce leakage errors. Leakage can, to some extent, be autonomously corrected or “cooled” by introducing any additional photon dissipation channel without adversely affecting the bias [37]. A two-photon dissipation source is preferred for such autonomous leakage correction.

The inherent nonlinearity in the Kerr cat enables a straightforward realization of two-photon dissipation for autonomous leakage correction. For example, consider the setup in Fig. 4, which is the building block of the surface-code architecture. Suppose that the resonance frequency of one of the data qubit SNAILs is  $\omega_l$ . Each SNAIL is capacitively coupled to a readout resonator of frequency  $\omega_r$ . A microwave drive applied to the SNAIL at  $2\omega_l - \omega_r$  can induce the two-photon dissipation as follows. Due to the Kerr nonlinearity in the SNAIL, the drive photons are consumed to convert two photons in the SNAIL to one photon in the readout mode. The photon in the low- $Q$  readout is rapidly lost to the environment, thus effectively inducing two-photon dissipation in the SNAIL. Depending on the drive amplitude, mixing between the readout and SNAIL mode, and the lifetime of the readout, the strength of the

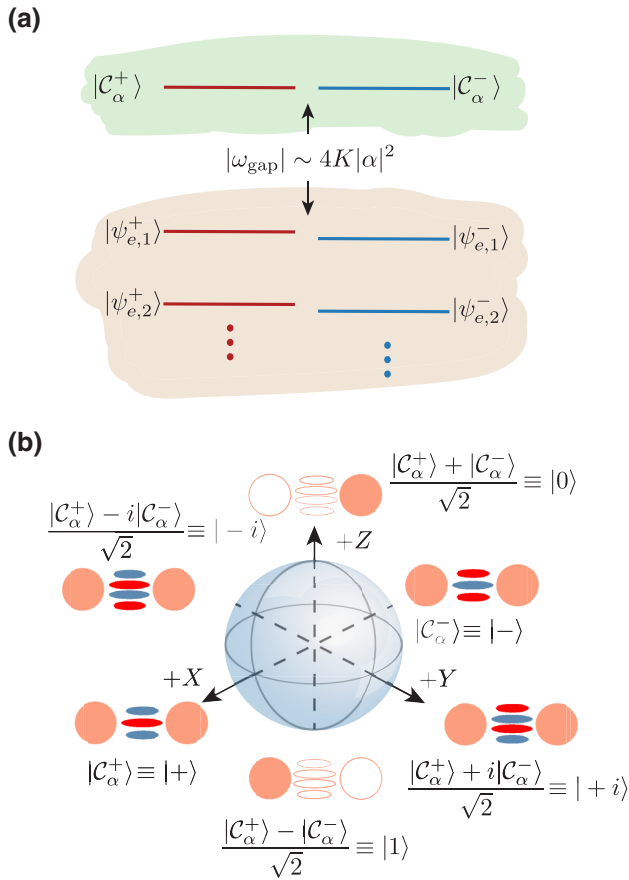


FIG. 3. (a) The eigenspectrum of the two-photon driven nonlinear oscillator in the rotating frame. The cat states  $|\mathcal{C}_\alpha^\pm\rangle$  are exactly degenerate. The eigenspectrum can be divided into an even- and an odd-parity manifold. The cat subspace, highlighted in green, is separated from the first excited state by an energy gap  $\omega_{\text{gap}} \sim 4K|\alpha|^2$ . (b) The Bloch sphere of the cat qubit. The figure also shows sketches of the Wigner functions corresponding to the eigenstates of Pauli- $X$ ,  $-Y$ , and  $-Z$  operators.

two-photon dissipation achieved can easily be engineered to be  $K/100 - K/10$ .

Single-photon loss and thermal noise have been observed to be the two main sources of errors in the experiment in Ref. [11]. Consequently, we also consider these noise sources for the simulations in the present work. Importantly, the experiment has found that non-phase-flip errors decrease with increasing cat size, confirming the theoretically predicted biased-noise property of the qubit.

### A. Kerr cats and dissipative cats

In our work, the choice of the Kerr cat as the biased-noise system used for the numerical analysis is motivated by several characteristics that make it an excellent candidate for large-scale quantum computation. First, the large nonlinearities inherently present in superconducting circuits ensure a large energy gap  $|\omega_{\text{gap}}|$ , typically in the range




FIG. 4. A drive at  $2\omega_t$  applied to the SNAIL implements the two-photon drive required to realize the Kerr-cat Hamiltonian in Eq. (1). Here,  $\omega_t$  is the resonance frequency of the SNAIL. It is also possible to activate the two-photon dissipation for autonomous leakage correction by applying a microwave drive at  $2\omega_t - \omega_r$ . Here,  $\omega_r$  is the resonance frequency of the readout resonator coupled to the SNAIL.

of 10–300 MHz. As further described in the next section, this large gap allows fast gates (10–100 ns). The typical lifetimes in these circuits, 10 – 100  $\mu\text{s}$  [4,56,57,62], are approximately 100 – 1000 times longer than the gate times, ensuring high-fidelity gate operations [11,37]. Consequently, it becomes much easier to reach the threshold parameters estimated in our work with current experimental techniques. The inherent nonlinearity can also be exploited to implement all the single- and multiqubit gate operations and measurements using microwave pulses, thus eliminating any need for external coupling elements and providing a hardware-efficient solution for scaling up [11,37].

We now contrast the Kerr-cat qubit with the purely-dissipative-cat qubit, which also has a strongly biased noise channel [25]. The dissipative cat is realized by engineering a two-photon drive and two-photon dissipation in a linear resonator via external nonlinear modes, such as a transmon [12,31,63]. So far, the strength of the two-photon dissipation induced this way has been limited to 10–100 KHz, leading to slow gate operations ( $\gtrsim 1 \mu\text{s}$ ) [64].

In addition, the lack of self-Kerr nonlinearity in purely dissipative cats makes implementation of gate operations, preparation, and measurements more challenging. This is evident, for example, from the scheme for  $X$ -basis measurement. The Kerr-cat readout (for more details, see Sec. III B) proceeds by effectively rotating the qubit around the  $Y$  axis, which maps the  $X$ -basis states to the  $Z$  basis. The  $Z$ -basis states are subsequently mapped onto a linear readout resonator and a homodyne measurement of its field reveals the state of the Kerr-cat qubit. A useful feature of this scheme is that the strength of the homodyne field, and hence the measurement fidelity, increases with the size of the Kerr cat. The  $Y$ -axis rotation is possible by freely evolving the qubit under the self-Kerr nonlinearity in the Kerr-cat mode. This rotation is not possible in purely

dissipative cats and hence this simple measurement scheme cannot be employed.

In the case of dissipative cats,  $X$ -basis measurement proceeds by adiabatic deflation of the even- and odd-parity cat states to vacuum and single-photon Fock states, respectively. The Fock states are subsequently read out via coupling to other nonlinear modes like a transmon or asymmetrically threaded SQUID [12,31]. Compared to the Kerr-cat readout scheme, this scheme is slow and its fidelity is effectively independent of cat size.

Another challenge with dissipative cats is that additional nonlinear coupling elements required to implement quantum operations may open up new sources of dissipation in the system, which can reduce the noise bias [63,64]. Finally, during gate operations, dissipative cats may suffer from nonadiabatic phase-flip errors (see Appendix B) [31]. Because this extra source of phase flips is absent in Kerr cats, it is easier to achieve the threshold parameters in practice.

Future theoretical and experimental research in optimizing the hardware and control architecture will certainly improve both biased-noise cat platforms and may even lead to development of new biased-noise qubits [48]. But because of the reasons outlined above and the current state of technology, this paper focuses on the XZZX code implementation with Kerr-cat qubits.

## B. Logical cat operations and noise channels

We assume that at the beginning of the computation, each SNAIL (data and ancilla) has been pumped into the cat-qubit basis by the adiabatic pumping technique described in Refs. [11,36]. Once in the logical cat subspace, the set of bias-preserving operations in the Kerr-cat platform, relevant for the XZZX code implementation, are the projective  $X$ -basis measurement,  $\mathcal{M}_X$ , and the gates  $S, S^\dagger, CZ$ , and  $CX$ . We use the projective measurement  $\mathcal{M}_X$  for state preparation  $\mathcal{P}_{|+\rangle}$ . We now describe how these operations can be implemented.

### 1. $S$ and $S^\dagger$ gates

Consider the Hamiltonian of a Kerr-cat qubit when, in addition to the two-photon drive, an additional resonant microwave drive of amplitude  $\varepsilon$  is applied  $H_Z = H_{\text{cat}} + (a\varepsilon^* + a^\dagger\varepsilon)$ . Projection of the interaction onto the logical subspace gives  $P_C H_Z P_C = \alpha(\varepsilon + \varepsilon^*)Z + \mathcal{O}^*(e^{-2\alpha^2})$ , where  $\mathcal{P}_C = |0\rangle\langle 0| + |1\rangle\langle 1| = |\mathcal{C}_\alpha^+\rangle\langle \mathcal{C}_\alpha^+| + |\mathcal{C}_\alpha^-\rangle\langle \mathcal{C}_\alpha^-|$  and  $Z$  is Pauli- $Z$  in the logical cat basis (we use  $\mathcal{O}^*[f(x)]$  notation to suppress polynomial factors in  $x$  [65]).  $S$  and  $S^\dagger$  gates are implemented by applying the microwave drive, in phase ( $\varepsilon = \varepsilon^* > 0$ ) and  $\pi$  out of phase ( $\varepsilon = \varepsilon^* < 0$ ) with the two-photon drive, respectively, for time  $T_S = \pi/8\alpha\varepsilon$ , see Fig. 5(a). Clearly, by increasing the cat size and the strength of the microwave drive  $|\varepsilon|$ , it is possible to decrease  $T_S$ , thereby decreasing the probability of errors

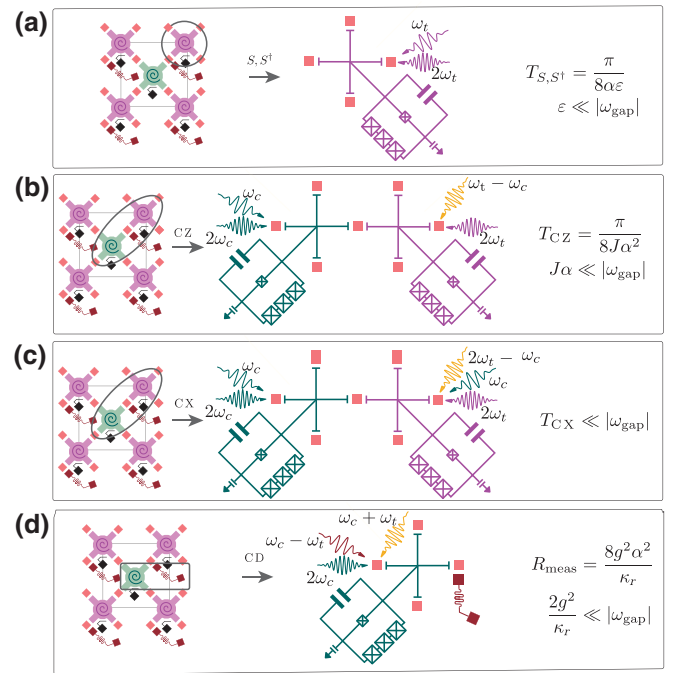


FIG. 5. A schematic showing the parametrically driven (a)  $S$  and  $S^\dagger$ , (b)  $CZ$ , (c)  $CX$ , and (d)  $CD$  operations required to realize syndrome-measurement circuits for the surface codes. The  $X(\pi/4)$  gate required for ancilla readout is not shown here. It simply requires turning off the two-photon drive on the ancilla Kerr cat for time  $\pi/2K$ . The speed of the homodyne measurement after the  $CD$  operation is characterized by the rate  $R_{\text{meas}} = 2\kappa_r|\beta|^2$ , where  $|\beta| = 2g\alpha/\kappa_r$  is the maximum of the readout-resonator displacement caused by the Kerr-cat qubit [66].

during the gate operation. Note that the drive can cause virtual transitions out of the computational cat subspace. Consequently, we must have  $|\varepsilon| \ll |\omega_{\text{gap}}|$  so that these virtual transitions and leakage outside the qubit manifold of approximately  $|\varepsilon/\omega_{\text{gap}}|^2$  remain negligible.

### 2. $CZ$ gate

The implementation of a  $CZ$  gate between two Kerr cats requires a Hamiltonian interaction of the form  $H_{ZZ} = H_{\text{rncat},c} + H_{\text{cat},t} - J(a_c^\dagger a_t + a_c a_t^\dagger) + J(a_c + a_c^\dagger)$  [for the corresponding drive scheme, see Fig. 5(b)]. Here, the subscripts  $c$  and  $t$  refer to the control and target qubits, respectively. Evolution under this Hamiltonian for time  $T_{CZ} = \pi/8J\alpha^2$  realizes the  $CZ$  gate. Clearly, larger  $\alpha$  and  $J$  lead to a faster gate. To minimize virtual transitions out of the computational cat subspace, we must have  $J\alpha \ll |\omega_{\text{gap}}|$  [36]. The term  $\propto (a_c^\dagger a_t + a_c a_t^\dagger)$  is a beam-splitter interaction between the two qubits and is realized by applying a drive to either of the two qubits at their frequency difference. The third-order nonlinearity in the SNAIL consumes the drive photons to convert photons in one SNAIL to photons in the other, effectively realizing the beam-splitter

Hamiltonian. The term  $\propto (a_c + a_c^\dagger)$  is realized by applying a resonant microwave drive to the control SNAIL.

### 3. CX gate

The diagonal  $S$ ,  $S^\dagger$ , and CZ gates are trivially bias preserving. A more challenging gate to realize, however, is the bias-preserving CX gate. Such a gate can be achieved by exploiting the continuous-variable nature of the cat encoding and temporarily leaving the cat-code subspace during the gate. The intuition for the CX gate is best understood by first considering a bias-preserving  $X$  (or NOT) gate. An  $X$  gate can be executed by adiabatically changing the phase of the two coherent states  $|\pm\alpha\rangle \rightarrow |\pm\alpha e^{i\phi(t)}\rangle$ , such that  $\phi(0) = 0$  and  $\phi(T) = \pi$ . At time  $T$ , we thus map  $|\pm\alpha\rangle \rightarrow |\mp\alpha\rangle$  and implement an  $X$  gate (in the large  $\alpha$  limit where  $|0/1\rangle \simeq |\pm\alpha\rangle$ ). Similarly, a CX gate can be executed if we can condition the phase  $\phi(T)$  on the state of the control qubit. As shown in Ref. [37], this can be achieved using a total Hamiltonian

$$\begin{aligned}
 H_{CX} = & -K(a_c^{\dagger 2} - \beta^2)(a_c^2 - \beta^2) \\
 & - K \left[ a_t^{\dagger 2} - \alpha^2 e^{-i2\phi(t)} \left( \frac{\beta - a_c^\dagger}{2\beta} \right) - \alpha^2 \left( \frac{\beta + a_c^\dagger}{2\beta} \right) \right] \\
 & \times \left[ a_t^2 - \alpha^2 e^{i2\phi(t)} \left( \frac{\beta - a_c}{2\beta} \right) - \alpha^2 \left( \frac{\beta + a_c}{2\beta} \right) \right] \\
 & - \frac{\dot{\phi}(t)}{4\beta} a_t^\dagger a_t (2\beta - a_c^\dagger - a_c). \quad (2)
 \end{aligned}$$

The interactions needed for the CX gate can be activated parametrically as shown in Fig. 5(c) (adapted from Ref. [37]). The CX-gate time is limited by the adiabaticity requirement  $\dot{\phi}(t) \ll |\omega_{\text{gap}}|$ .

### 4. $\mathcal{M}_X, \mathcal{P}_{|+\rangle}$

The remaining components required for the surface-code implementation are the projective  $X$ -basis measurement,  $\mathcal{M}_X$ , and preparation  $\mathcal{P}_{|+\rangle}$ . The projective measurement proceeds in two steps. The first step comprises of rotating the  $|\pm\rangle$  ancilla states to the computational basis states  $|0, 1\rangle$ , respectively, using an  $S$  gate followed by  $X(\pi/4) = (1 + iX)/\sqrt{2}$  operation (see Fig. 6).

The  $X(\pi/4)$  gate is a non-bias-preserving gate, but since this gate is being applied to the ancilla during measurement, the effect of a nondephasing error is the same as that of a dephasing error [38]. It is implemented by turning off the parametric drive and evolving the cats freely under the self-Kerr nonlinearity for time  $\pi/2K$  [11,67].

Once the  $|\pm\rangle$  states are rotated to  $|0, 1\rangle$ , a controlled-displacement (CD) operation between the Kerr-cat qubit and a linear readout resonator can be used to read out in the computational basis. The required (rotating frame) interaction Hamiltonian is  $H_{\text{readout}} = H_{\text{cat}} + g(a_1^\dagger + a_1)(a_r^\dagger +$




FIG. 6. The circuit for projective measurement and subsequent preparation of basis states in the  $X$  basis. The top line is the data qubit and the bottom line is the readout resonator. The CD gate maps the state of the data qubit to the displacement of the readout mode. A homodyne measurement of the signal from the readout resonator then reveals the state of the qubit.

$a_r)/2$ , where  $a_1$  and  $a_r$  are the annihilation operators for the Kerr-cat mode and the readout mode, respectively, and  $g$  is the beam-splitter coupling between the two modes [11,67]. Projecting onto the qubit subspace gives  $P_C H_{\text{readout}} P_C = g\alpha Z_1(a_r^\dagger + a_r)$ . From this projected Hamiltonian, it is clear that the readout resonator experiences a displacement conditioned on the state of the qubit.

If the single-photon decay rate of the readout mode is  $\kappa_r$ , then the qubit-dependent displacement amplitude is approximately  $\mp(2g\alpha/\kappa_r)(1 - e^{-\kappa_r t/2})$ . This displacement amplitude can be subsequently measured via a typical homodyne measurement and the outcome of the homodyne measurement can be used to infer the state of the ancilla qubit. Since the displacement of the readout resonator increases with  $\alpha$ , the subsequent homodyne measurement of its field becomes more efficient as the cat size is increased. Consequently, higher-fidelity measurements are possible by making the cat size larger.

The interaction Hamiltonian,  $H_{\text{readout}}$ , can be achieved by applying two microwave drives to the SNAIL: one at the frequency difference of the readout resonator and the SNAIL and the other at the sum of these two frequencies [see Fig. 5(d)]. The first drive gives rise to a beam-splitter-type coupling ( $a_1^\dagger a_r + a_r^\dagger a_1$ ), while the second drive gives the squeezing-type coupling ( $a_1^\dagger a_r^\dagger + a_r a_1$ ). The amplitudes of the two drives can be adjusted so that the strength of the beam-splitter and squeezing interactions are the same, leading to the desired Hamiltonian  $H_{\text{readout}}$ . In the limit of large-photon-number cat qubits, only one drive at the difference frequency is sufficient. In this limit, the resulting beam-splitter interaction projected in the cat basis is  $P_C(a_1^\dagger a_r + a_r^\dagger a_1)P_C \sim g\alpha Z_1(a_r^\dagger + a_r)$  to a very good approximation [ $\mathcal{O}^*(e^{-2\alpha^2})$ ]. This readout scheme, called *quadrature readout*, has been demonstrated experimentally [11]. After the measurement is complete, the ancilla is projected onto one of the computational basis states conditioned on the measurement outcome. In order to reinitialize the ancilla in  $|+\rangle$  or  $|-\rangle$  for the next round of syndrome extraction, the states can be rotated from

the computational basis by applying  $X(\pi/4)^\dagger = X(-\pi/4)$  followed by  $S^\dagger$ .

#### IV. THRESHOLDS

A crucial figure of merit for a quantum error-correcting code is the threshold: the maximum noise strength below which it becomes possible to efficiently reduce the noise strength on the logical level arbitrarily by increasing the number of physical qubits and gates. We perform numerical simulations to estimate the threshold noise strength using both a generic biased-circuit-level noise model in Sec. IV A, and a realistic Kerr-cat qubit noise model in Sec. IV B. The former demonstrates the generic performance increase of the XZZX code when using bias-preserving gates and operations. The latter provides a target noise strength in a Kerr-cat architecture based on a realistic noise model.

The threshold is calculated by numerically computing logical failure rates at different sizes and determining the maximum physical noise strength below which the logical error rate decreases exponentially with the system size. In each simulated run, the circuit implementing  $d_z$  rounds of noisy-syndrome extraction is applied, followed by a single round of noiseless-syndrome extraction. The final round can be thought of as a readout, in which each data qubit is measured individually, and without loss of generality can be considered noiseless, as any measurement error in this round can be equally regarded as having occurred in the previous (noisy) round. The decoder is run on the observed syndrome and a logical failure occurs if the physical error and the correction, computed by the decoder, are homologically distinct. Further details of the Pauli simulation method are provided in Appendix A 1.

As is standard in most numerical studies of quantum error-correcting codes, a Pauli-noise model is used in these threshold calculations, enabling efficient simulation of large system sizes. Although the noise model for Kerr-cat qubits obtained through numerically solving the Lindblad master equation is not a Pauli channel, we demonstrate with exact simulations of smaller systems, described in Sec. V B, that a Pauli approximation to the noise does not appear to significantly affect calculated logical error rates.

The results we present in this section for the XZZX code are obtained with the conventional open-boundary layout of the surface code, which has three-qubit checks on the boundary [52]. Alternative boundary conditions (including periodic in both in space and time) are also tested and similar threshold values are obtained in all cases (despite large differences in logical error rates).

As shown in Ref. [51], with the standard layout and open boundary conditions, the aspect ratio of the code can be tuned to the bias. By increasing the length of the logical  $Z$  operator compared to the logical  $X$  operator, an overall lower logical error rate can be obtained for a given

number of qubits. Unless otherwise specified, we determine the aspect ratio using small systems by varying the length of the code relative to its width until the logical  $X$  error rate is comparable to the logical  $Z$  error rate near threshold. This aspect ratio is then fixed when scaling up the code. To make more system sizes computationally accessible, we restrict to aspect ratios such that the length of the code is a multiple of its width. As with boundary conditions, varying the aspect ratio does not appear to appreciably change the value of the threshold, provided that the aspect ratio is fixed as the system is scaled up.

We also perform simulations of the standard CSS version of the surface code and an alternative variant of the surface code proposed for biased noise, which we call the tailored surface code (TSC). The TSC is simply the standard surface code but with  $Z$  checks replaced with  $Y$  checks, such that a single phase-flip error on a data qubit creates four adjacent defects. This version of the surface code has been studied in Refs. [41,42], where it has been found to have strong resilience to phase-flip noise. However, we find that the thresholds obtained with the TSC with circuit-level noise and a modified MWPM decoder are somewhat lower than those of the XZZX code, so we do not present them here. Nevertheless, the TSC may have some advantages for finite-sized systems, which we discuss in Sec. V.

##### A. Generic biased-circuit-noise thresholds

First, we estimate the threshold of the XZZX surface code using a generic biased-circuit-noise model. We compare this to the threshold of the standard CSS version of the surface code under the same noise model and find that the XZZX code has a significantly higher threshold. Surprisingly, we also find that the XZZX code maintains a large advantage over the CSS surface code even if a standard CX gate for two-level qubits, rather than a bias-preserving CX gate, is used.

The generic noise model has two parameters  $p_z$  and the noise bias  $\zeta$ . In the following, we assume, without loss of generality, that the Pauli errors are applied after the gate. On the CZ gates, the errors  $Z_a \otimes I_d$  and  $I_a \otimes Z_d$ , are each assigned the probability  $p_z$  and all other two-qubit Pauli errors are assigned the probability  $p_z/\zeta$ . We remark that in this simplistic toy model, we are assuming the probability of a correlated  $Z_a \otimes Z_d$  error to be the same as the probability of nondephasing errors  $p_z/\zeta$ . This gives a total error probability (or infidelity) of  $p_z(2 + 13/\zeta)$  per CZ gate. Here, the subscripts  $d$  and  $a$  refer to the data and ancilla qubits, respectively. The single-qubit noise, including idle noise on data qubits and noise during ancilla preparation, has a single-qubit Pauli error model with the  $X$ ,  $Y$ , and  $Z$  error probabilities given by  $p_x/\zeta$ ,  $p_y/\zeta$ , and  $p_z$ , respectively. Faulty measurements are modeled by



flipping the ancilla measurement outcomes with probability  $p_z + p_z/\zeta$ .

In the case of ordinary discrete two-level qubits, evolution under an interaction Hamiltonian of the form  $\chi [(1 + Z_a)I_d/2 + (1 - Z_a)X_d/2]$  for time  $\pi/2\chi$  would be required for the CX operation. It is easy to show that if the probability of a  $Z$  error during the gate is  $p_z$ , then the resulting error channel is (writing only the Pauli terms for brevity)  $p_z Z_a I_d \rho Z_a I_d + p_z' I_a Z_d \rho I_a Z_d + p_z'' Z_a Z_d \rho Z_a Z_d + p_z''' I_a Y_d \rho I_a Y_d + p_z'''' Z_a Y_d \rho Z_a Y_d$ , with  $p_z' = (2p_z/\pi) \int_0^{\pi/2} \cos^4(\phi) d\phi \sim 0.375p_z$  and  $p_z'' = (2p_z/\pi) \int_0^{\pi/2} \cos^2(\phi) \sin^2(\phi) d\phi \sim 0.125p_z$ . Consequently, we use this equation for probability assignment for the  $I_a Z_d$ ,  $Z_a I_d$ ,  $Z_a Z_d$ ,  $I_a Y_d$ , and  $Z_a Y_d$  errors, while all other two-qubit Pauli errors are assigned the probability  $p_z/\zeta$ . The total infidelity of the CX gate in this case is  $p_z(2 + 10/\zeta)$ .

For the bias-preserving CX gate in the generic noise model, unlike its ordinary counterpart described above, phase-flip errors on the data qubits may convert to  $Z_a \otimes Z_d$  but they do not convert to  $Z_a \otimes Y_d$  or  $I_a \otimes Y_d$  errors. Consequently, in this case,  $I_a \otimes Z_d$  and  $Z_a \otimes Z_d$  errors are assigned the probability  $p_z/2$ ,  $Z_a \otimes I_d$  is assigned the probability  $p_z$ , and all other two-qubit Pauli errors are assigned the probability  $p_z/\zeta$ . Note that the total infidelity of the CX gate in this case is larger than the infidelity of the generic non-bias-preserving CX gate by  $2p_z/\zeta$ . For large  $\zeta$ —for example,  $\zeta = 100$  used here—this difference becomes negligibly small. The resulting logical error probability versus  $p_z$  is plotted for the surface code with both bias-preserving and non-bias-preserving CX gates in Fig. 7(a).

For comparison, we also perform simulations of the conventional CSS surface code using the bias-preserving

CX gate and the same generic biased-error model. In this case, we use the “rotated” layout of Ref. [68] with open boundary conditions and two-qubit checks on the boundary, which has a code distance that is a factor of  $\sqrt{2}$  larger than the XZZX code on the unrotated geometry using an equal number of physical qubits. The resulting logical error rates are shown in Fig. 7(b). We do not use the rotated layout for the XZZX code, since the all- $Z$  logical operator in this case is oriented at  $45^\circ$  to the lattice boundaries and we therefore cannot improve the resilience of the code against phase flips by varying the aspect ratio of the code (unlike with the standard layout). We also present the logical failure rate as a function of  $\sqrt{n}$  for both the XZZX code and the CSS surface code at physical error rate  $p = 0.005$  in Fig. 7(c).

The procedure for choosing the aspect ratio for the XZZX and CSS codes using bias-preserving CX gates is as described above, which yields the same aspect ratio for both codes near threshold of around 3 : 1. To simplify comparison, we use the same aspect ratio for the XZZX code with standard CX gates; however, this may not be optimal.

As can be seen, the XZZX code using a bias-preserving gate achieves the highest threshold  $p_z = 0.98 \pm 0.05\%$ , corresponding to a two-qubit error rate of close to 2%. This is approximately double the threshold of the CSS surface code with the same noise model, which achieves a threshold of  $p_z = 0.50 \pm 0.05\%$ . The XZZX code with a standard non-bias-preserving CX gate performs worse than the XZZX code with a bias-preserving gate, achieving a threshold of  $0.80 \pm 0.05\%$ . However, this is clearly much higher than the CSS surface-code threshold, suggesting that the XZZX code may be a good choice in biased-noise

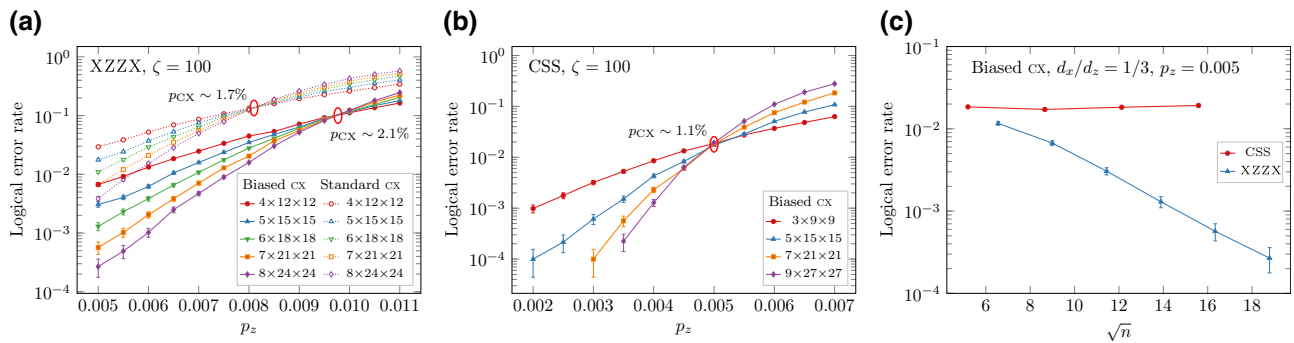


FIG. 7. The error thresholds of surface-code variants with a generic circuit-noise model and bias  $\zeta = 100$ . In (a), we compare the threshold of the XZZX code using bias-preserving CX gates (solid lines) to its threshold using standard non-bias-preserving CX gates (dotted lines). We find that the availability of bias-preserving CX gates leads to an approximately 25% improvement in the threshold. This result highlights the fact that along with total gate error, the structure of noise in the gate plays an important role in determining the performance of the surface-code variant. In (b), the threshold of the CSS code using bias-preserving CX gates is shown for comparison and is seen to be approximately half that of the XZZX code using the same bias-preserving gates. In this model, both types of CX gate have error rates approximately equal to  $2p_z$  (for details of the noise model, see the text). The crossing points are circled and the total CX error probability  $p_{CX}$  at these points is shown. In (c), we plot the logical error rate of both code variants against  $\sqrt{n}$  near the threshold of the CSS code, found in (b), where  $n$  is the number of data qubits. In all plots, the error bars indicate one standard deviation.

architectures even when a bias-preserving CX gate is not available.

### B. Kerr-cat noise thresholds

To obtain the noise models used for threshold estimates of the XZZX Kerr-cat code, we first perform a master-equation simulation of the operations needed for the syndrome-extraction circuit (Fig. 2): the CX and CZ gates, as well as the measurement and preparation schemes described above. The simulations are performed for two cat sizes,  $\alpha^2 = 1$  and  $\alpha^2 = 6.25$ .

The noise channel of the larger cat is more biased than that of the smaller one. Based on the observations in the recent experimental work [11], we subject the qubit to a white thermal-noise channel represented by the Lindbladian  $\kappa_1(1 + n_{\text{th}})\mathcal{D}[a]\rho + \kappa_1 n_{\text{th}}\mathcal{D}[a^\dagger]\rho$ . Thermal noise introduces leakage and in order to autonomously correct some of the leakage, we add an additional two-photon dissipation to the Lindbladian  $\kappa_2\mathcal{D}[a^2]\rho$ , where the rate of two-photon dissipation  $\kappa_2$  is fixed to  $\kappa_2 = K/10$ .

When the two-photon dissipation is added, care must be taken to adjust the phase of the two-photon drive to preserve the cat-qubit subspace [37]. In our simulations, the two-photon dissipation is added *while* the gates are being implemented for all but the CX gate. The CX gate requires a time-dependent two-photon dissipation, which becomes hard to numerically simulate. In practice, it also increases the number of drives that must be carefully phase tuned in an experiment.

In order to avoid these complications, we simulate the CX gate with the following steps. First, the master equation with the Hamiltonian in Eq. (2) and dissipative terms  $\mathcal{L}[\rho] = \sum_{i=c,t} \kappa_1(1 + n_{\text{th}})\mathcal{D}[a_i]\rho + \kappa_1 n_{\text{th}}\mathcal{D}[a_i^\dagger]\rho$  is simulated with  $\phi(t) = \pi t/T$  for time  $T$ . Next, the interactions between the control and target Kerr cats are removed and the Hamiltonian of the system is set to that of two uncoupled Kerr cats. Finally, the master equation with the two uncoupled Kerr cats is simulated with time-independent two-photon dissipation,  $\sum_{i=c,t} \kappa_2\mathcal{D}[a_i^2]\rho$ , in addition to  $\mathcal{L}[\rho]$  for time  $T$  [37].

Of course, more phase-flip errors are introduced in the second step of the simulation. The resulting error probabilities in this two-step CX gate, which takes time  $T_{\text{CX}} = 2T$ , will be roughly twice as high as that in the CX gate simulated with time-dependent two-photon dissipation in one step, which takes time  $T_{\text{CX}} = T$ .

To estimate the measurement errors, we perform a master-equation simulation of the  $X(\pm\pi/4)$  and  $S^{(\dagger)}$  gates. Subsequent steps of CD and finite-efficiency homodyne measurements are simulated by stochastic master-equation simulations. We assume a homodyne-measurement efficiency of 60% [69]. We find quantitative agreement between the analytic expression for homodyne-measurement fidelity in Ref. [70] and numerical results

from stochastic master-equation simulations. For an  $\alpha^2 = 6.25$  cat, we use the beam-splitter Hamiltonian  $H_{\text{readout}} = H_{\text{cat}} + g(a_1^\dagger a_r + a_r^\dagger a_1)$ , while for  $\alpha^2 = 1$ , we use  $H_{\text{readout}} = H_{\text{cat}} + g(a_1^\dagger + a_1)(a_r^\dagger + a_r)/2$ . Note that if the probability of bit-flip errors during the CD gate is small (a condition that is satisfied in the biased-noise cat qubit), this step in the measurement protocol is equivalent to the well-studied longitudinal readout scheme [70]. Consequently, known analytic expressions for longitudinal readout error may be used.

Observe from Fig. 5 that all gate operations [other than  $X(\pi/4)$ ] become faster with increasing cat size. For a given  $\alpha$ , we choose the gate speeds so that the probability of leakage outside the cat manifold is at least an order of magnitude smaller than the total infidelity. This choice is justified as the goal, when designing pulses for experimental implementation of gates, is to not limit the gate fidelity by the amount of leakage. While no pulse shaping is being used in our simulations, such techniques combined with the derivative removal by adiabatic gate (DRAG) method can be used to further reduce leakage without sacrificing the gate fidelity [71,72]. In our simulations, the leakage probability per qubit, per stabilizer round, is between  $10^{-3}$  and  $10^{-4}$ , which is consistent with recent experiments [73]. Additionally, the two-photon dissipation channel applied during or after the gates also checks the leakage growth. For the cat sizes considered here, the two-photon dissipation after every CX gate reduces the leakage by a factor of approximately 10.

In this way, in our simulations of the qubit operations, we suppress leakage to a reasonable amount at the physical level but then neglect the residual leakage in the further surface-code simulations. While this is the dominant practice in the literature for studying the performance of a code, further investigation must be carried out to model, quantify, and counteract the effect of leakage in the Kerr-cat surface-code architecture. As in other platforms, further suppression of leakage may require leakage reduction units (LRUs). These units reduce leakage errors to errors in the computational basis [74,75]. In the case of Kerr cats, leakage in one qubit is unlikely to propagate to others during two-qubit gates. Moreover, leakage is primarily confined to the lower-energy subspace, so that the computational-basis errors induced after LRUs remain biased. Of course, the additional circuit complexity of the LRUs may effect thresholds. Assessing the performance of the surface code with LRUs optimized for Kerr cats is beyond the scope of the current work but will be the topic of future study.

We simulate the XZZX surface code using the noise channel obtained numerically as described above. Tables I and II list the average fidelities, bias, and leakage in different operations obtained from numerical simulations for cat qubits with  $\alpha^2 = 1.0$  and  $\alpha^2 = 6.25$ , respectively.

TABLE I. The main parameters characterizing the operations used in syndrome extraction for a one-photon Kerr-cat qubit. The average gate bias is defined as the ratio of the sum of the probabilities of the  $I \otimes Z$ ,  $Z \otimes I$ , and  $Z \otimes Z$  errors and the sum of the remaining two-qubit Pauli errors. The effective Pauli-noise channel of the preparation and measurement operations is naturally biased. The total error in preparation is simply the sum of the probability of a  $Z$  and a  $Y$  error. As for  $\alpha^2 = 6.25$ , the CX is carried out in two steps, so that  $T_{CX} = 2T$ , where  $T$  is the time of each step. The leakage in all the operations is approximately  $10^{-4} - 10^{-5}$ .

$\alpha^2 = 1.0, \kappa_2 = K/10, n_{\text{th}} = 8\%$ CX gate, $T_{CX} = 112.6/K$		
Single-photon loss rate	Infidelity, biased-CX	Average bias, biased-CX
$0.67 \times 10^{-4}$	0.02	16.3
$0.83 \times 10^{-4}$	0.024	17.2
$1.11 \times 10^{-4}$	0.031	16.4
$1.45 \times 10^{-4}$	0.041	16.5
$1.67 \times 10^{-4}$	0.046	16.8

CZ gate, $J \sim 0.013K, T_{CZ} \sim 30.8/K$		
Single-photon loss rate	Infidelity	Average bias
$0.67 \times 10^{-4}$	0.006	9.6
$0.83 \times 10^{-4}$	0.008	10.6
$1.11 \times 10^{-4}$	0.01	12.5
$1.45 \times 10^{-4}$	0.013	13.2
$1.67 \times 10^{-4}$	0.014	14.5

Preparation, $\varepsilon = 0.026K, T_P \sim 16.5/K$	
Single-photon loss rate	Infidelity
$0.67 \times 10^{-4}$	0.0015
$0.83 \times 10^{-4}$	0.0019
$1.11 \times 10^{-4}$	0.0027
$1.45 \times 10^{-4}$	0.0032
$1.67 \times 10^{-4}$	0.004

Measurement, $\kappa_r = 0.1K, g \sim 0.06K$ $T = T_P + T_{CD}, T_{CD} = 42/K$ Measurement efficiency= 60%	
Single-photon loss rate	Total error
$0.67 \times 10^{-4}$	0.017
$0.83 \times 10^{-4}$	0.017
$1.11 \times 10^{-4}$	0.018
$1.45 \times 10^{-4}$	0.019
$1.67 \times 10^{-4}$	0.020

TABLE II. The main parameters characterizing the operations used in syndrome extraction for a 6.25-photon Kerr-cat qubit. The average gate bias is defined as the ratio of the sum of the probabilities of the  $I \otimes Z$ ,  $Z \otimes I$ , and  $Z \otimes Z$  errors and the sum of the remaining two-qubit Pauli errors. The leakage in all the operations is approximately  $10^{-4} - 10^{-5}$ . The effective Pauli-noise channel of the preparation and measurement operations is naturally biased. The total error in preparation is simply the sum of the probability of a  $Z$  and a  $Y$  error. The bias-preserving CX-gate simulation is carried out in two steps: first, the Hamiltonian in Eq. (2) with the Lindbladian  $\mathcal{L}[\rho] = \sum_{i=c,t} \kappa_1(1 + n_{\text{th}})\mathcal{D}[a_i]\rho + \kappa_1 n_{\text{th}}\mathcal{D}[a_i^\dagger]\rho$  is simulated for time  $T$ ; and in the second step, the master equation of two uncoupled Kerr cats with Lindbladian  $\sum_{i=c,t} \kappa_2\mathcal{D}[a_i^2]\rho + \mathcal{L}[\rho]$  is simulated for time  $T$ , so that  $T_{CX} = 2T$ . For fair comparison, a two-photon dissipation channel after the  $CX_R = RCZR^\dagger$  gate is also applied, for a time of approximately  $9/K$ , so that the residual leakage values after the CX and  $CX_R$  gates are the same. The time of the  $R$  gate is  $2.6/K$ ,  $CZ = 0.8/K$ , giving a total time of  $15/K$ .

$\alpha^2 = 6.25, \kappa_2 = K/10, n_{\text{th}} = 8\%$ CX gate, $T_{CX} = 18.47/K$ (bias-preserving) $T_{CX} = 15.0/K$ ( $R CZ R^\dagger$ )		
Single-photon loss rate	Infidelity, biased-CX ( $R CZ R^\dagger$ )	Average bias, biased-CX ( $R CZ R^\dagger$ )
$1.67 \times 10^{-4}$	0.047 (0.035)	319 (11.7)
$2.08 \times 10^{-4}$	0.054 (0.043)	337 (12.0)
$2.77 \times 10^{-4}$	0.071 (0.057)	357 (12.4)
$4.16 \times 10^{-4}$	0.104 (0.084)	375 (12.7)

CZ gate, $J \sim 0.08K, T_{CZ} \sim 0.8/K$		
Single-photon loss rate	Infidelity	Average bias
$1.67 \times 10^{-4}$	0.0029	$1.7 \times 10^3$
$2.08 \times 10^{-4}$	0.0036	$1.9 \times 10^3$
$2.77 \times 10^{-4}$	0.0039	$2.3 \times 10^3$
$4.16 \times 10^{-4}$	0.0056	$3.1 \times 10^3$

Preparation, $\varepsilon \sim 0.15K, T_P = 2.6/K$	
Single-photon loss rate	Infidelity
$1.67 \times 10^{-4}$	0.003
$2.08 \times 10^{-4}$	0.004
$2.77 \times 10^{-4}$	0.005
$4.16 \times 10^{-4}$	0.007

Measurement, $\kappa_r = 1.0K, g \sim 0.5K$ $T_M = T_P + T_{CD}, T_{CD} = 2.5/K$ Measurement efficiency= 60%	
Single-photon loss rate	Infidelity
$1.67 \times 10^{-4}$	0.008
$2.08 \times 10^{-4}$	0.009
$2.77 \times 10^{-4}$	0.011
$4.16 \times 10^{-4}$	0.012

Figure 8 shows the logical failure rate of the XZZX surface code as a function of the ratio of the single-photon loss rate and Kerr nonlinearity  $\kappa/K$ . The total gate error scales as  $\propto \kappa/K$  and so this is the important metric for experimentalists aiming to design the system.

The threshold values obtained for  $\kappa/K$  are given by  $(\kappa/K)_{\text{thres}} \sim 1.1 \times 10^{-4}$  for  $\alpha^2 = 1.0$  and  $(\kappa/K)_{\text{thres}} \sim$

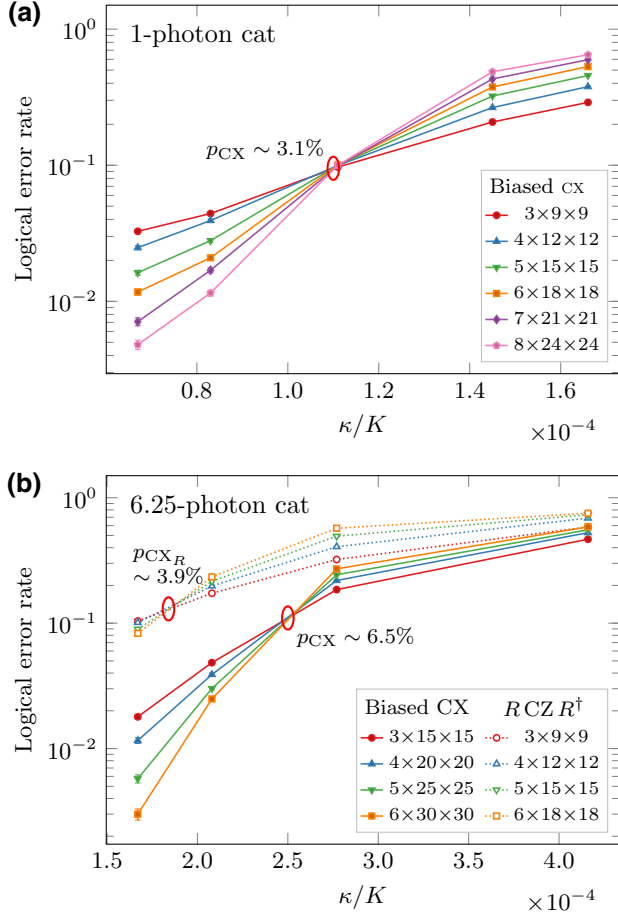


FIG. 8. The error thresholds of the XZZX surface code with the circuit-noise model for Kerr-cat qubits with a cat size (a)  $\alpha^2 = 1$  photon and (b)  $\alpha^2 = 6.25$  photons. The bottom axis shows the ratio of the single-photon loss rate  $\kappa$  or, equivalently, the inverse  $T_1$  time of the SNAIL used to realize the cat and the Kerr-nonlinear constant. The simulations also take into account a constant 8% thermal noise and two-photon dissipation  $\kappa_2 = K/10$  added for autonomous leakage correction. The noise channel in (b) is significantly more biased than in (a). We find the threshold  $\kappa/K$  of  $1.1 \times 10^{-4}$  and  $2.5 \times 10^{-4}$  in (a) and (b), respectively, which corresponds to a CX error threshold of approximately 3.1% and approximately 6.5%, respectively. For comparison, the threshold with an  $\alpha^2 = 6.25$  photon cat and a non-bias-preserving gate  $CX_R = RCZR^\dagger$  with  $R = \exp(i\pi Z_t/4) \exp(i\pi X_t/4)$  is also shown in (b). Even though the gate infidelity of  $CX_R$  is lower than that of the bias-preserving CX in Eq. (2), the threshold is lower by approximately  $1.84 \times 10^{-4}$  because the bias is smaller.

$2.5 \times 10^{-4}$  for  $\alpha^2 = 6.25$ . To illustrate the physical implication of these thresholds, consider a SNAIL-based Kerr cat with a Kerr nonlinearity of  $K/2\pi = 10$  MHz. Then  $(\kappa/K)_{\text{thres}} \sim 1.1 \times 10^{-4}$  for  $\alpha^2 = 1.0$  implies that the lifetime of the SNAIL ( $1/\kappa$ ) should be longer than  $144.7 \mu\text{s}$  for error correction with the XZZX surface code to be successful. For  $\alpha^2 = 6.25$ ,  $(\kappa/K)_{\text{thres}} \sim 2.5 \times 10^{-4}$  implies that the lifetime of the SNAIL ( $1/\kappa$ ) should be longer than  $63.6 \mu\text{s}$  for error correction with the XZZX surface code to be successful.

The most noisy operation in the system is the CX gate and we see that, with the  $\alpha^2 = 1.0$  cat, the CX-gate infidelity must be lower than 3.1%. On the other hand, with the  $\alpha^2 = 6.25$  cat, the CX-gate infidelity must be lower than approximately 6.5%. The average CX-gate bias at threshold for the  $\alpha^2 = 6.25$  cat is approximately 351, while that for  $\alpha^2 = 1$  cat is approximately 16. Clearly, as expected, we find that the XZZX surface code is able to tolerate more errors in the gates when the cat size is large because the noise bias increases.

The SNAIL lifetime in the first experimental demonstration of the Kerr cat was approximately  $15.5 \mu\text{s}$ , which shows that an improvement in the lifetime is required to operate below threshold with Kerr-cat qubits. Lifetimes of 50–100  $\mu\text{s}$  are routinely achieved in superconducting Josephson-junction circuits [4,56,57,62], so that reaching a lifetime above the  $T_1$  threshold appears very reasonable using state-of-the-art superconducting-circuit devices and fabrication. Optimizing the gate implementations and incorporating a one-step CX will further improve the threshold requirements. Finally, we note that the first Kerr-cat experiment [11] showed a strong suppression of frequency fluctuations due to  $1/f$  noise for the pumped cat, compared to when operating the SNAILmon as a conventional transmon qubit without pumping. In practice, frequency fluctuations of the SNAIL are near negligible when pumped to the cat subspace. This suggests that, although the Kerr cat is a flux-tunable device, recent improvements in coherence times for fixed-frequency qubits may also carry over to Kerr-cat qubits [76].

### 1. Comparison with non-bias-preserving CX gate

Note that with Kerr-cat qubits, it is possible to implement a non-bias-preserving CX gate by implementing the gate sequence  $CX_R = RCZR^\dagger$  with  $R = \exp(i\pi Z_t/4) \exp(i\pi X_t/4)$ . The subscript  $t$  denotes that the single-qubit gate  $R$  acts on the target of the CZ. We have already discussed how  $\exp(i\pi Z_t/4) = S^\dagger$ ,  $\exp(i\pi X_t/4) = X(\pi/4)$ , and CZ gates are simulated. Unlike the bias-preserving CX-gate implementation given in Eq. (2),  $CX_R$  does not preserve the bias in noise. However, the gate can be implemented faster, resulting in higher total fidelity. Table II also lists the gate fidelity and bias for the  $CX_R$  gate for the  $\alpha^2 = 6.25$  photon cat. For fair comparison, a

two-photon dissipation channel after the  $CX_R$  gate is also applied, so that the residual leakage after the  $CX$  and  $CX_R$  gates are the same for a given  $\kappa_1$ ,  $n_{\text{th}}$ , and  $\kappa_2$ .

To compare the two different gates, we repeat the simulation of the XZZX surface code but this time using the noise channel of the  $CX_R$  gate. As shown in Fig. 8(b), the threshold in this case is lowered by approximately 40% compared to the threshold with the bias-preserving  $CX$ . Moreover, the below-threshold logical error rate for the XZZX code is significantly higher with the  $CX_R$  gate than the bias-preserving  $CX$  gate, despite the  $CX_R$  gate having a lower total infidelity.

## 2. Comparison with the CSS surface code

We find that the CSS surface code with Kerr-cat noise in this parameter region does not perform as well as the XZZX code. At  $\alpha = 2.5$ , over the range of  $\kappa/K$  shown in Table II, we observe a logical error rate of close to 50%, which does not decrease with system size, indicating that the CSS code is far above threshold. At lower bias,  $\alpha = 1$ , we also do not observe a threshold in the overall error rate over the range of  $\kappa/K$  used in our XZZX-code simulations.

We also consider the scenario of Ref. [31], where the width of the CSS code is fixed  $d_x = 3$  and only the length of the code  $d_z$  is scaled up. The same above-threshold behaviour is observed for  $\alpha = 2.5$ , even when  $X$  logical errors are completely neglected. For  $\alpha = 1$ , if we neglect  $X$  logical errors, we observe a threshold in the  $Z$  logical error rate comparable to the threshold in the overall error rate of the XZZX code at  $\kappa/K = 1.1 \times 10^{-4}$ . Given, however, that  $X$  errors are non-negligible at this bias, the overall failure rate still appears above threshold and so error correction with the CSS surface code does not appear to be functional in the parameter regions considered.

The high thresholds that we obtain with the XZZX code are due to the fact that it appears very well suited to Kerr-cat qubit noise. One of the major advantages of this pairing is that a full round of syndrome measurements can be performed in a substantially shorter time than with the standard CSS surface code and therefore it is far less prone to error. This is because, for large  $\alpha$ , the gate that takes the longest to execute is the  $CX$  gate, while the relative time required to perform a  $CZ$  gate compared to  $CX$  tends to zero as  $\alpha$  increases. For the CSS surface code, readout of an  $X$  check requires four data qubits to be coupled to the ancilla via  $CX$  gates, which takes at least time  $4T_{CX}$ . In contrast, every check in the XZZX code requires two  $CX$  gates and two  $CZ$  gates, and given that these can be parallelized, the entangling operations for a full round of syndrome measurements can be performed in time  $2(T_{CX} + T_{CZ})$ . In the limit of large  $\alpha$ , the total time taken to apply the gates for a round of syndrome measurements with the XZZX code is approximately  $2T_{CX}$ , which is roughly equal to that of a

one-dimensional (1D) repetition code and half that of the CSS surface code.

## V. SMALL CODES

Here, we perform exact simulations of different quantum error-correcting codes for systems of small size undergoing biased noise. With reference to the previous section, we are able to analyze the scalability of different quantum error-correcting codes by simulating systems of a large number of qubits. We obtain these results using an efficient decoder and by approximating quantum noise channels as a Pauli-noise channel to make simulations tractable. Exact simulations enable us to investigate other details of various codes that are specialized to biased noise.

We perform several analyses using small-scale simulations using the Kerr-cat qubit noise model described in Sec. III B and an optimal decoding algorithm. By focusing on small codes, we can perform optimal decoding without concern about the efficiency of the decoder. We describe the optimal decoder and the methods we use to conduct exact simulations in Appendix A.

The smallest realization of a surface code that can correct at least one  $X$  error and one  $Z$  error has  $n = 9$  data qubits arranged on a  $3 \times 3$  lattice and uses the ‘‘rotated’’ layout of Ref. [68]. In our small-scale simulations, we compare the XZZX code on this layout to the TSC [42] on the same layout and to the 1D repetition code with  $n = 9$  qubits. The TSC is an alternative surface-code variant where we measure weight-four Pauli- $X$  stabilizer checks and weight-four Pauli- $Y$  stabilizer checks. This code has been demonstrated to have favorable properties to correct for biased noise [42,45]. We use open boundary conditions for all simulations of the  $n = 9$  codes with two-qubit checks on the boundaries of the surface-code lattices.

While, in Sec. IV, we show that the XZZX code has impressive fault-tolerant thresholds with finite noise biases, the TSC with open boundary conditions has the advantage that it can tolerate up to  $(n - 1)/2$  dephasing errors. In contrast, a logical failure can be introduced to the  $3 \times 3$  XZZX code with as few as two Pauli- $Z$  errors [41].

Like the TSC, the  $n = 9$  repetition code can also correct up to four dephasing errors with certainty. Moreover, it has the additional advantage that it requires only two entangling gates, rather than four, in the circuits used for stabilizer parity measurements. As such, the stabilizer readout circuits for the repetition code are easier to implement and less error prone than surface-code stabilizers. We might therefore expect that the repetition code is more accessible to realization using near-term experiments.

This convenience of the repetition code comes at the expense that it has no protection against bit-flip errors. Indeed, in Ref. [31], numerical simulations have shown that the repetition code does not increase its suppression

of logical failures for a realistic biased-noise model when increasing the system size beyond  $n \sim 10$  qubits. Nevertheless, for a small number of qubits and a very large noise bias such that the bit-flip errors are very rare, repetition codes may outperform surface codes. We can therefore expect variations in the performance of these codes in different parameter regimes.

Our exact simulation results highlight different features of the TSC, XZZX, and repetition code when using Kerr-cat qubits. The XZZX code in particular exhibits high performance in both high- and low-bias regimes, it is relatively insensitive to decoder miscalibration, and it does not suffer from significant increases in the logical error rate when standard CX gates are used instead of bias-preserving CX. In contrast, the repetition code will perform better when the bias is high relative to the system size, and the TSC may have an advantage in architectures where CX gates are not so noisy (compared to other operations).

In Fig. 9, we plot the logical error probabilities obtained using Kerr-cat qubits and bias-preserving operations alongside those obtained using non-bias-preserving CX and a miscalibrated decoder (which we describe in Sec. V D). The probabilities of the  $X$ ,  $Y$ , and  $Z$  logical errors separately are shown as well as the total logical error probability (which is the sum of the former three). The noise parameters are chosen to be slightly below the threshold estimated in Fig. 8, which is a regime relevant to near-term experiments.

Certain properties of these codes are reflected in these separate logical error probabilities. For instance, the  $Z$  logical is the only one that would be affected in the presence of pure phase-flip noise (i.e., noise where the bit-flip probability is exactly zero). With the rotated layout on a  $3 \times 3$  lattice, the weight of the smallest  $Z$  logical of the XZZX code is 3, while for the TSC and repetition codes it is 9 [42]. The  $Z$  logical is therefore much more protected than the  $X$  or  $Y$  logicals for the TSC and repetition codes compared to the XZZX code, as seen in Fig. 9. We see that the failures of the repetition code are almost exclusively due to bit-flip errors rather than phase-flip errors, since in Fig. 9 the  $Z$  logical errors are nearly undetectable.

We point out that, as in the Pauli simulations used to determine the threshold in Sec. V, we approximate the cat qubits in the exact surface-code simulations as strict two-level systems and neglect additional levels into which the system may leak. While the leakage suppression with engineered two-photon dissipation (and noise caused by this process) is still accounted for in the qubit noise model, there remains a small amount of residual leakage that is ignored. While we believe this effect to be small (for the CX gate, the probability of leakage is an order of magnitude smaller than that of a nondominant error), we leave a detailed study of this effect to future research.

Under the following subheadings, we discuss our small-scale analysis under several different criteria. We begin by

examining the break-even point of different codes. This is a pseudothreshold that tells us the noise parameters below which a small quantum error-correcting code can outperform the constituent parts from which it is built. These numbers give us an indication of the quality of the qubits that we need to demonstrate quantum error correction using small-scale experiments. In addition to this, we discuss the change in performance obtained by using bias-preserving gates as compared with a more conventional gate-noise model. We also use exact simulations to compare the true noise model extracted from master-equation simulations (but still neglecting leakage) to the Pauli-twirl approximation, which is used in Sec. IV B. Finally, we also use small-scale simulations to investigate the loss of performance due to a miscalibrated decoder.

### A. Break-even error rates

An important milestone for experimental demonstrations of error-correcting codes is reaching the pseudothreshold or break-even point. We say that a code is below the pseudothreshold if the probability of a logical error on the encoded qubit after a fixed number of rounds of error correction is less than the error probability of the noisiest element in the syndrome circuit, which in this case is the CX gate. The number of rounds is set to three for all of the following simulations.

As shown in Fig. 9(b), at high bias  $\alpha = 2.5$  and  $\kappa = 2.08 \times 10^{-4}$  with a bias-preserving CX and optimal decoding, the logical error rate for the XZZX code is  $5.41 \pm 0.09\%$  and for the TSC is  $5.53 \pm 0.10\%$ , which for both codes is equal, within statistical error, to the error rate of the CX gate of  $5.43\%$ . This shows that the pseudothreshold can be achieved with the XZZX and the TSC at remarkably high CX-gate infidelity. At this level of bias, where the probability of bit-flip errors is extremely low, the repetition code has an error rate of  $1.62 \pm 0.02\%$ , which is lower than both surface-code variants. This is due to its overall shorter circuit, having fewer qubits idle during CX-gate execution, and also its ability to correct up to four dephasing errors with certainty.

In contrast, at a lower bias of  $\alpha = 1$  and with  $\kappa = 8.33 \times 10^{-5}$  with a bias-preserving CX, the XZZX code is the only code of the three to be below the pseudothreshold. The advantage of the XZZX code over the repetition code here is that it can detect and correct bit-flip errors, which are non-negligible at this level of bias.

The drawback of the TSC in a Kerr-cat architecture comes from the fact that it requires twice as many CX gates compared to the XZZX code (which are far noisier than CZ gates). While we do not present the results here, we find that for simple biased-noise models where CZ and CX gates have roughly equal error probabilities, the  $3 \times 3$  TSC with optimal decoding can outperform the XZZX code in terms of logical error rate for a wide range of biases. However, as




FIG. 9. The logical error probabilities for low-distance ( $n = 9$ ) codes obtained by exact simulation. (a) A moderate bias, corresponding to  $\alpha = 1$ , is used with  $\kappa/K = 8.33 \times 10^{-5}$ . (b) A high bias, corresponding to  $\alpha = 2.5$ , is used with  $\kappa/K = 2.08 \times 10^{-4}$ . In both cases, the overall noise strength  $\kappa$  is chosen to be slightly below the threshold estimated in Fig. 8. The optimal performance (blue) is compared to the performance obtained with a miscalibrated decoder (red) and a non-bias-preserving CX gate (orange) for the XZZX surface code, the TSC, and the repetition code. Each data point is evaluated over 12 000 samples and the error bars represent the standard error of the mean. Aside from the red bars, which are obtained using a miscalibrated decoder, optimal decoding is used in all of these simulations.

efficient near-optimal decoding appears more complicated for the TSC than the XZZX code for larger systems, more work is still required before the TSC can be regarded as a practical or scalable alternative to the XZZX code under this kind of noise.

### B. The necessity of bias-preserving CX gates

One striking difference between the different codes is the relative importance of using bias-preserving CX gates. To test this, we compare the performance of each code using the bias-preserving CX gate to a non-bias-preserving CX gate (for a description of the latter CX gate, see Sec. IV A). We define this gate with the same infidelities as the bias-preserving CX gate.

The blue and orange bars in Fig. 9 represent the logical error probabilities obtained using bias-preserving CX gates and standard (non-bias-preserving) CX gates, respectively. While both the TSC and XZZX surface codes have a higher logical error rate when not using a bias-preserving CX gate, the increase in error rate is much larger for the TSC. At high bias,  $\alpha = 2.5$ , for the XZZX code, the error rate increases by around 70% when using non-bias-preserving CX, while the TSC sees a huge increase of nearly 400% compared to when using a bias-preserving CX gate. This is likely due to fact that the TSC uses twice

as many CX gates as the XZZX code and so using non-bias-preserving CX gates reduces the overall noise bias less in the XZZX code than the TSC. This effect is also observed, although it is less pronounced, at lower bias ( $\alpha = 1$ ). Like the threshold results of Sec. IV, these results suggest that high performance may be achieved with Kerr cats and the XZZX code even if bias-preserving gates are not used (although the best performance is likely obtained with bias-preserving gates).

Unlike the surface-code variants, the repetition code is entirely dependent on a bias-preserving CX gate. Since it is only designed to correct phase-flip errors, the additional bit-flip errors introduced with a non-bias-preserving CX gate drastically reduce its performance. This can be seen in both the  $\alpha = 1$  and the  $\alpha = 2.5$  plots, where the logical error rate is roughly equal to the probability of an  $X$  error occurring on a data qubit somewhere in the error-correction circuit.

### C. Interrogating the Pauli-twirl approximation

The threshold simulations in Sec. IV are performed with a Pauli-noise model, which allows efficient simulation of large systems. This Pauli-noise model is obtained by taking the Pauli-twirl approximation of the full noise model, which has non-Pauli or coherent terms. The Pauli-twirl

approximation effectively sets non-Pauli terms in the noise to zero. We do not need to apply the Pauli twirl to the noise in the exact simulations (although we still neglect leakage to other states). However, our exact simulation data suggest that non-Pauli terms in the physical noise make a negligible contribution to the performance of the surface-code variants.

The use of the Pauli twirl of the noise maps on the physical level in the exact simulation produces an average logical  $\chi$  matrix with diagonal terms (which can be interpreted as logical Pauli error probabilities) within statistical error of those obtained using the full (non-Pauli) noise maps on the physical level. Furthermore, off-diagonal terms in the logical  $\chi$  matrix are at least 2 orders of magnitude smaller than the diagonal terms for the averaged logical channel, suggesting that the averaged logical channel is close to a stochastic Pauli channel. This provides some justification for the use of the Pauli-twirl approximation in the threshold calculations in Sec. IV.

In contrast, the repetition code preserves coherence on the logical level. The full noise model results in slightly worse performance than its Pauli twirl, with about 10% lower fidelity in the logical channel for  $\alpha = 1$  and  $\kappa/K = 8.33 \times 10^{-5}$ . Furthermore, off-diagonal terms in the logical  $\chi$  matrix are larger than the dominant diagonal terms. This is not particularly surprising, since the  $X$  component of the noise is not transformed from the physical to the logical, e.g., a coherent  $X$  rotation on the first physical qubit of the repetition code results in an undetectable coherent  $X$  rotation by an equal angle of the logical qubit.

#### D. Miscalibrated decoding algorithms

An exact density-matrix simulator can also function as an optimal decoder. It can thus be used to probe the performance gains possible with improved decoding or, conversely, performance loss if the decoder is not optimally calibrated to the noise. To demonstrate the importance of calibrating the decoder to biased-gate noise, we compare the performance of a miscalibrated decoder to the optimal decoder. With the physical CX gate chosen to be bias preserving, the miscalibrated decoder is tuned to the noise model of the standard CX (rather than the bias-preserving CX). Thus the miscalibrated decoder is not properly adapted to the noise on the CX gate; however, it still correctly takes into account the bias on all other elements. The difference in performance observed between the miscalibrated decoder and optimal decoder tells us how important it is for the decoder to have an accurate characterization of bias on the CX gate.

As shown in Fig. 9, at high bias,  $\alpha = 2.5$ , the XZZX code is much less sensitive to decoder miscalibration than the TSC code at this bias. The use of a miscalibrated decoder results in an increase of less than 20% in the logical error rate for the XZZX code, compared to over 100%

for the TSC. The insensitivity of the XZZX code to miscalibrated decoding suggests that near-optimal performance may be achieved with this code using heuristic decoding methods (such as minimum-weight perfect matching) where a precise characterization of the noise model does not need to be used.

At lower bias,  $\alpha = 1$ , we see a smaller increase in the logical error rate when the decoder is miscalibrated for both the XZZX and the TSC surface codes. This suggests that the importance of tailored decoding increases as the bias increases. In contrast, the logical error rate of the repetition code, due to its relatively simple decoding requirements, appears almost completely unaffected by decoder miscalibration at both low and high bias.

## VI. DISCUSSION AND CONCLUSIONS

We propose a robust architecture for fault-tolerant quantum computing in which Kerr-cat qubits are used as the elementary qubits in an XZZX surface code. The resulting system demonstrates exceptional thresholds that can be met using a superconducting-circuit architecture with parameters that appear fairly modest compared to the current state of the art. We make conservative assumptions about the Kerr-cat photon number and imperfections such as thermal noise and measurement inefficiency. For the parameters used in our simulations, the Kerr cat has an anharmonicity, given by  $\omega_{\text{gap}}$ , comparable to that of a transmon, leading to fast gates. Note, however, that our simulations are performed without pulse optimization, which may yield significant improvements [77]. Further optimization is also possible by improving the decoding algorithm, as our results are obtained using a standard minimum-weight perfect-matching decoder [60].

It remains to make a detailed analysis of the resource cost to use our error-correcting system to perform a computation below some target logical failure rate. This will be important to compare the performance of our proposal with others [31,38–49]. Resource estimates are highly sensitive to the details of noise in laboratory systems [78] and can be improved dramatically with optimizations to the code (e.g., to the code layout and decoder) and hardware. Hence, we leave such calculations to future work. Nevertheless, given the remarkable threshold that our system exhibits, we expect we will be able to demonstrate a scalable system by realizing our proposal using modern technology.

Theoretical research should progress hand in hand with experimental development. To this end, details in our simulations have been motivated by the dominant sources of noise in Kerr-cat qubits that have been observed in recent experiments [11]. We should expect that as the system is scaled up in terms of both the size of the cats and the number of qubits on a chip, other sources of errors such as crosstalk, drive-induced heating, and leakage must be accounted for. It will be important to develop better



methods to characterize the noise in physical realizations of our error-correcting systems [79] and to refine our quantum error-correcting codes accordingly.

The architecture motivating our present study is based on the Kerr-cat qubit realized with on-chip superconducting SNAILs. This is a promising platform for realizing a large-scale surface-code architecture using systems presently available in the laboratory. The principles of the design and fabrication of a large-scale SNAIL chip are very similar to other circuit-QED-based approaches; for instance, quantum processors based on tunable transmons [80]. The scalability of our proposal with existing technologies suggests a practical solution to achieve fault-tolerant quantum computing in the near term.

### ACKNOWLEDGMENTS

We would like to thank S. Bartlett, P. Bonilla Ataides, and S. Flammia for conversations about decoders and tailored codes. S.P. would like to thank N. Frattini for numerous discussions, which were also helpful in the production of the conceptual figure for the Kerr-cat–SNAIL–surface code. A.S.D. was supported by the Japan Science and Technology Agency (JST), Precursory Research for Embryonic Science and Technology (PRESTO) Grant No. JPMJPR1917, Japan. B.J.B., A.L.G., and D.K.T. were supported by the Australian Research Council via the Centre of Excellence in Engineered Quantum Systems (EQUS) Project No. CE170100009 and by the Army Research Office (ARO) under Grant No. W911NF-21-1-0007. B.J.B. also received support from the University of Sydney Fellowship Programme, and A.L.G. is also supported by the Australian Research Council via a Discovery Early Career Research Award, Project No. DE190100380. S.P. is supported by the ARO under Grant No. W911NF-18-1-0212. Most of the numerical computation in this work was carried out at the Yukawa Institute Computer Facility. Access to high-performance computing resources was also provided by the National Computational Infrastructure (NCI), which is supported by the Australian Government.

### APPENDIX A: SURFACE-CODE SIMULATION METHOD

We use two types of numerical simulation to evaluate the performance of cat-surface codes. For large systems, in order to determine error thresholds, we use circuit-level Pauli simulations and for small codes more accessible to near-term implementation, we use exact state-vector simulations. The use of two simulations methods provides complementary insights into surface-code performance: the Pauli simulations allow system-size scaling analysis to be performed, while the exact simulations allow general types of error and to see the potential performance achievable with optimal decoding.

### 1. Pauli simulation

The Pauli simulations rely on the fact that the surface code is implemented with a Clifford circuit, which transforms Pauli operators to Pauli operators. One can therefore simulate a noisy circuit by sampling Pauli errors in the circuit (according to an error model of each circuit element) and then, by commuting the Pauli errors through to the end of the circuit, determine which ancilla measurements are flipped (i.e., the syndrome) and the residual error on the data qubits.

The Pauli simulation used for threshold calculation assumes that the initial state is a noiseless surface-code state, to which a number of rounds of noisy-syndrome measurements are applied, followed by a single final round of noiseless-syndrome measurements. The assumption that the final round of check measurements is noiseless is justified in practice, since the final readout of a surface-code qubit is performed by single-qubit measurements on the data qubits for which there is no distinction between a measurement error and a data qubit error. After sampling an error, the decoding algorithm is run on the corresponding syndrome to determine a correction. If the computed correction and the sampled error belong to the same logical equivalence class, we regard the error correction as successful. If not, a logical failure is said to have occurred. By repeating the simulation many times and counting the number of logical failures, the logical failure rate can be determined. To determine the error threshold, the simulation is performed at varying error rates as the code distance  $d$  is increased. For code distance  $d$ , the number of rounds of noisy measurements is set to  $d$ . The threshold is determined to be the noise strength below which increasing  $d$  leads to an exponential decrease in the logical failure rate.

### 2. Exact simulation

The exact simulation method involves storing the  $N$ -qubit density operator of the noisy code state,  $\rho$ , as a  $4^N$  component vector in memory and updating it as gates, noise, and measurements are applied with matrix-vector multiplication. Given the exponential scaling in required memory, the method is only suitable for small system sizes. However, it has the advantage of allowing arbitrary noise models to be studied (rather than restricting to Pauli noise). Keeping track of the full density operator also means that the simulation algorithm may be used for optimal decoding. This type of decoding is similar to the approach described in Ref. [81,82]. By calculating how a linearly independent set of encoded states (represented as  $2 \times 2$  matrices) transforms under error correction for a fixed syndrome  $s$ , it is possible to calculate the logical channel  $\mathcal{E}_s$ , represented as a  $4 \times 4$   $\chi$  matrix, conditioned on the syndrome. The optimal Pauli correction is then the one that minimizes the distance of the logical channel to

the identity, i.e.,

$$\operatorname{argmin}_{L \in \mathcal{L}, X, Y, Z} \|L \circ \mathcal{E}_s - I\|. \quad (\text{A1})$$

For Pauli noise, this is equivalent to maximum-likelihood decoding. One simplification made for the exact simulation is that the check measurements are performed sequentially, i.e., all the gates and the ancilla of a check are applied before measuring the next check. This is in contrast to the Pauli simulation, where all checks are measured simultaneously with an interleaved two-qubit gate pattern (described in Sec. II). This simplification to the exact simulation is made to save computational time, since performing check measurements in parallel necessitates keeping all ancilla qubits in memory (in addition to the data qubits), rather than a single ancilla qubit at a time. We remark that there is a simulation method, described in Ref. [83], based on novel tensor-network contraction techniques that can obviate the large memory cost of the simple exact simulation method that we use; however, the overall computational costs in this method still remain daunting.

There is some choice in where to put idle noise in the simulation when checks are measured sequentially. We apply idle noise to data qubits and ancilla qubits such that the total strength of the idle noise on each qubit is the same as if the gates were interleaved to measure the checks in parallel. This idle noise includes the noise on qubits left idle during measurement and preparation of ancilla as well as noise on qubits left idle while CX gates are applied to other qubits. Noise on qubits left idle during rounds of CZ gates is relatively small compared to CX gates and is neglected in our exact simulations (although it is included in our Pauli simulations). Thus the total amount of noise in the circuit is roughly equivalent, in our exact simulations, to the circuit with parallelized check readout that would be used in practice; however, due to the different gate order, there will be some differences in how errors propagate between data and ancilla qubits.

## APPENDIX B: KERR-CAT AND PURELY-DISSIPATIVE-CAT GATE FIDELITIES

In this appendix, with the help of a simple numerical example, we outline an important difference between Kerr cats and dissipative cats that leads to higher-fidelity gate operations with Kerr-cat qubits. Consider the implementation of the  $S$  gate in the Kerr-cat qubit. The required Hamiltonian is

$$\begin{aligned} H &= H_{\text{cat}} + \varepsilon(a^\dagger + a), \\ H_{\text{cat}} &= -K a^{\dagger 2} a^2 + K \alpha^2 (a^{\dagger 2} + a^2). \end{aligned} \quad (\text{B1})$$

In order to understand the effect of the microwave drive  $\varepsilon(a^\dagger + a)$ , let us look at how the photon annihilation and creation operators affect the computational cat states. The

action of the photon annihilation operator causes transitions between the cat states,  $a|C_\alpha^\pm\rangle = \pm\alpha r^{\pm 1}|C_\alpha^\mp\rangle$ . Here,  $r = \sqrt{(1 - e^{-2\alpha^2})/(1 + e^{-2\alpha^2})}$ . The creation operator, on the other hand, not only causes transitions between the cat states but it can also cause transitions outside the cat subspace. In fact, in the limit of large  $\alpha$ ,  $a^\dagger|C_\alpha^\pm\rangle \sim \pm\alpha|C_\alpha^\mp\rangle + |\psi_1^\mp\rangle$ , where  $|\psi_1^\mp\rangle$  are the first excited states with odd- and even-photon-number parity, respectively. These states are separated from the cat subspace by the energy gap  $|\omega_{\text{gap}}|$ . While the applied drive is resonant with the cat-qubit frequency, it is off resonant from the excited states by an amount  $|\omega_{\text{gap}}|$ . Consequently, this drive can only cause a small amount of population transfer to the leakage state. More importantly, the amount of leakage on average,  $\propto (\varepsilon/\omega_{\text{gap}})^2$ , decreases with  $|\omega_{\text{gap}}|$  [37].

An  $S$  gate with the purely dissipative cat requires the Hamiltonian  $H = i(\kappa_2/2)\alpha^2(a^{\dagger 2} - a^2) + \varepsilon(a^\dagger + a)$  and the engineered dissipation  $\kappa_2\mathcal{D}[a^2]\rho$ . The gate time is  $\pi/8\alpha\varepsilon$ , which is the same as the  $S$ -gate time with Kerr cats. In this case,  $a^\dagger$  also causes transition to the first excited manifold  $|C_\alpha^\pm\rangle \rightarrow |\psi_1^\mp\rangle$ . The excited-state population quickly decays back to the cat subspace at a rate of approximately  $2\kappa_2\alpha^2$ . The quantity of approximately  $2\kappa_2\alpha^2$  is also called the Lindbladian gap. The residual excited-state population is  $(\varepsilon/2\kappa_2\alpha^2)^2$ . Because of the transition-matrix elements  $\langle C_\alpha^\pm|a^2|\psi_1^\pm\rangle \sim 2\alpha$  and  $\langle C_\alpha^\pm|a^2|\psi_1^\mp\rangle = 0$ , the dissipation causes transitions within the even-parity and odd-parity subspaces, respectively. On the other hand, the leakage causes transitions between the even- and odd-parity subspace. Hence, cooling down of the excited-state population leads to additional phase-flip errors. These additional phase-flip errors in the purely dissipative cats are also termed as nonadiabatic errors in the literature [29,84]. The amount of phase flips,  $\propto \varepsilon/\kappa_2\alpha^3$ , increases with  $\varepsilon$  [84]. Fortunately, the additional phase-flip errors are noticeably absent in Kerr cats. Consequently, gates with Kerr cats have higher fidelity than those with purely dissipative cats.

We demonstrate the effect of additional phase-flip errors in purely dissipative cats by simulating the  $S$  gate and comparing its fidelity with that of the  $S$  gate with Kerr cats. Figure 10 shows this comparison. To keep the analysis simple, we neglect other sources of decoherence in either of the two cat qubits. In the simulations, we fix  $\varepsilon$  and  $\alpha = 2$  and hence the gate time  $\pi/8\alpha\varepsilon$  is the same for the two cats. We choose  $\kappa_2 = 2K$  so as to keep the energy and Lindbladian gaps (or, equivalently, the residual leakage) approximately the same in the two cats. Figure 10 clearly shows that the  $S$ -gate infidelity can be significantly higher for Kerr cats than for purely dissipative cats for the same amount of leakage. We also show the infidelity of the  $S$  gate in the hybrid approach with small two-photon dissipation  $\kappa_2 = K/10$  included in addition to the Hamiltonian Eq. (B1). The result shows that the addition of small




FIG. 10. The (a)  $S$ -gate infidelity and (b) leakage with Kerr cats (blue), purely dissipative cats (red), and Kerr cats with two-photon dissipation (orange),  $\kappa_2 = K/10$ . In (b), the gray dashed line indicates the approximate analytic expression for average residual leakage  $(\epsilon/4K\alpha^2)^2, (\epsilon/2\kappa_2\alpha^2)^2$ . In these simulations, we take  $\alpha = 2$  and all other sources of decoherence, such as single-photon loss and thermal excitations, are neglected. Note that the oscillations in the leakage for Kerr cats results from off-resonant population transfer back and forth between the cat and leakage space due to the Hamiltonian dynamics. These oscillations disappear when a smoother time-dependent pulse  $\epsilon(t)$  is used for implementing the  $S$  gate. Nonetheless, the average behavior agrees well with the simple analytic estimate given by the gray line.

amount of two-photon dissipation during the Kerr-cat  $S$  gate does not degrade the gate fidelity in the parameter regime used in this paper.

[1] 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).

[2] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, and B. Campbell, *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature* **508**, 500 (2014).

[3] D. Nigg, M. M. 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).

[4] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, and Y. Chen, *et al.*, State preservation by repetitive error detection in a superconducting quantum circuit, *Nature* **519**, 66 (2015).

[5] 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, *Nat. Commun.* **6**, 6979 (2015).

[6] M. Takita, A. D. Córcoles, E. Magesan, B. Abdo, M. Brink, A. W. Cross, J. M. Chow, and J. M. Gambetta, Demonstration of Weight-Four Parity Measurements in the Surface Code Architecture, *Phys. Rev. Lett.* **117**, 210505 (2016).

[7] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, and L. Jiang, *et al.*, Extending the lifetime of a quantum bit with error correction in superconducting circuits, *Nature* **536**, 441 (2016).

[8] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, and C. L. Zou, *et al.*, Quantum error correction and universal gate set operation on a binomial bosonic logical qubit, *Nat. Phys.* **15**, 503 (2019).

[9] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. P. Home, Encoding a qubit in a trapped-ion mechanical oscillator, *Nature* **566**, 513 (2019).

[10] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, and R. J. Schoelkopf, *et al.*, Quantum error correction of a qubit encoded in grid states of an oscillator, *Nature* **584**, 368 (2020).

[11] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, Stabilization and operation of a Kerr-cat qubit, *Nature* **584**, 205 (2020).

[12] R. Lescanne, M. Villiers, T. Peronin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas, Exponential suppression of bit-flips in a qubit encoded in an oscillator, *Nat. Phys.* **16**, 509 (2020).

[13] Y. Ma, Y. Xu, X. Mu, W. Cai, L. Hu, W. Wang, X. Pan, H. Wang, Y. P. Song, C. L. Zou, and L. Sun, Error-transparent operations on a logical qubit protected by quantum error correction, *Nat. Phys.* **16**, 827 (2020).

[14] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, and P. V. Klimov, *et al.*, Structures of rhodopsin in complex with G-protein-coupled receptor kinase 1, *Nature* **595**, 383 (2021).

[15] P. W. Shor, in *Proceedings of 37th Conference on Foundations of Computer Science*, FOCS '96 (IEEE Computer Society, USA, 1996), p. 56.

[16] A. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys.* **303**, 2 (2003).

- [17] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys. (NY)* **43**, 4452 (2002).
- [18] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
- [19] B. J. Brown, D. Loss, J. K. Pachos, C. N. Self, and J. R. Wootton, Quantum memories at finite temperature, *Rev. Mod. Phys.* **88**, 045005 (2016).
- [20] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature* **549**, 172 (2017).
- [21] D. Aharonov and M. Ben-Or, in *STOC '97 Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 1997), p. 176.
- [22] I. L. Chuang, D. W. Leung, and Y. Yamamoto, Bosonic quantum codes for amplitude damping, *Phys. Rev. A* **56**, 1114 (1997).
- [23] P. T. Cochrane, G. J. Milburn, and W. J. Munro, Macroscopically distinct quantum-superposition states as a bosonic code for amplitude damping, *Phys. Rev. A* **59**, 2631 (1999).
- [24] D. Gottesman, A. Kitaev, and J. Preskill, Encoding a qubit in an oscillator, *Phys. Rev. A* **64**, 012310 (2001).
- [25] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, Dynamically protected cat-qubits: A new paradigm for universal quantum computation, *New J. Phys.* **16**, 045014 (2014).
- [26] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, Protecting a bosonic qubit with autonomous quantum error correction, *Nature* **590**, 243 (2021).
- [27] C. Vuillot, H. Asasi, Y. Wang, L. P. Pryadko, and B. M. Terhal, Quantum error correction with the toric Gottesman-Kitaev-Preskill code, *Phys. Rev. A* **99**, 032344 (2019).
- [28] K. Noh and C. Chamberland, Fault-tolerant bosonic quantum error correction with the surface-Gottesman-Kitaev-Preskill code, *Phys. Rev. A* **101**, 012316 (2020).
- [29] J. Guillaud and M. Mirrahimi, Repetition Cat Qubits for Fault-Tolerant Quantum Computation, *Phys. Rev. X* **9**, 041053 (2019).
- [30] A. L. Grimsmo, J. Combes, and B. Q. Baragiola, Quantum Computing with Rotation-Symmetric Bosonic Codes, *Phys. Rev. X* **10**, 011058 (2020).
- [31] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, and A. Keller, *et al.*, Building a fault-tolerant quantum computer using concatenated cat codes, [arXiv:2012.04108 \[quant-ph\]](https://arxiv.org/abs/2012.04108) (2020).
- [32] K. Noh, C. Chamberland, and F. G. S. L. Brandao, Low overhead fault-tolerant quantum error correction with the surface-GKP code, [arXiv:2103.06994 \[quant-ph\]](https://arxiv.org/abs/2103.06994) (2021).
- [33] B. M. Terhal, J. Conrad, and C. Vuillot, Towards scalable bosonic quantum error correction, *Quantum Sci. Technol.* **5**, 043001 (2020).
- [34] A. Joshi, K. Noh, and Y. Y. Gao, Quantum information processing with bosonic qubits in circuit QED, *Quantum Sci. Technol.* **6**, 033001 (2021).
- [35] H. Goto, Universal quantum computation with a nonlinear oscillator network, *Phys. Rev. A* **93**, 050301 (2016).
- [36] S. Puri, S. Boutin, and A. Blais, Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving, *npj Quantum Inf.* **3**, 18 (2017).
- [37] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, and A. Blais, *et al.*, Bias-preserving gates with stabilized cat qubits, *Sci. Adv.* **6**, eaay5901 (2020).
- [38] P. Aliferis and J. Preskill, Fault-tolerant quantum computation against biased noise, *Phys. Rev. A* **78**, 052331 (2008).
- [39] P. Aliferis, F. Brito, D. P. DiVincenzo, J. Preskill, M. Steffen, and B. M. Terhal, Fault-tolerant computing with biased-noise superconducting qubits: A case study, *New J. Phys.* **11**, 013061 (2009).
- [40] A. M. Stephens, W. J. Munro, and K. Nemoto, High-threshold topological quantum error correction against biased noise, *Phys. Rev. A* **88**, 060301 (2013).
- [41] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultra-high Error Threshold for Surface Codes with Biased Noise, *Phys. Rev. Lett.* **120**, 050505 (2018).
- [42] D. K. Tuckett, A. S. Darmawan, C. T. Chubb, S. Bravyi, S. D. Bartlett, and S. T. Flammia, Tailoring Surface Codes for Highly Biased Noise, *Phys. Rev. X* **9**, 041031 (2019).
- [43] X. Xu, Q. Zhao, X. Yuan, and S. C. Benjamin, High-Threshold Code for Modular Hardware With Asymmetric Noise, *Phys. Rev. Appl.* **12**, 064006 (2019).
- [44] M. Li, D. Miller, M. Newman, Y. Wu, and K. R. Brown, 2D Compass Codes, *Phys. Rev. X* **9**, 021041 (2019).
- [45] D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, Fault-Tolerant Thresholds for the Surface Code in Excess of 5% Under Biased Noise, *Phys. Rev. Lett.* **124**, 130501 (2020).
- [46] S. Huang, M. Newman, and K. R. Brown, Fault-tolerant weighted union-find decoding on the toric code, *Phys. Rev. A* **102**, 012419 (2020).
- [47] J. Guillaud and M. Mirrahimi, Error rates and resource overheads of repetition cat qubits, *Phys. Rev. A* **103**, 042413 (2021).
- [48] L. Hänggeli, M. Heinze, and R. König, Enhanced noise resilience of the surface-GKP code via designed bias, *Phys. Rev. A* **102**, 052408 (2020).
- [49] O. Higgott and N. P. Breuckmann, Subsystem codes with high thresholds by gauge fixing and reduced qubit overhead, [arXiv:2010.09626 \[quant-ph\]](https://arxiv.org/abs/2010.09626) (2020).
- [50] X.-G. Wen, Quantum Orders in an Exact Soluble Model, *Phys. Rev. Lett.* **90**, 016803 (2003).
- [51] J. P. Bonilla Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, The XZZX surface code, *Nat. Commun.* **11**, 2172 (2021).
- [52] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, [arXiv:quant-ph/9811052 \[quant-ph\]](https://arxiv.org/abs/quant-ph/9811052) (1998).
- [53] J. Edmonds, Paths, trees, and flowers, *Can. J. Math.* **17**, 449 (1965).
- [54] V. Kolmogorov, BLOSSOM V: A new implementation of a minimum cost perfect matching algorithm, *Math. Prog. Comput.* **1**, 43 (2009).
- [55] B. J. Brown and D. J. Williamson, Parallelized quantum error correction with fracton topological codes, *Phys. Rev. Res.* **2**, 013303 (2020).

- [56] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. Borgorin, M. Brink, L. Capelluto, O. Gunluk, T. Itoko, and N. Kanazawa, *et al.*, Demonstration of quantum volume 64 on a superconducting quantum computing system, *Quantum Sci. Technol.* **6**, 025020 (2021).
- [57] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [58] Z. Nussinov and G. Ortiz, A symmetry principle for topological quantum order, *Ann. Phys.* **324**, 977 (2009).
- [59] B. J. Brown, W. Son, C. V. Kraus, R. Fazio, and V. Vedral, Generating topological order from a two-dimensional cluster state using a duality mapping, *New J. Phys.* **13**, 065010 (2011).
- [60] 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).
- [61] Y. Tomita and K. M. Svore, Low-distance surface codes under realistic quantum noise, *Phys. Rev. A* **90**, 062320 (2014).
- [62] E. A. Sete, N. Didier, A. Q. Chen, S. Kulshreshtha, R. Manenti, and S. Poletto, Parametric-resonance entangling gates with a tunable coupler, [arXiv:2104.03511 \[quant-ph\]](https://arxiv.org/abs/2104.03511) (2021).
- [63] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, and M. J. Hatridge, *et al.*, Confining the state of light to a quantum manifold by engineered two-photon loss, *Science* **347**, 853 (2015).
- [64] S. Touzard, A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, and K. M. Sliwa, *et al.*, Coherent Oscillations Inside a Quantum Manifold Stabilized by Dissipation, *Phys. Rev. X* **8**, 021005 (2018).
- [65] G. J. Woeginger, Open problems around exact algorithms, *Discret. Appl. Math.* **156**, 397 (2008).
- [66] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting, *Phys. Rev. A* **74**, 042318 (2006).
- [67] 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, *Phys. Rev. X* **9**, 041009 (2019).
- [68] H. Bombin and M. A. Martin-Delgado, Optimal resources for topological two-dimensional stabilizer codes: Comparative study, *Phys. Rev. A* **76**, 012305 (2007).
- [69] S. Touzard, A. Kou, N. Frattini, V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. Devoret, Gated Conditional Displacement Readout of Superconducting Qubits, *Phys. Rev. Lett.* **122**, 080502 (2019).
- [70] N. Didier, J. Bourassa, and A. Blais, Fast Quantum Non-demolition Readout by Parametric Modulation of Longitudinal Qubit-Oscillator Interaction, *Phys. Rev. Lett.* **115**, 203601 (2015).
- [71] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 14 (2021).
- [72] Z. Chen, J. Kelly, C. Quintana, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, and E. Lucero, *et al.*, Measuring and Suppressing Quantum State Leakage in a Superconducting Qubit, *Phys. Rev. Lett.* **116**, 020501 (2016).
- [73] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. Satzinger, C. Quintana, P. Klimov, D. Sank, C. Gidney, and A. Fowler, *et al.*, Removing leakage-induced correlated errors in superconducting quantum error correction, *Nat. Commun.* **12**, 1761 (2021).
- [74] P. Aliferis and B. M. Terhal, Fault-tolerant quantum computation for local leakage faults, *Quant. Inf. Comput.* **7**, 139 (2007).
- [75] M. Suchara, A. W. Cross, and J. M. Gambetta, Leakage suppression in the toric code, *Quant. Inf. Comput.* **15**, 997 (2015).
- [76] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, and S. Sussman, *et al.*, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nat. Commun.* **12**, 1779 (2021).
- [77] F. Motzoi, J. M. Gambetta, P. Reberstrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [78] M. E. Beverland, B. J. Brown, M. J. Kastoryano, and Q. Marolleau, The role of entropy in topological quantum error correction, *J. Stat. Mech.* **2019**, 073404 (2019).
- [79] R. Harper, S. T. Flammia, and J. J. Wallman, Efficient learning of quantum noise, *Nat. Phys.* **16**, 1184 (2020).
- [80] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, and D. A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [81] A. S. Darmawan and D. Poulin, Tensor-Network Simulations of the Surface Code under Realistic Noise, *Phys. Rev. Lett.* **119**, 040502 (2017).
- [82] A. S. Darmawan and D. Poulin, Linear-time general decoding algorithm for the surface code, *Phys. Rev. E* **97**, 051302 (2018).
- [83] C. Huang, X. Ni, F. Zhang, M. Newman, D. Ding, X. Gao, T. Wang, H.-H. Zhao, F. Wu, and G. Zhang, *et al.*, Alibaba cloud quantum development platform: Surface code simulations with crosstalk, [arXiv:2002.08918 \[quant-ph\]](https://arxiv.org/abs/2002.08918) (2020).
- [84] C. Chamberland, A. Kubica, T. J. Yoder, and G. Zhu, Triangular color codes on trivalent graphs with flag qubits, *New J. Phys.* **22**, 023019 (2020).