

Hybrid Cat-Transmon Architecture for Scalable, Hardware-Efficient Quantum Error Correction

Connor T. Hann^{1,*}, Kyungjoo Noh¹, Harald Putterman¹, Matthew H. Matheny¹, Joseph K. Iverson¹, Michael T. Fang¹, Christopher Chamberland^{1,†}, Oskar Painter^{1,2,3} and Fernando G.S.L. Brandão^{1,2,‡}

¹*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

²*IQIM, California Institute of Technology, Pasadena, California 91125, USA*

³*Thomas J. Watson, Sr., Laboratory of Applied Physics, California Institute of Technology, Pasadena, California 91125, USA*

 (Received 22 November 2024; revised 1 March 2025; accepted 22 May 2025; published 11 July 2025)

Dissipative cat qubits are a promising physical platform for quantum computing, since their large noise bias can enable more hardware-efficient quantum error correction. In this work we theoretically study the long-term prospects of a hybrid cat-transmon quantum computing architecture where dissipative cat qubits play the role of data qubits, and error syndromes are measured using ancillary transmon qubits. The cat qubits' noise bias enables more hardware-efficient quantum error correction, and the use of transmons allows for practical, high-fidelity syndrome measurement. While correction of the dominant cat Z errors with a repetition code has recently been demonstrated in experiment, here we show how the architecture can be scaled beyond a repetition code. In particular, we propose a cat-transmon entangling gate that enables the correction of residual cat X errors in a thin rectangular surface code, so that logical error can be arbitrarily suppressed by increasing code distance. We numerically estimate logical memory performance, finding significant overhead reductions in comparison to architectures without biased noise. For example, with current state-of-the-art coherence, physical error rates of 10^{-3} and noise biases in the range 10^3 – 10^4 are achievable. With this level of performance, the qubit overhead required to reach algorithmically relevant logical error rates with the cat-transmon architecture matches that of an unbiased-noise architecture with physical error rates in the range 10^{-5} – 10^{-4} .

DOI: [10.1103/75x7-5ysv](https://doi.org/10.1103/75x7-5ysv)

I. INTRODUCTION

Bosonic quantum error-correcting codes offer a promising path towards error correction with lower qubit overhead. In contrast to *qubit codes*—like the surface code [1–3]—that redundantly encode quantum information in the Hilbert space of multiple physical qubits, *bosonic codes* [4–9] encode information in the infinite-dimensional Hilbert space of a single quantum harmonic oscillator, such as an electromagnetic or mechanical resonator. The ability to already correct or suppress some errors with only one

physical component makes bosonic codes naturally hardware efficient. Residual errors can then be corrected by concatenating a bosonic code with an outer qubit code. In these *concatenated bosonic codes* [10–13], the reduced error of the bosonic qubits reduces the overhead to achieve a target logical error with the outer qubit code. A host of recent experiments have demonstrated the ability of various bosonic codes to correct or suppress errors, including cat codes [14–22], binomial codes [23,24], and Gottesman-Kitaev-Preskill codes [25–28]. Some experiments have even approached [24,25,29] or exceeded [27,30] the so-called “breakeven” point, where the lifetime of the bosonic qubit exceeds the bosonic mode’s coherence time.

Among bosonic codes, cat codes are especially promising candidates for hardware-efficient error correction because they can exhibit biased noise. In two-component cat codes, typical physical errors, such as photon loss and dephasing, result in a probability of X errors on the encoded bosonic qubit that is exponentially suppressed relative to the probability of Z errors [5]. The large ratio of the Z to X error probabilities in the encoded qubit, referred to

*Contact author: connohan@amazon.com

†Current affiliation: NVIDIA.

‡Contact author: fbrandao@amazon.com

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.

as a noise bias, can then be exploited by the outer qubit code in a concatenated bosonic code. For example, the presence of a noise bias can lead to increased error correction thresholds [31–34] and reduced overheads [13,35,36]. Biased-noise cat qubits have been realized in several recent experiments, with the implementations relying on either engineered Hamiltonians [18,19,21] or dissipation [14,16,17,20] to stabilize the code space and achieve a noise bias. Notably, Refs. [17,20,37] demonstrated that the cat’s noise bias can be increased exponentially by increasing the cat’s mean photon number.

In order for the outer qubit code to take full advantage of the cats’ noise bias, it is critical that this bias be preserved as gates are applied to the cat qubits. While proposals have been put forth for bias-preserving gates between cat qubits (referred to here as cat-cat gates) [35,38], these implementations can have onerous experimental requirements. For example, the infidelity of the proposal of Guillaud and Mirrahimi [35] for a cat-cat controlled- X (CX) gate scales as $1 - F \sim \sqrt{\kappa_1/\kappa_2}$, where κ_1 is the single-photon loss rate and κ_2 is the engineered dissipation rate of the cats [5]. This unfortunate square-root scaling necessitates both strong engineered dissipation and high coherence for good performance. Indeed, a recent experiment [22] achieved an impressive $\kappa_1/\kappa_2 \approx 1/150$, but κ_1/κ_2 orders of magnitude smaller still may be needed to operate a concatenated cat code at algorithmically relevant error rates [13,36,39,40]. Such high coherence together with strong engineered dissipation remains to be demonstrated experimentally.

In this work, we propose to circumvent the unfavorable error scaling of the bias-preserving cat-cat entangling gates—while still benefiting from the cats’ noise bias—by using transmon qubits, rather than cat qubits, as ancillae to measure error syndromes. In the resulting hybrid *cat-transmon architecture*, only cat-transmon entangling gates, rather than cat-cat gates, are required to perform error correction. The fidelities of these cat-transmon gates scale with κ_1/χ , where χ is the cat-transmon dispersive coupling strength. This scaling both enables higher fidelity gates at given κ_1 , and removes the need for strong engineered dissipation. Moreover, because the gates’ implementations rely on the natural qubit-oscillator dispersive interaction, there is no need for complicated Hamiltonian engineering, and fast gates are enabled by easily engineered megahertz-scale dispersive couplings. These practical benefits allow for fast, high-fidelity cat-transmon gates to be readily implemented, enabling near-term experimental demonstrations of concatenated cat codes, including the recent realization of a repetition cat code in Ref. [41].

The price we pay for these practical benefits is an upper limit on the achievable noise bias. This limit is due to the propagation of transmon ancilla errors into cat data qubit X errors via the cat-transmon gates. Steps are taken to mitigate this error propagation [42–44], such that no

single transmon loss or dephasing error can induce a cat X error on its own. We predict that significant biases, in the range of 10^3 – 10^4 , are achievable with current state-of-the-art parameters as a result. Even so, a way to correct the residual cat X errors is still necessary for the cat-transmon architecture to be fully scalable, i.e., to achieve arbitrarily low logical error. We propose a cat-transmon entangling gate, the CR_X gate, that enables this correction while enjoying all of the practical benefits mentioned above.

With this blueprint for a scalable architecture in hand, we proceed to quantify the hardware efficiency of the cat-transmon architecture relative to a generic unbiased-noise architecture. We find that significant reductions in logical memory overhead are possible despite the limited bias. For example, physical error rates of 10^{-3} are accessible with current state-of-the-art coherence, $\kappa_1/\chi \sim 10^{-5}$ – 10^{-4} . With this level of error, we predict that a thin rectangular surface code with logical memory error $< 10^{-10}$ could be built with only 200 qubits (100 cats, 100 transmons), thanks to the sizable noise bias (10^3 – 10^4) of the cat-transmon architecture. In contrast, with the same physical error but without noise bias, nearly 1000 qubits would be required to construct a comparably performant square surface code. Alternatively, without bias, physical error rates would need to be reduced to $< 10^{-4}$ to match the qubit overhead of the cat-transmon architecture. Of course, while the larger biases that cat-cat gates offer could yield even further improvement, we emphasize that the primary appeal of the cat-transmon architecture is its practicality; significant noise bias and hardware efficiency in error correction are achievable using only well-established experimental platforms and capabilities in superconducting circuits.

The rest of this manuscript is organized as follows. In Sec. II we provide a high-level overview of the cat-transmon architecture, focusing on how different cat-transmon gates enable the correction of both the dominant cat Z errors and the suppressed cat X errors. Next, in Sec. III we present a detailed description of the bias-preserving CR_X gate that enables correction of the suppressed cat X errors. We numerically benchmark the performance of these gates in Sec. IV, and then we use the results to make projections for the cat-transmon architecture’s performance as a logical memory in Sec. V. Finally, in Sec. VI we conclude by discussing several directions for future research to further improve the performance of the architecture.

II. HYBRID CAT-TRANSMON ARCHITECTURE

In this section, we provide an overview of the cat-transmon architecture (Fig. 1), which is an example of a hybrid oscillator-qubit processor [45]. We first discuss the




FIG. 1. Hybrid dissipative cat-transmon QEC architecture. (a) Cat-transmon unit cell. Dissipatively stabilized cats act as data qubits, and transmons are employed to extract error syndromes. The cats and transmons are coupled via tunable dispersive couplings. (b) Surface codes built from the cat-transmon unit cell. By constructing rectangular codes that offer more protection against the cats' dominant Z errors than their suppressed X errors, the cats' biased noise can allow for more hardware-efficient error correction. The surface code logical operators are shown in red and blue, and example stabilizers are shown in black with the order in which entangling gates are performed during syndrome measurement indicated by arrows. (c) Syndrome measurement circuits. To measure X -type syndromes, transmon-controlled CX gates are implemented using the native dispersive couplings. To measure Z -type syndromes, the architecture employs a proposed cat-controlled entangling gate, termed CR_X . To prevent leakage out of the cat-qubit code space, engineered dissipation is applied to the cats during the transmon ancilla readout and reset steps (indicated by boxes labeled “D”).

architecture's basic building blocks: the dissipatively stabilized cat data qubits and the dispersively coupled transmon ancilla qubits. Then, we describe how cat-transmon entangling gates enable error correction in an outer surface code.

A. Physical components

The data qubits in our architecture are encoded in two-component cat codes, which store information in superpositions of coherent states $|\pm\alpha\rangle$. The logical X -basis states of the cat code, $|\pm\rangle_C$, are defined as the even- and odd-parity superpositions of these coherent states,

$$|\pm\rangle_C = \mathcal{N}_\pm (|\alpha\rangle \pm |-\alpha\rangle), \quad (1)$$

where the $\mathcal{N}_\pm = 1/\sqrt{2(1 \pm e^{-2|\alpha|^2})}$ are normalization constants. The logical Z -basis states of the cat code, $|0/1\rangle_C$, are then given by the superpositions

$$\begin{aligned} |0/1\rangle_C &= \frac{1}{\sqrt{2}} (|+\rangle_C \pm |-\rangle_C) \\ &= |\pm\alpha\rangle + O(e^{-2|\alpha|^2}) |\mp\alpha\rangle \end{aligned} \quad (2)$$

with $|0/1\rangle_C$ approaching the coherent states $|\pm\alpha\rangle$ in the limit of large $|\alpha|^2$. The cat code logical Pauli operators are given by $X = |0\rangle\langle 1|_C + |1\rangle\langle 0|_C$ and $Z = |0\rangle\langle 0|_C - |1\rangle\langle 1|_C$. We refer to the bosonic modes that host cat qubits in our architecture as “storage modes.”

When the state of a storage mode is confined to the cat code space, the cat qubit exhibits biased noise. In this work, we focus on the paradigm of dissipative stabilization: confinement to the code space is engineered by a driven-dissipative process where the storage mode exchanges pairs of photons with its environment, which can be engineered, e.g., by nonlinearly coupling the mode

to a separate lossy mode [5,14,16,17]. The dynamics of the dissipatively stabilized system can be described by the Lindblad master equation

$$\dot{\rho}(t) = \kappa_2 \mathcal{D}[a^2 - \alpha^2] \rho(t), \quad (3)$$

where a is the storage mode's annihilation operator, κ_2 is the engineered dissipation rate, and $\mathcal{D}[L]\rho(t) = L\rho L^\dagger - (L^\dagger L\rho + \rho L^\dagger L)/2$. The steady-state manifold of these dynamics is the cat code space. Once confined to the code space, typical error mechanisms like single-photon loss ($L = \sqrt{\kappa_1}a$) and white-noise dephasing ($L = \sqrt{\kappa_\phi}a^\dagger a$) result in a noise bias that grows exponentially with $|\alpha|^2$. For example, single-photon loss will induce cat X errors at a rate $\propto e^{-c|\alpha|^2}$, where the exponential prefactor c depends on the parameter regime [46], and Z errors at a rate $\propto |\alpha|^2$. The exponential suppression of the X error rate with $|\alpha|^2$ relative to the linear increase in the Z error rate results in a tunable bias that can be deliberately made large, as demonstrated in recent experiments [17,20,22,47]. We note that the Kerr cat [7] also constitutes a promising paradigm for cat stabilization, as demonstrated in recent experiments [18,19,21]. However, here we restrict our attention to dissipative cats, since large Kerr nonlinearities are not easily compatible with cat-transmon gate implementations discussed below.

The ancillary qubits in our architecture are standard transmon qubits and are coupled dispersively to the storage modes. The joint system is modeled via the Hamiltonian

$$H = H_{\text{storage}} + H_{\text{transmon}} + H_{\text{disp}}. \quad (4)$$

Here, the bare Hamiltonians of the storage mode and transmon are given by

$$H_{\text{storage}} = \omega_a a^\dagger a - K a^\dagger a^\dagger a a, \quad (5)$$

$$H_{\text{transmon}} = \omega_q |e\rangle \langle e| + (2\omega_q - \nu) |f\rangle \langle f|, \quad (6)$$

where ω_a is the frequency of the storage mode and K its self-Kerr nonlinearity, and ω_q is the frequency of the transmon and ν its anharmonicity. We exploit the fact that a qubit can be encoded either in the subspace spanned by the transmon's ground and first excited states, $|g\rangle$ and $|e\rangle$, or, alternatively, in its ground and second excited states, $|g\rangle$ and $|f\rangle$. The dispersive coupling is given by the effective Hamiltonian [48]

$$H_{\text{disp.}} = \chi_e |e\rangle \langle e| a^\dagger a + \chi_f |f\rangle \langle f| a^\dagger a, \quad (7)$$

where $\chi_{e,f}$ denote the dispersive frequency shifts for the transmon $|e, f\rangle$ states. Additionally, resonant drives can be applied to the storage and transmon to actuate gates. These drives are described by the Hamiltonian

$$H_{\text{drive}} = \epsilon(t) e^{-i\omega_a t} a^\dagger + \text{H.c.} + \frac{\Omega(t)}{2} e^{-i\omega_q t} |e\rangle \langle g| + \text{H.c.}, \quad (8)$$

where, under the assumption of large anharmonicity $|\nu| \gg |\Omega|$, only the $|g\rangle$ and $|e\rangle$ states are included in the transmon drive, since only transitions between these two states will be resonantly driven during cat-transmon gates. See Appendix E for a discussion of this and other assumptions made in the modeling of the cat-transmon system.

Our implementations of cat-transmon gates place certain requirements on the device parameters. First, we require that both the engineered dissipation strength κ_2 and the dispersive couplings $\chi_{e,f}$ can be quickly toggled on and off, with $\chi_{e,f}^{(\text{on})} \gg \kappa_2^{(\text{off})}$ and $\kappa_2^{(\text{on})} \gg \chi_{e,f}^{(\text{off})}$. This requirement derives from the fact that stabilization must be turned off during gates—since, by design, the state of the storage mode temporarily leaves the cat code space—but should be turned on otherwise. The required tunability has been demonstrated in recent experiments, where, for example, κ_2 can be controlled via external flux drive [17], and $\chi_{e,f}$ can be controlled via a tunable coupler [37]. Second, when the transmon is operated in the $|g\rangle, |f\rangle$ manifold, we require $\chi_e \approx \chi_f$. This approximate “chi matching” realizes a form of error transparency [44,49] that enables one to achieve a large noise bias. Chi matching can be realized either actively via Hamiltonian engineering [42–44] or passively in design via a careful arrangement of underlying device parameters [41]. We emphasize that the chi matching here needs only be approximate, as the cat's dissipative stabilization can correct the effects of a small mismatch (i.e., $|\chi_e/\chi_f - 1| \lesssim 0.1$ [41]). This is in contrast to earlier experiments [42,43] that lacked dissipative stabilization and so required essentially perfect matching. Beyond these requirements, we do not otherwise specify implementation details of how the dissipative stabilization or dispersive interactions are engineered, so that the

results below are not restricted to a particular implementation. That said, we note that Ref. [41] constitutes a concrete example of an architecture that simultaneously satisfies all of these requirements. In that work, the engineered dissipation is implemented using an asymmetrically threaded superconducting quantum interference device [17], tunable dispersive couplings are realized using tunable transmon couplers, and chi matching is realized natively via a careful choice of the storage mode and transmon frequencies.

B. Gates and error correction

The basic unit cell of our architecture, the cat-transmon unit cell, is shown in Fig. 1(a). Surface codes can be obtained by tiling the cat-transmon unit cell, as shown in Fig. 1(b). Importantly, we employ rectangular codes to capitalize on the cats' biased noise. Rectangular codes are parameterized by two different code distances, d_X and d_Z , that separately characterize the degree of protection against X and Z errors, respectively. Because the cat qubits' X errors are already strongly suppressed, it suffices to use relatively small d_X , reducing the qubit overhead relative to a square code. We note that there are other ways to leverage noise bias to improve hardware efficiency (e.g., by employing Clifford-deformed surface codes like the XZZX [33] or XY [32] codes), but we focus on the standard (i.e., rotated, nondeformed) rectangular surface code because we find that it offers a larger hardware-efficiency improvement once errors are far below threshold (see Appendix I for analysis of the XZZX code and comparison).

Circuits for measuring the surface code stabilizers are shown in Fig. 1(c). The two types of stabilizer, X type and Z type, enable the correction of Z and X errors, respectively. Correcting both types of errors is crucial for full scalability, given the finite noise bias. Two distinct gates are required to measure the syndromes, both because we use different implementations for the data and ancilla qubits, and also because the use of operations that convert Z errors to X errors, like Hadamard gates, must be avoided on the cat qubits to maintain their noise bias. The X -type syndrome measurement requires a transmon-controlled X operation on the cat, while the Z -type measurement requires a cat-controlled X operation on the transmon.

To correct the cats' dominant Z errors, a transmon-controlled CX gate is implemented following the approach of Rosenblum *et al.* [42]. Entanglement between the cat and transmon is generated simply by letting the system evolve freely under the dispersive coupling, with the cat's engineered dissipation turned off. For a transmon encoded in the $|g\rangle, |e\rangle$ (respectively, $|g\rangle, |f\rangle$) manifold, evolving for a time π/χ_e (π/χ_f) results in the storage mode acquiring a π -phase shift conditioned on the transmon being in the excited state. This phase shift effectively

swaps $|+\alpha\rangle \leftrightarrow |-\alpha\rangle$, equivalent to a transmon-controlled CX gate.

If the transmon is operated in the $|g\rangle, |e\rangle$ manifold, the cat's noise bias is substantially compromised by this transmon-controlled CX gate, but if the transmon is operated instead in the $|g\rangle, |f\rangle$ manifold [42–44], a large noise bias can be retained. For $|g\rangle, |e\rangle$, the bias is compromised because single transmon errors can readily propagate to cat X errors. For example, if the transmon suffers a T_1 decay error, $|e\rangle \rightarrow |g\rangle$, immediately before the gate, the storage mode does not acquire the desired phase shift, amounting to a cat X error. This problem can be mitigated by instead operating the transmon in the $|g\rangle, |f\rangle$ manifold while employing chi matching $\chi_e \approx \chi_f$. Now, the resonator acquires the desired phase shift even if the transmon decays from $|f\rangle \rightarrow |e\rangle$ during the gate. This error transparency [44,49,50] means that single transmon decay errors do not propagate to cat X errors. However, less probable double-decay ($|f\rangle \rightarrow |e\rangle \rightarrow |g\rangle$) or heating ($|g\rangle \rightarrow |e\rangle$) errors can still propagate to cat X errors, and these errors place an upper limit on the achievable bias. We thus refer to this transmon-controlled CX gate as only *moderately* noise biased, and we reserve the term *exponentially* noise biased for implementations where the cat's bias grows arbitrarily and exponentially with $|\alpha|^2$. In Sec. IV below, we numerically demonstrate that significant bias is nevertheless achievable with this CX gate.

To correct the cats' suppressed X errors, we propose an implementation of a cat-controlled transmon X rotation, referred to as the CR_X gate, that is exponentially noise biased. The CR_X gate is implemented via a composite pulse sequence, involving monochromatic drives applied to the transmon and storage mode, where again entanglement is generated by the dispersive coupling and engineered dissipation is turned off during the gate. The duration of the CR_X gate can be tuned, and we choose to fix the duration to match the CX duration so no additional delay padding is required during syndrome measurement cycles of Fig. 1(c). Notably, the CR_X gate can be exponentially noise biased, despite only requiring operating the transmon in the $|g\rangle, |e\rangle$ manifold. We discuss the implementation in detail in Sec. III below, and we show that the gate is exponentially noise biased in Sec. IV below.

Having summarized the CX and CR_X implementations, we highlight the practical appeal of these gates. Both gates rely on native dispersive coupling to generate entanglement, so complicated Hamiltonian engineering is not required. This simplicity has already enabled multiple experimental demonstrations of the CX gate [41–43], along with a demonstration of a rudimentary CR_X gate in Ref. [41] (where it was used to enable Z -basis cat readout; see Appendix A5). Moreover, both gates also avoid the use of engineered dissipation, so that the dissipation strength does not contribute to gate fidelity. Instead, as observed in Ref. [51], it suffices to apply

the engineered dissipation only intermittently, while the cat qubits idle [see Fig. 1(c)]. As demonstrated experimentally in Ref. [37,41], the required dissipation strength to preserve X error suppression with this “pulsed” dissipation is quite modest (the requirement on dissipation strength is quantified in Appendix B). Both the simplicity of implementation and modest requirement on engineered dissipation strength for these cat-transmon gates are points of contrast with known implementations of cat-cat gates, and, as we show in Sec. IV below, the former offer high fidelity and significant bias with current state-of-the-art parameters.

III. CR_X GATE: CAT-CONTROLLED TRANSMON ROTATION

In order to measure Z -type syndromes and hence correct cat X errors, we require a cat-controlled entangling gate. For this purpose, we propose a gate implementation, dubbed the CR_X gate, where the transmon and cat are entangled via an alternating sequence of storage and transmon drives. The gate is exponentially noise biased, and incorporates both dynamical decoupling and pulse shaping to minimize error. The CR_X gate builds on ideas from the gate implementations in Refs. [50,52–54] to realize an exponentially noise-biased cat-controlled gate.

In the following subsections, we first describe the gate's implementation, then provide further detail on two aspects—storage echoing and selective pulse shaping—that are crucial to achieving high fidelity. Numerical simulations of the gate are presented in Sec. IV below. Additionally, while here we focus on the CR_X gate's use to correct cat X errors, in Appendix A5 we discuss other applications of the gate that may be of independent interest. These include realizations of high-fidelity cat Z and controlled- Z (CZ) gates (i.e., with infidelity scaling as approximately κ_1/χ rather than as approximately $\sqrt{\kappa_1/\kappa_2}$), as well as single-shot Z -basis cat readout with readout error exponentially suppressed in $|\alpha|^2$.

A. Pulse sequence

The main idea behind the gate is sketched in Fig. 2. First, a drive on the storage mode is applied to displace the mode by $+\alpha$, such that the coherent states $|\alpha\rangle$ and $|\alpha\rangle$ are mapped to $|2\alpha\rangle$ and $|0\rangle$, respectively. Next, a *number-selective* transmon pulse is applied, such that the transmon is flipped from $|g\rangle \leftrightarrow |e\rangle$ only if the storage is in $|0\rangle$, and the transmon is ideally unaffected otherwise. This selectivity is enabled by the cat-transmon dispersive coupling [Eq. (7)]. This coupling shifts the transmon frequency conditioned on the number of photons in the storage, so a number-selective pulse is achieved simply by applying a frequency-selective pulse to the transmon. Finally, the storage mode is displaced back by $-\alpha$, returning the system to the cat code space.




FIG. 2. Simplified schematic of the CR_X gate. To implement the cat-controlled transmon rotation, the cat is first displaced by $+\alpha$, such that $|- \alpha\rangle \rightarrow |0\rangle$. A number-selective drive is then applied to the transmon, with the drive frequency resonant with the transmon frequency when the storage is in vacuum. The storage mode is subsequently displaced back, and the net effect is a rotation of the transmon that occurs only if the cat was initially in $|- \alpha\rangle$. Note that this simplified diagram neglects transmon-dependent rotation of the storage mode during the selective pulse due to the dispersive coupling.

The net result of this procedure is to flip the state of the transmon only if the cat qubit was initially in state $|- \alpha\rangle$. More precisely, the gate implements the unitary

$$\text{CR}_X = |\alpha\rangle \langle \alpha| \otimes I + |- \alpha\rangle \langle - \alpha| \otimes R_X(\pi), \quad (9)$$

where $R_X(\pi) = e^{-i\pi X/2}$, with $X = |g\rangle \langle e| + \text{H.c.}$ We refer to this unitary as the CR_X gate. The CR_X gate is locally equivalent to a cat-controlled CX gate, the difference being that an additional phase gate, $S^\dagger = \text{diag}(1, -i)$, is applied to the cat by the CR_X gate. This local phase must be accounted for, but does not inhibit our ability to use the CR_X gate for syndrome measurements (see Appendix H).

Why is the CR_X gate exponentially noise biased? The intuition is that a cat X error requires a large unintentional rotation of the storage mode. If the CR_X gate's constituent pulses can be applied quickly enough, such that there is not sufficient time for the cat-transmon dispersive coupling to generate such a large rotation, then the probability of a cat X error will be minimal. Indeed, through numerical simulations (discussed below) and analytical arguments (see Appendix A 2), applying the pulses in a time $T_{\text{CR}_X} \leq \pi/\chi$ is sufficient to achieve a low cat X error probability. We emphasize that, in contrast to the CX gate, the CR_X gate does not require that we use higher transmon levels to achieve a large cat noise bias.

The above description of the CR_X gate is a simplification. In practice, inserting additional pulses is both necessary to keep the cat phase well defined and beneficial to protect against dephasing as discussed below. The gate sequence we actually propose is as follows:

- (1) storage displacement: $D(+\alpha)$;
- (2) selective transmon pulse: $R_X(\pi/4)$;
- (3) unselective transmon pulse: $R_X(\pi)$;
- (4) selective transmon pulse: $R_X(\pi/4)$;
- (5) storage displacement: $D(-2\alpha e^{i\phi})$;

- (6) selective transmon pulse: $R_X(-\pi/4)$;
- (7) unselective transmon pulse: $R_X(-\pi)$;
- (8) selective transmon pulse: $R_X(-\pi/4)$;
- (9) storage displacement: $D(\alpha e^{2i\phi})$.

Here, $D(\alpha)$ denotes a storage displacement with amplitude α , and $\phi = -\chi T_{\text{sel.}}$, where $T_{\text{sel.}}$ denotes the duration of a selective $R_X(\pm\pi/4)$ pulse. We emphasize that though the sequence comprises nine different pulses, all of the constituent pulses are routine and easily implementable. In Appendix A, we show that this sequence is equivalent to the CR_X gate up to an $R_X(\pi/2)$ rotation on the transmon, and we also show that, if desired, the gate can be reduced to five pulses at the price of increased sensitivity to dephasing. We also note that, by substituting the selective $R_X(\pi/4)$ rotations with $R_X(\theta/4)$ rotations, the above pulse sequence also suffices to implement a general cat-controlled $R_X(\theta)$ rotation of the transmon.

B. Mitigating dephasing with transmon and storage echoes

The pulse sequence incorporates echoes on both the transmon and storage in order to protect against low-frequency dephasing noise. The transmon echoes, implemented via the two unselective transmon π pulses (steps (3) and (7)), are a standard technique to suppress the impacts of low-frequency dephasing noise [55]. Such echoes are similarly employed in the echoed conditional displacement gate [53] for example.

Importantly, though, in the case of the CR_X gate these transmon π pulses are necessary to ensure that the phase of the cat is well defined, as in Ref. [52]. During the selective pulses, the dispersive coupling causes the storage mode to accumulate phase conditioned on the transmon state. The π pulse at step (3) [(7)] ensures that, absent transmon decay or heating, the storage mode accumulates a fixed phase $\phi = \chi T_{\text{sel.}}$ over steps (2)–(4) [(6)–(8)] that is independent of the initial transmon state. Without echoes, the storage phase accumulated by the end of the gate would be dependent on the transmon state, making the final displacement back to the code space (step 3 of Fig. 2) impossible. With echoes, the total phase accumulated by the storage mode is fixed at 2ϕ , and this phase accumulation can be tracked classically in software.

The “storage echo,” implemented via the intermediate displacement [step (5)], is a nonstandard technique to suppress the impacts of low-frequency storage dephasing, and we illustrate the idea conceptually in Figs. 3(a) and 3(b). Photon-number-preserving perturbations, like dephasing or self-Kerr, do not normally induce cat Z errors, since they commute with the cat X operator (the photon-number parity operator). However, these same perturbations can induce Z errors in a displaced cat. Indeed, in the displaced frame the dephasing operator becomes $O_{\text{dephase}} =$




FIG. 3. Storage “echo” to mitigate dephasing-induced phase flips during the CR_X gate. (a) Without the echo displacement, spurious rotations in a displaced cat can result in parity flips (cat Z errors) once displaced back, graphically indicated by the positioning of the red interference fringe at the origin in the right diagram. (b) To mitigate these Z errors, an “echo displacement” is applied in between successive selective pulses, such that the cat spends equal time on the left and right halves of the phase space. (c) The final probability of a cat Z error is plotted as a function of the storage dephasing strength for both white and $1/f$ dephasing noise. The inclusion of the storage echo significantly reduces the error probability for the low-frequency-dominated $1/f$ noise, but not for white noise. Markers indicate trajectory simulation results, and solid lines are analytical predictions. See Appendix A 3 for derivations and numerical details.

$a^\dagger a \rightarrow (a^\dagger + \alpha^*)(a + \alpha)$, which contains terms that do not preserve photon-number parity. To mitigate these errors, the storage echo displaces the cat to the opposite side of phase space midway through the gate, so that on average the parity-non-preserving components of perturbations like low-frequency dephasing or self-Kerr cancel. Explicitly, one can see this cancellation by moving to the interaction picture with respect to the displacements, then computing the time-averaged dephasing operator $\bar{O}_{\text{dephase}}$ (i.e., a leading-order Magnus expansion [56]),

$$\begin{aligned} \bar{O}_{\text{dephase}} &\approx \frac{1}{2}[(a^\dagger + \alpha^*)(a + \alpha) + (a^\dagger - \alpha^*)(a - \alpha)] \\ &= a^\dagger a + \text{const.} \end{aligned} \quad (10)$$

We demonstrate the efficacy of this storage echoing technique in Fig. 3(c). Following Ref. [57], in Appendix A, we show that the cat Z error probability as a function of time is given by $p_Z(t) = 1 - \frac{1}{2}(1 + e^{-\Gamma_\phi(t)})$, where

$$\Gamma_\phi(t) = 16|\alpha|^4 \int_0^\infty \frac{d\omega}{2\pi} S(\omega) \frac{F(\omega t)}{\omega^2}. \quad (11)$$

Here $S(\omega)$ is the frequency-noise power spectral density, and $F(\omega t)$ is a filter function that takes different forms depending on whether the echo is applied. Without the echo, $F(\omega t) = 2 \sin^2(\omega t/2)$, and with the echo, $F(\omega t) = 8 \sin^4(\omega t/2)$. In Fig. 3(c) we plot these analytical expressions for the dephasing rate alongside the results of numerical simulations incorporating dephasing noise with either white-noise [$S(\omega) = \text{constant}$] or $1/f$ [$S(\omega) \propto 1/\omega$] spectra [58] (see Appendix A 3). As expected, we find that the storage echo significantly reduces in the probability of a cat Z error due to the low-frequency-dominated $1/f$

dephasing noise, but does not offer any protection against white-noise dephasing.

C. Mitigating coherent error with pulse shaping

Coherent errors arise in the CR_X gate due to the finite duration of the selective transmon pulses. As an example, consider step (2) of the gate, where an ideal selective pulse should enact an $R_X(\pi/4)$ rotation of the transmon if the storage mode is in $|0\rangle$, and act trivially on the transmon if the storage mode is in $|2\alpha\rangle$. In the former case, the transmon frequency is not dispersively shifted and takes its bare value ω_0 . In the latter case, there is an indeterminate dispersive shift; the transmon frequency is $\omega_0 + n\chi$, where the probability of having n photons in the storage is $p(n) = |\langle 2\alpha|n\rangle|^2$. Coherent errors necessarily arise in finite-duration selective pulses because the desired process at ω_0 cannot be fully spectrally resolved from undesired processes detuned by $n\chi$, which is especially challenging at small cat mean photon numbers. The probability of an unwanted transition process detuned by Δ is related to $p(n = \Delta/\chi)$, together with the spectral content of the selective pulse shape, $\Omega(t)$, at the corresponding detuning, $|S_\Omega(\Delta)|^2$, where $S_\Omega(\Delta) = \int dt \Omega(t) e^{i\Delta t}$.

Conveniently, these coherent errors can be easily suppressed through pulse shaping. Figure 4(a) conceptually illustrates how a simple DRAG (derivative removal by adiabatic gate) pulse [59,60] mitigates these errors. The DRAG pulse augments the original pulse $\Omega_0(t)$ by adding an out-of-phase component proportional to its derivative, $\Omega(t) = \Omega_0(t) - \beta i \dot{\Omega}_0(t)$, where the constant β is specifically chosen to suppress spectral content at $\Delta = |2\alpha|^2 \chi$, i.e., at the photon number where $p(n)$ is maximal. Figure 4(b) numerically illustrates the efficacy of this




FIG. 4. Mitigating CR_X coherent error with pulse shaping. (a) Frequency spectra of standard and shaped pulses. The shaped DRAG pulse is designed to minimize spectral content at the frequency of the transmon when $|2\alpha|^2$ photons are in the storage, i.e., $\omega_q + 2|\alpha|^2\chi$, where ω_q is the unshifted transmon frequency. To guide the eye, the photon-number distribution for the coherent state $|2\alpha\rangle$ is plotted (on a linear scale, not shown), with $|\langle 2\alpha|n\rangle|^2$ positioned at $\omega_q + n\chi$ on the horizontal axis. At the frequencies where $|2\alpha\rangle$ has significant support, the spectral content of the DRAG pulse is significantly suppressed relative to the case without DRAG. (b) Coherent error reduction via DRAG. The coherent error of different selective pulse ansatzes is plotted as a function of $|\alpha|^2$, where at each point the free parameters of a pulse ansatz are optimized to minimize coherent error. The coherent error is computed as the average gate infidelity of the Pauli error channel extracted from simulations of the CR_X gate without decoherence (see Appendix D).

simple technique. Replacing a simple Gaussian $\Omega_0(t)$ with this “standard” DRAG pulse results in nearly an order-of-magnitude reduction in the coherent error of a CR_X gate. Figure 4(b) also shows that CR_X coherent errors are suppressed with increasing $|\alpha|^2$. This suppression is a consequence of the increasing average dispersive shift $|2\alpha|^2\chi$, so that it becomes easier to spectrally resolve transmon frequencies corresponding to the storage $|0\rangle$ and $|2\alpha\rangle$ states, regardless of whether pulse shaping is used.

The standard DRAG protocol is designed to suppress an unwanted transition at only a single frequency, but the CR_X coherent error can be even further reduced by generalized DRAG pulses that suppress multiple unwanted transitions simultaneously [61,62]. For example, the pulse

$$\Omega(t) = \Omega_0(t) - i \left[\frac{1}{\Delta_1} + \frac{1}{\Delta_2} \right] \dot{\Omega}_0(t) - \frac{1}{\Delta_1 \Delta_2} \ddot{\Omega}_0(t) \quad (12)$$

approximately suppresses transitions at two frequencies, Δ_1 and Δ_2 ; see Appendix A for a derivation. By choosing to distribute the Δ_1 and Δ_2 over the range of frequencies where $p(n)$ is appreciable, it is possible to achieve significantly lower coherent error with this “multicomponent DRAG” than with standard DRAG, as demonstrated in Fig. 4(b). Moreover, while, in general, DRAG’s addition of out-of-phase components can degrade the fidelity of the desired transition, in Appendix A we show that this degradation can be entirely avoided in the context of the CR_X gate. In particular, adding a time-dependent pulse detuning,

$$\delta(t) = -\text{Re } \Omega(t) \tan \left[\int_0^t dt' \text{Im } \Omega(dt') \right] \quad (13)$$

and adjusting the amplitude of $\Omega(t)$ to satisfy

$$\theta = \int_0^T dt \text{Re } \Omega(t) \sec \left[\int_0^t dt' \text{Im } \Omega(t') \right] \quad (14)$$

exactly implements an $R_X(\theta)$ rotation at the desired transition, regardless of any DRAG correction. Pulse shaping thus enables negligible CR_X coherent errors even at modest cat mean photon numbers. For example, in Fig. 4(b) the multicomponent DRAG pulses achieve coherent errors $< 10^{-3}$ already for $|\alpha|^2 \gtrsim 4$. Even so, the residual CR_X coherent error disincentivizes operation at small $|\alpha|^2$. See Appendix A for derivations and details on the numerics.

IV. ERROR AND BIAS OF CAT-TRANSMON GATES

In this section, we numerically benchmark the performance of the two entangling gates, CX and CR_X , that are required for error correction in the cat-transmon architecture. Despite the limitations imposed by transmon decoherence, we find that the cat-transmon gates can be implemented with high fidelity and bias for currently achievable coupling strengths and coherence times. We begin by first describing the noise model used in the simulations, then present projections for the fidelity and bias of the two gates.

A. Noise model and assumptions

The noise model used in simulations of the cat-transmon gates is summarized in Table I. The model is parametrized

TABLE I. Single-parameter noise model used for gate simulations in the main text. All decoherence rates for the storage and transmon are varied in proportion with parameter q , with $q \sim 10^{-5}$ – 10^{-4} achievable assuming current state-of-the-art coherence.

Decoherence source	Value
Storage single-photon loss (κ_1/χ)	q
Transmon loss ($\gamma_\downarrow/\chi$)	$3q$
Storage pure dephasing (κ_ϕ/χ)	$0.01q$
Transmon pure dephasing (γ_ϕ/χ)	$1.5q$
Transmon heating ($\gamma_\uparrow/\chi$)	$0.015q$

by a single dimensionless parameter, q , that serves as a measure of decoherence in both the storage mode and transmon. In our gate simulations, the only physically relevant quantities are ratios of decoherence rates to χ , since both the CX and CR_X gate times scale inversely with χ , so these ratios are all specified in terms of q . This single-parameter model is both convenient to analyze and embodies the fact that, in practice, the coherence of resonators and transmons in a device often vary together as a function of the quality of materials and fabrication.

For context, values of q in the range of 10^{-4} – 10^{-5} are realistically achievable with state-of-the-art devices in the near term. For example, the device of Ref. [63] simultaneously exhibited near-millisecond storage $T_1 = \kappa_1^{-1}$ and megahertz scale χ , achieving $\kappa_1/\chi = q \approx 10^{-4}$. With few-millisecond storage T_1 , as has already been demonstrated in similar devices [64], q closer to 10^{-5} is achievable. Realistic values of q for near-term two-dimensional (2D) devices are higher by roughly an order of magnitude, $q \sim 10^{-3}$ – 10^{-4} , with $q \approx 10^{-3}$ representative of the device in Ref. [41] as an example. In the long term, transmon coherence could limit the ability of the cat-transmon architecture to capitalize on ultrahigh intrinsic storage T_1 (> 10 ms [65]); expectations for achievable q in the long term for both 2D and 3D platforms are discussed in Sec. VI below.

We note some of the physical considerations that motivate the specific values in Table I. The model assumes that the storage mode $T_1 = \kappa_1^{-1}$ is 3 times larger than the transmon $T_1 = \gamma_\downarrow^{-1}$, which is roughly representative of recent 3D devices [27,63] as well as the 2D device of Ref. [41]. This separation between storage and transmon lifetimes may be less representative of 2D devices generally, however, so in Appendix C we consider an alternative noise model where storage and transmon lifetimes are comparable. For the transmon, the model of Table I assumes an effective thermal population of $\gamma_\uparrow/\gamma_\downarrow = 0.5\%$, and $T_1 = T_2$, where $T_1 = \gamma_\downarrow^{-1}$ and $T_2 = (\gamma_\downarrow/2 + \gamma_\phi)^{-1}$, both of which are reasonable assumptions for typical

transmon devices. For the storage mode, while dephasing noise spectra of superconducting resonators have been observed to exhibit a $1/f$ frequency dependence [66], we simulate white rather than $1/f$ dephasing noise for the sake of numerical efficiency. The choice of a small white-noise pure dephasing rate $\kappa_\phi \ll \kappa_1$ is then motivated by two factors: (1) it is representative of 3D cavities, which exhibit little intrinsic dephasing [53,65], and (2) as Fig. 3 shows, the choice of a small white noise dephasing rate ($\kappa_\phi \ll \kappa_1$) is representative of a much larger $1/f$ dephasing rate ($\kappa_\phi \sim \kappa_1$) in terms of the impact on gate performance.

Finally, we emphasize that our simulations assume that the cat-transmon interaction is described by the simple, effective dispersive coupling Hamiltonian in Eq. (7). In practice, to produce quantitative agreement with a specific experimental realization of a cat-transmon system, it would likely be necessary to employ a more detailed model that incorporates, e.g., details of how the tunable coupling and dissipative stabilization are realized. However, we deliberately adopt this simple effective model so that we can assess the potential of the proposed architecture generally, without restricting our focus to a particular hardware realization. See Appendix E for further discussion of the simulations and assumptions.

B. Gate performance

To quantify the gates' performance, we perform master equation simulations with decoherence rates as specified in Table I. From these simulations we extract a Pauli error channel describing each gate's errors, as discussed in Appendix D. We use these Pauli channels to compute the average gate infidelity $1 - F$ [67], cat bit-flip probability $p_{\text{bit}}^{(\text{cat})}$, and the cat noise bias $\eta^{(\text{cat})}$ as a function of $|\alpha|^2$ and q .

The average gate infidelity of the CX and CR_X gates is shown in Fig. 5(a). At sufficiently large $|\alpha|^2$ the gate infidelities of the two gates are similar. Both become dominated by cat Z errors, which occur with probability $\kappa_1|\alpha|^2(\pi/\chi) = \pi|\alpha|^2q$. At smaller $|\alpha|^2$, the infidelities of the two gates display qualitatively different behavior. While the CX gate is only limited by decoherence, the CR_X gate is also limited by coherent errors in the selective transmon pulses (caused by the difficulty in spectrally resolving dispersive shifts of the $|0\rangle$ and $|2\alpha\rangle$ states at small $|\alpha|^2$; see Sec. III C). Low infidelities, $1 - F < 10^{-3}$, can be achieved for both gates simultaneously for realistically achievable values of q in the range 10^{-5} – 10^{-4} .

Figure 5(b) shows $p_{\text{bit}}^{(\text{cat})}$, the probability of cat bit-flip error (i.e., the probability of an X - or Y -type Pauli error on the cat) for each gate. The clear qualitative difference between $p_{\text{bit}}^{(\text{cat})}$ between the CX and CR_X gates highlights that the CX gate is only moderately noise biased, while the CR_X gate is exponentially noise biased. In particular, for the CX




FIG. 5. Error and bias of transmon-controlled CX (top row) and cat-controlled CR_X (bottom row) gates. Pauli error channels are extracted from master equation simulations of the gates, and the resultant (a) average gate infidelity, (b) cat bit-flip probability, and (c) cat noise bias are plotted as a function of q (which specifies the decoherence rates of the cat and transmon per Table I) and $|\alpha|^2$. In (a), red dashed lines are infidelity contours (shown only for $|\alpha|^2 \geq 6$) for an error model with only cat Z errors, at rate $\kappa_1 |\alpha|^2$. The close agreement of these contours with the full contours (white) in both the CX and CR_X plots indicates that infidelities of both gates are dominated by cat Z errors at sufficiently large $|\alpha|^2$.

gate, as $|\alpha|^2$ increases, $p_{\text{bit}}^{(\text{cat})}$ approaches the lower bound

$$p_{\text{bit,CX}}^{(\text{cat})} \geq \frac{1}{2} \left[\frac{\gamma_{\uparrow}}{2} \frac{\pi}{\chi} + \left(\frac{\gamma_{\downarrow}}{2} \frac{\pi}{\chi} \right)^2 \right], \quad (15)$$

where the two terms correspond to the probabilities of cat bit flips induced by transmon heating and double decay, respectively. Despite the lack of exponential suppression, this probability can be small because of the quadratic suppression of the transmon's dominant loss errors and small heating rates at low temperature. In contrast, $p_{\text{bit},CR_X}^{(\text{cat})}$ continues to decrease exponentially with $|\alpha|^2$ because transmon errors cannot propagate through the CR_X gate to large phase space rotations of the storage mode. We note that it is difficult to write down a simple analytical expression like Eq. (15) for $p_{\text{bit},CR_X}^{(\text{cat})}$ since $p_{\text{bit},CR_X}^{(\text{cat})}$ is a multiexponential function containing contributions from both coherent pulse errors and transmon decoherence; see Appendix A 2 for a detailed discussion.

Finally, Fig. 5(c) shows the cat noise bias $\eta^{(\text{cat})}$ for each gate. We define $\eta^{(\text{cat})} \equiv p_Z^{(\text{cat})}/p_{\text{bit}}^{(\text{cat})}$, where $p_Z^{(\text{cat})}$ is the probability of a Z-type Pauli error on the cat. Importantly, while only the CR_X gate exhibits exponentially increasing bias with $|\alpha|^2$, both gates nevertheless exhibit large bias at modest $|\alpha|^2$. In particular, for realistic q , $\eta^{(\text{cat})} > 10^3$ is simultaneously achievable for both gates. The combination of large bias and low infidelity can enable significant

hardware-efficiency improvements for error correction, as we show in the next section.

V. LOGICAL MEMORY PERFORMANCE PROJECTIONS

By leveraging the CX and CR_X gates, hybrid cat-transmon surface codes can be constructed that are capable of correcting both the cats' dominant Z errors and their suppressed X errors (see Fig. 1). In this section, we simulate error correction with these codes to quantify their performance as logical memories. We first assess how realistically achievable error rates compare to the error threshold, then we quantify the hardware-efficiency benefit of the architecture's noise bias by comparing logical memory overhead to the case without biased noise. All simulations in this section employ circuit-level noise, with d_Z rounds of noisy syndrome measurements simulated using Stim [68], and minimum weight perfect matching decoding performed using PyMatching [69].

A. Error threshold

To determine the error threshold, we continue with the noise model in Table I and directly use the CX and CR_X Pauli error channels from Sec. IV in our circuit-level noise model. In addition, we also specify error channels for cat idling and transmon readout and reset operations (see the syndrome extraction circuits in Fig. 1). We take 1%




FIG. 6. Logical Z memory error rates (per d_Z rounds) of the cat-transmon architecture. The logical error rate for different code distances is plotted as a function of q at fixed $d_X = 5$ and $|\alpha|^2 = 6$. These fixed values are sufficient to achieve $p_X^{(L)} \ll p_Z^{(L)}$ for all q and d_Z shown (see the main text). The “threshold” $q \approx 5 \times 10^{-4}$ is indicated by the dashed line. Accessible q values for current state-of-the-art devices are indicated by the gray region, which is defined here by computing the max and min achievable q for κ_1 and χ in the ranges $\kappa_1^{-1} \in [1, 2]$ ms and $\chi/2\pi \in [1, 3]$ MHz. We take 10^7 samples for each data point, and error bars indicate 95% confidence intervals.

combined error probability for transmon readout and reset [70]. We assume that the total duration of the readout and reset operations is the same as the duration of the CX and CR_X gates, i.e., π/χ , so that the noise model in Table I can also be directly used in simulations of cat idling during transmon readout and reset. We then extract a Pauli error channel for cat idling subject to dissipative stabilization. We assume only a modest engineered dissipation rate relative to the dispersive coupling, $\kappa_2 = \chi/10$, which is sufficient to prevent accumulation of leakage out of the cat code space (see Appendix B).

In Fig. 6, we plot the logical Z error probability after d_Z rounds of noisy syndrome measurement, $p_Z^{(L)}$, as a function of the d_Z and q , at fixed $d_X = 5$ and $|\alpha|^2 = 6$. We observe exponential suppression of $p_Z^{(L)}$ with increasing d_Z below a “threshold” value of $q = 5 \times 10^{-4}$. To place this value in context, the region of realistically achievable q for near-term state-of-the-art devices is shaded in gray. We see that error rates more than 10 times below threshold—the “deep subthreshold regime”—are achievable in near-term devices, even with engineered dissipation of only modest strength.

It should be noted that we use the term threshold loosely here. The logical X error rate after d_Z syndrome measurement rounds, $p_X^{(L)}$, is lower bounded at fixed d_X and $|\alpha|^2$, so identifying the proper asymptotic threshold would eventually also require increasing d_X . In practice, however,

the code distances and error rates required to reach algorithmically relevant logical error rates are more relevant than the asymptotic threshold. Accordingly, $d_X = 5$ and $|\alpha|^2 = 6$ were chosen because they are sufficiently large to ensure that $p_X^{(L)} \ll p_Z^{(L)}$, with $p_X^{(L)} < 10^{-10}$ in the deep subthreshold regime (see Appendix F). Additionally, we note that, while the XZZX surface code actually offers a higher threshold, we focus on the standard surface code here because it offers lower overheads in the deep subthreshold regime (see Appendix I).

B. Comparison to unbiased noise

To more readily enable the comparison with the unbiased-noise case, in this subsection we adopt a simplified circuit-level noise model to describe the cat-transmon architecture. In this simplified model, a biased Pauli error channel is applied after every two-qubit gate, and after data qubit idling during ancilla qubit readout. The biased error channels are specified in terms of two parameters, p_Z and η : all Z-type errors have equal probability, with the total probability of a Z-type error given by p_Z , and the total probability of any other error given by p_Z/η . This biased error model is compared with an unbiased error model where the same two-qubit gates and data idling operations are instead subject to a uniform depolarizing channel with entanglement infidelity [71] $1 - F_e = p$. The biased error channel’s entanglement fidelity is given by $1 - F_e = p_Z(1 + 1/\eta)$, which approaches p_Z for large bias η , so that comparing p and p_Z is reasonable. Both models assume 1% combined ancilla qubit readout and reset error probability.

We find that this simplified biased-noise error model is sufficient to provide quantitatively accurate predictions for the cat-transmon architecture’s logical error probabilities. In particular, in Appendix G we show how the parameters p_Z and η can be extracted from the gates’ full Pauli error channels, such that the logical errors with the simplified error model and with the full error model used in Sec. V A agree well. To show what values of p_Z and η are achievable in the cat-transmon architecture, in Fig. 7 we plot the achievable values of these parameters as a function of q . Consistent with the findings in Sec. IV, $p_Z \sim 10^{-3}$ and $\eta \sim 10^{3.5}$ are simultaneously achievable with state-of-the-art coherence, i.e., q between 10^{-4} and 10^{-5} . In Appendix C, we present analogous plots of the achievable p_Z and η for alternative noise models. There we find that slightly larger bias $\eta \sim 10^4$ is achievable when the storage and transmon lifetimes are comparable. In the analysis below, we thus consider the range $\eta \sim 10^3$ – 10^4 as broadly representative of what is achievable with the cat-transmon architecture.

The metric we use to compare the biased- and unbiased-noise cases is the logical memory overhead, defined as the minimum number of qubits (data and ancilla) required to




FIG. 7. Achievable bias and Z error probability. Colors indicate the maximum q sufficient to achieve a given value of p_Z (Z-type error probability) and η (bias). To guide the eye, contours of constant q are shown: any point in parameter space above a given contour is achievable with that value of q . The black dashed line indicates the maximum beneficial bias (i.e., the point beyond which the bias only increases because of increasing cat Z error probability without also decreasing X error probability; see Appendix A 2 for details). The plot is produced by simulating CX and CR_X gates at different values of $|\alpha|^2$ and q , then computing p_Z and η from their Pauli error channels as described in Appendix G. For visual clarity, Voronoi cells are constructed around these nonuniformly distributed parameter space points, and each cell is colored according to the corresponding q .

reach a given target total logical error $p^{(L)} = p_X^{(L)} + p_Z^{(L)}$. For a rectangular surface code with X and Z distances d_X, d_Z , the total number of qubits is given by $2d_X d_Z - 1$. To compute the overhead, we thus find the values d_X and d_Z that minimize $2d_X d_Z - 1$ subject to the constraint that logical error falls below the given $p^{(L)}$. In the unbiased-noise case, square codes $d_X = d_Z$ are optimal, but in the biased-noise case rectangular codes $d_X < d_Z$ can offer lower overheads (see Appendix F).

In Fig. 8, we plot logical memory overheads for the unbiased and biased error models as a function of p and p_Z , respectively, for different choices of the bias η and target logical error $p^{(L)}$. Even for the modest bias achievable with the cat transmon architecture, $\eta \sim 10^3$ – 10^4 , there is significant reduction in the overhead relative to the unbiased-noise case. Indeed, in the deep subthreshold regime, we find that realistically achievable noise biases can provide up to a factor 5 reduction in logical memory overhead relative to the case without noise bias. For example, with $\eta = 10^{3.5}$ and $p_Z = 10^{-3}$ for the biased-noise case, the overhead to achieve $p^{(L)} = 10^{-10}$ is about 200 qubits, while the overhead to achieve the same $p^{(L)}$ with $p = 10^{-3}$ is nearly 1000 qubits for the unbiased case.

With Fig. 8, one can also compare the physical error required to achieve the target logical error at a given

overhead. Consider again the values $p_Z \sim 10^{-3}$ and $\eta \sim 10^3$ – 10^4 that are achievable with the cat-transmon architecture assuming state-of-the-art coherence. Across the different $p^{(L)}$, we see that the overhead of the biased-noise architecture at these p_Z, η values is comparable to the overhead of the unbiased-noise architecture with $p \sim 10^{-5}$ – 10^{-4} . That is, the logical memory performance of the cat-transmon architecture with realistically achievable physical error and bias is roughly equivalent to that of a generic unbiased-noise architecture with physical error in the range 10^{-5} – 10^{-4} .

To provide context for this comparison, we consider a realization of an unbiased-noise architecture with transmon qubits and crudely estimate the transmon lifetimes that would be required to achieve $p \sim 10^{-5}$ – 10^{-4} . Our unbiased error model assumes that two-qubit gates and data qubit idling during readout each have error p , so that the average probability of a data qubit error is roughly $5p$ per error-correction cycle. Assuming an error-correction cycle duration of $T_{\text{cycle}} = 1 \mu\text{s}$, and that the performance of transmon-transmon gates is limited only by decoherence, the range $p \sim 10^{-5}$ – 10^{-4} thus corresponds to transmon lifetimes in the range of $T_{\text{cycle}}/5p = 2$ – 20 ms. Current state-of-the-art transmon lifetimes lie in the range of about 0.3–1 ms [72–74], approaching the minimum lifetime required for $p = 10^{-4}$ but below that required for $p = 10^{-5}$. That said, we emphasize that this is only a crude estimate, and indeed making a more accurate comparison would require additional considerations beyond coherence limits. For example, the cats’ dissipative stabilization also serves as built-in means of leakage reduction, so one should also account for any additional cost associated with leakage reduction in the unbiased case [75–77]. An accurate comparison will also require accounting for integration challenges (and opportunities) that arise when building larger devices. Identifying and quantifying these integration aspects for the cat-transmon architecture is thus an important direction for future research.

VI. DISCUSSION AND OUTLOOK

In this work, we have proposed and analyzed the hybrid cat-transmon architecture—a scalable architecture for hardware-efficient error correction. In this architecture, the dissipative cats’ noise bias enables hardware efficiency, and the use of ancillary transmons enables syndrome extraction with simple, high-fidelity cat-transmon entangling gates. To correct the cats’ dominant Z errors, we employ a moderately noise-biased CX gate that relies only on free evolution under the native dispersive coupling—a highly simple implementation that has enabled multiple experimental demonstrations [41–43]. To correct the cats’ suppressed X errors and hence enable full scalability, we propose the exponentially noise-biased CR_X gate that relies primarily on simple number-selective




FIG. 8. Logical memory overheads ($= 2d_X d_Z - 1$) for different target $p^{(L)}$ per d_Z rounds are computed for unbiased noise (black lines) and biased noise (gray, orange lines) as a function of the respective error parameters, p and p_Z . Different biases in the achievable range $\eta \sim 10^3$ – 10^4 are indicated by different shades of orange. Regions between these curves are shaded to visually highlight how varying the bias over this range impacts the overhead. Overheads for lower bias, $\eta = 10^2$, are plotted in gray for reference. For the biased-noise curves, small jumps in overhead correspond to increases in d_Z , while large jumps correspond to increases in d_X . Overheads are computed by extrapolating from fits, as described in Appendix F.

pulses, again enabled by the native dispersive coupling. In comparison to cat-cat gates, these cat-transmon gates can be fast ($\pi/\chi \sim 200$ ns) and offer high fidelity ($F \sim 99.9\%$), assuming megahertz-scale dispersive couplings and millisecond-scale storage T_1 . They also achieve sizable noise bias ($\eta \sim 10^3$ – 10^4), and though the noise bias of this architecture is ultimately limited by transmon errors, we find that significant hardware efficiency improvements are possible. For realistically achievable noise bias and error rates, projected logical memory overheads for the cat-transmon architecture are comparable to those of a generic unbiased-noise architecture with physical error probability in the range of about 10^{-5} – 10^{-4} .

The performance of the cat-transmon architecture is largely decoherence limited, so that the ratios of the cat and transmon's decoherence rates relative to the dispersive coupling strength are the main figures of merit that dictate logical performance. On the one hand, in 3D devices, ratios of the storage mode loss rate to dispersive coupling, $\kappa_1/\chi = q$, in the range $q \sim 10^{-5}$ – 10^{-4} are realistically achievable assuming megahertz-scale dispersive couplings and millisecond-scale storage T_1 , which would place such devices in the deep subthreshold regime. Recent work [65] demonstrating ultrahigh storage mode lifetimes of tens of milliseconds suggests further room for improvement. However, increasing storage lifetime by itself without also increasing transmon lifetime has diminishing returns for the performance of the cat-transmon architecture. In particular, cat Z error rates are dictated by the T_1 of dressed storage modes (i.e., storage modes hybridized with transmons), rather than the intrinsic T_1 of isolated storage modes, and the transmon lifetime ultimately limits the dressed storage T_1 via the inverse Purcell effect. Moreover,

transmon errors can directly propagate to cat Z errors through the CR_X gate, such that transmon coherence limits the achievable p_Z (see Appendix C). While currently achievable transmon lifetimes of the order of about 100 μs already suffice to access the deep subthreshold regime with millisecond-scale storage lifetimes, improvements in transmon lifetime would thus be required for this architecture to fully capitalize on ultrahigh intrinsic storage lifetimes.

On the other hand, 2D devices are naturally compact and scalable, but further improvements in coherence are required to push them into the deep subthreshold regime. Quantum acoustic devices have the potential to offer both high coherence and high scalability [78,79], but preserving the high coherence in integrated systems remains an outstanding challenge. Independent of the storage modes' physical implementation, efforts to increase the dispersive coupling strength [80] also provide a promising avenue to improved performance. A key challenge in pushing beyond megahertz-scale dispersive couplings will be suppressing the presence or mitigating the impacts of spurious nonlinearities.

While we have focused on the specific case of a transmon ancilla, we emphasize that the entangling gates we employ apply equally well to any type of dispersively coupled ancilla qubit, and alternative choices of the ancilla qubit could potentially enable further performance improvements. One intriguing possibility in this regard would be to substitute the transmon for a fluxonium. Fluxonium qubits are appealing generally because of their long coherence times and large anharmonicities [81–83]. In the context of this architecture, the ability to engineer a large fluxonium noise bias by operating off sweet spot

[84], and to engineer large dispersive couplings thanks to the large anharmonicity are appealing properties. In particular, for a fluxonium ancilla to offer an improvement over a transmon ancilla, the fluxonium-induced cat bit-flip probability during a CX gate, $p_{\text{bit}}, \text{CX}^{(\text{cat})} = \frac{1}{4}(\gamma_{\uparrow} + \gamma_{\downarrow})/\chi$, should be smaller than the corresponding probability for the transmon ancilla; cf. Eq. (15).

Modifications to either the CX or CR_X gates offer another promising avenue for further improvement of the architecture. Indeed, these two gates are somewhat mismatched; where the CX gate is only moderately noise biased yet capable of achieving high fidelity at small $|\alpha|^2$, the CR_X gate is exponentially noise biased but limited by coherent error at small $|\alpha|^2$. Instead, one can consider variations of the architecture where the benefits and drawbacks of the entangling gates are in better alignment. For example, one could retain the current CX gate and modify the CR_X gate to allow for better small- $|\alpha|^2$ performance by sacrificing its exponentially noise-biased nature. This modification could be achieved, for instance, simply by extending the CR_X gate time beyond π/χ , thereby mitigating coherent error at the price of more transmon-induced cat X errors. The resultant small- $|\alpha|^2$ architecture would likely enjoy even higher thresholds against photon loss at the price of reduced noise bias. Alternatively, one could retain the current CR_X gate and replace the CX gate with an exponentially noise-biased gate. For example, one could employ a cat ancilla and exponentially noise-biased cat-cat gates for X -type syndrome measurements, while continuing to use the transmon ancilla and cat-transmon CR_X gate for Z -type syndrome measurements. This approach has the advantage of enabling the optimizations proposed in Ref. [40] that mitigate the impact of the unfavorable $\sqrt{\kappa_1/\kappa_2}$ infidelity scaling of cat-cat entangling gates only in the context of X -type syndrome measurement. Combining this with CR_X -enabled Z -type syndrome measurements means that one need not rely exclusively on the cats' X error suppression to reach algorithmically relevant logical error rates.

Finally, while in this work we have focused on logical memory performance of the surface code as the main figure of merit, the performance of fault-tolerant logical operations will need to be evaluated to provide a complete picture of the architecture's potential. In the cat-transmon architecture, fault-tolerant logical operations can be implemented via lattice surgery [85]. Because syndrome measurement is the basic primitive involved in lattice surgery, the use of cat-cat entangling gates can still be avoided in favor of syndrome measurement with ancillary transmons, so that the practical benefits of the cat-transmon architecture persist at the level of logical operations. While we leave detailed study of fault-tolerant operations via lattice surgery in the cat-transmon architecture to future work, we note that the underlying physical noise bias has the potential to offer even further improvement in the context of logical operations. For example, this noise bias can be

exploited in the context of magic state distillation to reduce overheads further [86]. Furthermore, alternative code constructions that exploit the noise bias could be explored. As an example, one could consider a hierarchical construction [87] where some bias is retained at the level of encoded surface code logical qubits, which are then concatenated into a quantum low-density parity-check (LDPC) code [88–91] (using lattice surgery to enable any requisite long-range interactions). Further work on LDPC codes tailored to exploit the noise bias [34,92] is needed to evaluate the potential of this approach.

ACKNOWLEDGMENTS

We thank Arne L. Grimsmo, Gregory S. MacCabe, Shahriar Aghaeimeibodi, Yufeng Ye, Catherine Leroux, and Aleksander Kubica for useful discussions, as well as Colm A. Ryan for helpful feedback on the manuscript. We also thank Simone Severini, Bill Vass, James Hamilton, Nafea Bshara, and Peter DeSantis at AWS for their involvement and support of the research activities at the AWS Center for Quantum Computing.

APPENDIX A: DETAILS OF THE CR_X GATE

In this appendix, we discuss certain features of the CR_X gate in further detail. First, we show that the pulse sequence presented in the main text indeed implements the CR_X operation. Next, we demonstrate that cat bit-flip errors induced by transmon errors are exponentially suppressed with increasing $|\alpha|^2$. This allows us to identify the optimal choice of $|\alpha|^2$ for the cat-transmon architecture by weighing this exponential suppression against the bit-flip error saturation of the CX gate. Next, we derive the expressions given in the main text for the storage dephasing rate in the presence of the storage echo, and describe the generalized DRAG pulse shaping used to improve performance. Finally, we discuss additional applications of the CR_X gate that may be of independent interest. In particular, we show how the CR_X gate can be used to implement fast, high-fidelity cat Z and CZ rotations in an exponentially noise-biased manner, as well as single-shot cat Z -basis readout with readout error exponentially suppressed in $|\alpha|^2$.

1. Unpacking the pulse sequence

In this section, we show that the pulse sequence presented in Sec. III implements the CR_X unitary in Eq. (9). We begin by first showing that this unitary can also be implemented by a simplified sequence that omits the storage echo:

- (1) storage displacement: $D(+\alpha)$;
- (2) selective transmon pulse: $R_X(\pi/2)$;
- (3) unselective transmon pulse: $R_X(\pi)$;

- (4) selective transmon pulse: $R_X(\pi/2)$;
- (5) storage displacement: $D(-\alpha e^{i\phi})$.

Here $\phi = -\chi T_{\text{sel}}$, where T_{sel} is the selective pulse duration. This simplified sequence may be more convenient than the full CR_X sequence in situations where the storage echo is not needed or not helpful (e.g., if storage dephasing is negligible or is dominated by high-frequency components not mitigated by the echo).

Consider first the case where the cat qubit is initially in state $|- \alpha\rangle$ prior to the gate. The displacement in step (1) enacts $|- \alpha\rangle \rightarrow |0\rangle$. Since the selective pulses act nontrivially when the storage mode is in vacuum, a net rotation $R_X(2\pi)$ is applied to the transmon in steps (2)–(4). The final displacement enacts $|0\rangle \rightarrow |-\alpha e^{i\phi}\rangle$. Next, consider the initial state $|\alpha\rangle$. The displacement in step (1) enacts $|\alpha\rangle \rightarrow |2\alpha\rangle$. As a result, only the unselective pulse in step (3) acts nontrivially, and a net rotation of $R_X(\pi)$ is applied to the transmon in steps (2)–(4). During these steps, the transmon-storage dispersive coupling causes the storage mode state to rotate, $|2\alpha\rangle \rightarrow |2\alpha e^{i\phi}\rangle$. Note that this rotation is deterministic: the unselective transmon pulse ensures that, regardless of its initial state, the transmon is excited for time T_{sel} during these steps, so that the dispersive phase accumulated is $-\chi T_{\text{sel}} = \phi$. The final displacement enacts $|2\alpha e^{i\phi}\rangle \rightarrow |\alpha e^{i\phi}\rangle$. The net effect is thus to implement the operation

$$|\alpha e^{i\phi}\rangle \langle \alpha| \otimes R_X(\pi) + |-\alpha e^{i\phi}\rangle \langle -\alpha| \otimes R_X(2\pi).$$

Up to a local operations [a rotation of the storage mode phase by $-\phi$ and a rotation of $R_X(-\pi)$ on the transmon], this operation is equivalent to the CR_X unitary in Eq. (9).

Now, we consider the full CR_X sequence:

- (1) storage displacement: $D(+\alpha)$;
- (2) selective transmon pulse: $R_X(\pi/4)$;
- (3) unselective transmon pulse: $R_X(\pi)$;
- (4) selective transmon pulse: $R_X(\pi/4)$;
- (5) storage displacement: $D(-2\alpha e^{i\phi})$;
- (a) storage displacement: $D(-\alpha e^{i\phi})$;
- (b) storage displacement: $D(-\alpha e^{i\phi})$;
- (6) selective transmon pulse: $R_X(-\pi/4)$;
- (7) unselective transmon pulse: $R_X(-\pi)$;
- (8) selective transmon pulse: $R_X(-\pi/4)$;
- (9) storage displacement: $D(\alpha e^{2i\phi})$.

To aid in the explanation of this sequence, we have divided the storage displacement $D(-2\alpha e^{i\phi})$ in step (5) into two separate displacements $D(-\alpha e^{i\phi})$ in steps (5)(a) and (5)(b). With this division, this sequence can be roughly understood as two repetitions of the simplified sequence above. By an analogous argument, then, steps (1)–(5)(a)

implement

$$|\alpha e^{i\phi}\rangle \langle \alpha| \otimes R_X(\pi) + |-\alpha e^{i\phi}\rangle \langle -\alpha| \otimes R_X(3\pi/2),$$

and steps (5)(b)–(9) implement

$$|\alpha e^{2i\phi}\rangle \langle \alpha e^{i\phi}| \otimes R_X(-3\pi/2) \\ + |-\alpha e^{2i\phi}\rangle \langle -\alpha e^{i\phi}| \otimes R_X(-\pi).$$

Applied in sequence, the net operation is

$$|\alpha e^{2i\phi}\rangle \langle \alpha| \otimes R_X(-\pi/2) + |-\alpha e^{2i\phi}\rangle \langle -\alpha| \otimes R_X(\pi/2).$$

Up to local operations [a rotation of the storage mode phase by -2ϕ and a rotation of $R_X(\pi/2)$ on the transmon], this operation is equivalent to the CR_X unitary in Eq. (9).

2. Exponential suppression of cat bit flips and optimal $|\alpha|^2$

Like in the CX gate, the transmon-storage dispersive coupling allows transmon errors during the CR_X gate to propagate to spurious rotations of the storage mode, which can cause cat X errors. However, so long as these spurious storage mode rotations δ are sufficiently small, i.e., $\delta < \pi/2$, the probability of an X error will be exponentially suppressed with $|\alpha|^2$. In what follows we demonstrate that so long as the CR_X gate time is sufficiently short, roughly $T_{\text{CR}_X} \lesssim \pi/\chi$, the amount of spurious rotation is sufficiently small to achieve this exponential suppression.

The amount of storage mode rotation induced by transmon errors can be determined from the CR_X pulse sequence. To calculate the worst-case spurious phase accumulation due to transmon errors, consider the example of a CR_X gate applied to the initial state $|+\alpha\rangle$. During the first half of the gate, the largest spurious rotations of the storage mode are induced by transmon bit-flip errors ($|g\rangle \leftrightarrow |e\rangle$) immediately before or after the unselective pulse of step (3). Such an error effectively cancels the action of the unselective pulse, so that the transmon remains in the same state during at least the entire first half of the CR_X pulse sequence. In this situation, the transmon-dispersive coupling thus results in a storage mode phase accumulation of either 0 or 2ϕ , depending on the transmon state, each of which differs by $\pm\phi$ from the expected phase accumulation of ϕ during these steps.

Continuing with this example, the subsequent displacement in step (5) no longer maps the storage mode to the vacuum state, so the storage continues to accumulate spurious phase during the second half of the gate. Crudely speaking, we expect that yet another transmon error immediately before or after the unselective pulse in step (7) could result in additional spurious storage phase accumulation of approximately ϕ by the same logic as above. The total spurious phase accumulation over the full gate, δ ,

would thus be $\delta \approx 2\phi$. Requiring $T_{\text{sel.}} \lesssim \pi/4\chi$ (equivalent to $T_{\text{CR}_X} \lesssim \pi/\chi$ when the selective pulses dominate the total gate duration) thus ensures that $\delta \lesssim \pi/2$, so that bit-flip errors would be exponentially suppressed.

The above requirement is only approximate since we have neglected the fact that the spurious rotations and CR_X displacements do not commute. One can account for this the spurious phase accumulation more precisely as follows. Consider the specific instance of the above example where the transmon errors are such that the transmon remains in $|g\rangle$ throughout the whole gate. In this case, the only nontrivial operations are the three displacements, which combine to map initial state $|\alpha\rangle$ to $|\beta\rangle$, where $\beta = \alpha(2 - 2e^{i\phi} + e^{2i\phi})$. Relative to the expected final state of $|\alpha e^{2i\phi}\rangle$, the final state $|\beta\rangle$ corresponds to a spurious rotation of

$$\delta = \arctan \left[\frac{2(\sin \phi - \sin 2\phi)}{1 - 2\cos \phi + 2\cos 2\phi} \right]. \quad (\text{A1})$$

Note that, to leading order in ϕ , we have $|\delta| \approx 2\phi$, consistent with the simplified argument above. Using the more precise expression, we find that $|\delta| < \pi/2$ for $\phi < \pi/5$, corresponding to the requirement that $T_{\text{sel.}} \leq \pi/5\chi$ for exponential suppression of cat bit-flip errors.

We illustrate this exponential bit-flip suppression numerically in Fig. 9, where the simulated cat bit-flip probability of the CR_X gate is shown (solid lines) as a function of $|\alpha|^2$ for different choices of q . As in the main text, we have chosen pulse durations so that $T_{\text{CR}_X} = \pi/\chi$ (and, in particular, $T_{\text{sel.}} = \pi/5\chi$), so that we expect exponential suppression with increasing $|\alpha|^2$.




FIG. 9. Cat bit-flip probability of CR_X (solid) and CX (dashed) gates. For a given q , the sum of the CX and CR_X bit-flip probabilities initially decreases exponentially with $|\alpha|^2$, since CR_X bit flips are suppressed exponentially with $|\alpha|^2$. The sum will saturate once a sufficiently large α (black dots) is reached, owing to CX bit flips induced by transmon double decay and heating. Operating at $|\alpha|^2$ larger than the black dots is thus not beneficial.

The desired exponential suppression is indeed observed in Fig. 9, with different suppression exponents at different values of $|\alpha|^2$. This behavior is explained by the fact that transmon errors can be induced by different mechanisms, namely, decoherence and selective pulse errors. At small $|\alpha|^2$, transmon errors are dominated by selective pulse errors, which are themselves suppressed with $|\alpha|^2$. As a result, in this regime we see a rapid exponential suppression, since both the probability of the selective pulse errors and the cat phase-space separation are suppressed with $|\alpha|^2$. At sufficiently large $|\alpha|^2$, the suppression exponent changes, as the transmon errors begin to be dominated by decoherence (controlled by parameter q). The suppression exponent is smaller because, unlike selective pulse errors, the transmon decoherence errors are not themselves suppressed with $|\alpha|^2$. At lower transmon decoherence rates (smaller q), the rapid exponential suppression associated with selective pulse errors persists for longer before the more moderate suppression associated with transmon decoherence takes over.

Finally, for the sake of comparison, in Fig. 9 we also plot the cat bit-flip probability of the CX gate (dashed lines) at the same values of q . As discussed in the main text, the CX bit-flip probabilities saturate to a limit dictated by transmon double-decay and heating errors [see Eq. (15)]. In practice, there are diminishing returns to operating at a value of $|\alpha|^2$ larger than that where the CR_X and CX bit flips intersect (black dots). Beyond this point, the total cat bit-flip probability during error correction remains essentially constant, limited by the CX gate, while the cat phase-flip probability increases linearly. Moving beyond this point would still increase the noise bias, but since the physical error probabilities are only increasing, there would be no accompanying improvement in logical performance. These intersection points therefore specify the maximum beneficial bias achievable in the cat-transmon architecture, and they are indicated as the black dashed line in Fig. 7 of the main text.

3. Analytical expressions for storage dephasing

In this section, we derive the expression for the storage dephasing rate, Eq. (11) of the main text, and apply it to calculate the dephasing rate for different types of noise with and without the storage echo.

Equation (11) follows directly from Eq. (18) of Ref. [57], with the modification that the dephasing rate is enhanced in the case of the displaced cat in the CR_X gate, owing to the larger photon number of the displaced coherent state $|2\alpha\rangle$. In particular, the enhancement factor is given by the squared mean photon number of this state $(\langle 2\alpha|a^\dagger a|2\alpha\rangle)^2 = 16|\alpha|^4$. Multiplying the result of Ref. [57] by this enhancement factor yields Eq. (11), which

we repeat here for convenience:

$$\Gamma_\phi(t) = 16|\alpha|^4 \int_0^\infty \frac{d\omega}{2\pi} S(\omega) \frac{F(\omega t)}{\omega^2}. \quad (\text{A2})$$

We proceed to use this expression to compute the dephasing rate $\Gamma_\phi(t)$ for different dephasing noise spectra, comparing the cases with and without the storage echo. Without an echo, the filter function is given by $F(\omega t) = 2 \sin^2(\omega t/2)$, and with the echo, $F(\omega t) = 8 \sin^4(\omega t/4)$ [57]. We begin with the case of white noise, defined by a constant noise spectrum $S(\omega) = A^2$ for some constant A . Regardless of whether the echo is applied, we obtain

$$\Gamma_\phi^{(\text{white})}(t) = 16|\alpha|^4 A^2 t/2. \quad (\text{A3})$$

The case of $1/f$ noise, or pink noise, is defined by $S(\omega) = 2\pi A^2/\omega$. Without the echo, this yields

$$\Gamma_\phi^{(1/f, \text{noecho})}(t) = 16|\alpha|^4 (A t)^2 \log \frac{e^{3/2} - \gamma}{\omega_{\text{ir}} t}, \quad (\text{A4})$$

where γ is the Euler constant and ω_{ir} is the infrared cutoff (see, e.g., Ref. [58]). With the echo, we obtain a different result,

$$\Gamma_\phi^{(1/f, \text{echo})}(t) = 16|\alpha|^4 (A t)^2 \log 2. \quad (\text{A5})$$

At time scales that are short relative to the cutoff, $\omega_{\text{ir}} t \ll 1$, the echo can thus provide a significant reduction in the dephasing rate, though the magnitude of this improvement depends on the precise value of the cutoff.

These expressions are plotted alongside corresponding numerical results in Fig. 3, showing good agreement. In the simulations, we have chosen values for the constants A such that $\Gamma_\phi(t = 1/\kappa_\phi) = -16|\alpha|^2$. This way, for both white and $1/f$ noise, the time it takes for the single-photon coherence, $|n=0\rangle \langle n=1|$, to decay by a factor of $1/e$ is the same, and this time defines the decay rate κ_ϕ . To simulate $1/f$ noise, we perform trajectory simulations following the approach in Ref. [58]. In particular, for each choice of noise strength, the dephasing rate is computed from 10^4 individual trajectories, each of length T and down-sampled from a longer single trajectory, such that the infrared cutoff is given by $\omega_{\text{ir}} = 2\pi/(10^4 T)$.

4. Pulse shaping

In this section we describe the shaped pulses used to improve the performance of the CR_X gate. We consider four different types of selective pulses: standard DRAG pulses, exact one-component DRAG pulses, approximate two-component DRAG pulses, and semiclassical two-component DRAG pulses. We define each pulse type in turn below, and the CR_X coherent error for each pulse is plotted in Fig. 10. These




FIG. 10. Comparison of CR_X error with different selective pulse shapes. Lines are colored according to whether they are designed to suppress one (orange), two (green), or three (red) undesired transitions. Here, for numerical expediency, the CR_X coherent error is computed by averaging the infidelity over a small handful of initial cat-transmon states. In contrast, in Fig. 4 the error is computed from the full CR_X Pauli error channel, and hence the values in this figure differ slightly from the corresponding ones in Fig. 4. In the one-, two-, and three-component pulses the Δ_i are numerically optimized over a finite domain, $|\Delta_i| < 3n \times (2|\alpha|)^2 \chi$, for an n -component pulse.

pulse shapes are all derived directly following the approach in Ref. [61]. For numerical simulations of CR_X gates in the main text, we optimize over these four pulse types. Separately, we note that all pulse shapes in Fig. 10 exhibit reduced coherent error at larger $|\alpha|^2$, a consequence of the fact that the transmon dispersive shifts from the $|0\rangle$ and $|2\alpha\rangle$ states can be better resolved at larger $|\alpha|^2$.

We consider a transmon drive of the form $H_{\text{drive}} = \Omega(t)/2 |e\rangle \langle g| + \text{H.c.}$ The validity of this two-level approximation is discussed in Appendix E. For our baseline pulse shapes, we consider a family of truncated Gaussian pulses [62]

$$\Omega_0(t) = \exp\left(\frac{(t - T_{\text{sel}}/2)^2}{2\sigma^2}\right) - \exp\left(\frac{(T_{\text{sel}}/2)^2}{2\sigma^2}\right), \quad (\text{A6})$$

where T_{sel} denotes the selective pulse duration, we fix $\sigma = T_{\text{sel}}$, and the first $m - 1$ derivatives of the pulse vanish at $t = 0, T_{\text{sel}}$.

Given a base pulse shape Ω_0 (e.g., a simple Gaussian pulse), we define the standard DRAG pulses as

$$\Omega(t) = \Omega_0(t) + ip_1 \dot{\Omega}_0(t), \quad (\text{A7})$$

along with a time-dependent pulse detuning

$$\delta(t) = p_2 \Omega^2(t) + p_3, \quad (\text{A8})$$

where, for $\Omega_0(t)$, we take the $m = 1$ truncated Gaussian pulse. The parameters p_i are numerically optimized to

maximize the performance of the pulse (see, e.g., Refs. [93, 94]).

The exact one-component DRAG pulse exactly implements a given rotation at a desired transition without inducing any rotation at a single undesired transition detuned by Δ_1 . That is, consider a simplified CR_X -like scenario where the storage mode state is either vacuum $|0\rangle$ or a single Fock state $|n\rangle$ (rather than the coherent state $|2\alpha\rangle$). In this simplified scenario, the goal is to drive the transmon nontrivially when the storage is in $|0\rangle$, and to not drive the transmon when the storage is $|n\rangle$. We thus have one desired transition at frequency ω_q , and one undesired transition at $\omega_q + \Delta_1$, where $\Delta_1 = n\chi$. This simplified problem is solved exactly (i.e., a perfectly selective drive is implemented) by the pulse shape

$$\Omega(t) = \Omega_0(t) + i \frac{\Omega_0(t)\Delta(t) - \dot{\Delta}(t)\dot{\Omega}_0(t)}{\Omega_0(t)^2 + \Delta(t)} \quad (\text{A9})$$

with the time-dependent pulse detuning

$$\delta(t) = -\text{Re } \Omega(t) \tan \left[\int_0^t dt' \text{Im } \Omega(dt') \right], \quad (\text{A10})$$

implementing a rotation by angle

$$\theta = \int_0^T dt \text{Re } \Omega(t) \sec \left[\int_0^t dt' \text{Im } \Omega(t') \right]. \quad (\text{A11})$$

See Eq. (5.4) of Ref. [61] [which we note drops the $\propto \dot{\Delta}(t)$ term in the numerator]. In the above, $\Delta(t) = \delta(t) + \Delta_1$, and for $\Omega_0(t)$, we take the $m = 1$ truncated Gaussian pulse. In the context of the CR_X gate, this pulse shape is not perfectly selective, as there is more than just a single undesired transition [see Fig. 4(a)]. Despite this, this pulse shape can perform well in practice when Δ_1 is treated as an optimization parameter. In fact, in Fig. 10 we see that these exact one-component DRAG pulses with only the one optimization parameter achieve very similar performance to the standard DRAG pulses with three optimization parameters.

The approximate two-component DRAG pulse exactly implements a given rotation at a desired transition, while approximately suppressing rotation at two undesired transitions with detunings Δ_1, Δ_2 . The pulse is defined as

$$\begin{aligned} \Omega(t) = \Omega_0(t) - i \frac{[E_1^2(t) - E_2^2(t)]\Delta_1\Delta_2}{E_1^2(t)E_2^2(t)(\Delta_1 - \Delta_2)} \dot{\Omega}_0(t) \\ - \frac{E_1^2(t)\Delta_2 - E_2^2(t)\Delta_1}{E_1^2(t)E_2^2(t)(\Delta_1 - \Delta_2)} \ddot{\Omega}_0(t). \end{aligned} \quad (\text{A12})$$

In the above, for $\Omega_0(t)$, we take the $m = 2$ truncated Gaussian pulse, $E_i^2(t) = \Omega_0(t)^2 + \Delta_i^2$, and the time-dependent

pulse detuning and rotation angle are again given by Eqs. (A10) and (A11); cf. Eq. (5.15) of [61], where here we have simplified by neglecting terms $\propto \dot{E}_i(t)$, which is justified in the far-detuned limit $\Delta_i \gg \Omega_0$. In the context of the CR_X gate, Δ_1, Δ_2 are treated as optimization parameters.

Finally, the semiclassical two-component DRAG pulses are obtained by taking the $\Omega_0(t)/\Delta_i \rightarrow 0$ limit of the approximate two-component DRAG pulses, yielding

$$\Omega(t) = \Omega_0(t) - i \left(\frac{1}{\Delta_1} + \frac{1}{\Delta_2} \right) \dot{\Omega}_0(t) - \frac{1}{\Delta_1\Delta_2} \ddot{\Omega}_0(t) \quad (\text{A13})$$

with time-dependent pulse detuning and rotation angle given by Eqs. (A10) and (A11). In Fig. 10 we see that the optimized approximate two-component DRAG and semiclassical two-component DRAG pulses significantly outperform the optimized standard DRAG and exact one-component DRAG pulses. The CR_X gate performance when using the semiclassical two-component DRAG pulse is plotted as the orange “standard DRAG” line in Fig. 4.

Finally, for the sake of completeness, we note that the two-component pulses can be extended to more than two undesired transitions. The full expression for the three-component approximate pulse is quite involved, but the corresponding three-component semiclassical pulse takes the comparatively simple form

$$\begin{aligned} \Omega(t) = \Omega_0(t) - i \left(\frac{1}{\Delta_1} + \frac{1}{\Delta_2} + \frac{1}{\Delta_3} \right) \dot{\Omega}_0(t) \\ - \frac{\Delta_1 + \Delta_2 + \Delta_3}{\Delta_1\Delta_2\Delta_3} \ddot{\Omega}_0(t) + i \frac{1}{\Delta_1\Delta_2\Delta_3} \frac{d^3\Omega_0(t)}{dt^3}, \end{aligned} \quad (\text{A14})$$

where, for $\Omega_0(t)$, we take the $m = 3$ truncated Gaussian pulse. However, in Fig. 10 we see that these three-component pulses do not offer any further improvement over the relevant range of parameters.

5. Additional applications of the CR_X gate

Here we show that the CR_X gate naturally enables other useful capabilities, including high-fidelity, bias-preserving cat Z and cat-cat CZ rotations, as well as single-shot cat Z -basis readout.

Cat Z rotations can be implemented by exploiting the phase accumulation associated with the CR_X gate’s selective pulses. For example, when the transmon is initialized in $|+\rangle$, the CR_X unitary

$$\text{CR}_X = |\alpha\rangle\langle\alpha| \otimes I + |-\alpha\rangle\langle-\alpha| \otimes R_X(\pi) \quad (\text{A15})$$

acts trivially on the transmon. Since the transmon and cat are not entangled, the transmon can be discarded, and the net effect of the CR_X gate is to enact an $S = \text{diag}(0, i)$ gate on the cat. The S gate is equivalent to a Z rotation by $\pi/2$. In fact, it is possible to implement arbitrary cat Z rotations in this manner. Tuning the selective pulses so that the transmon starts and ends in the same state while acquiring a phase ϕ (see, e.g., Refs. [95,96]) is equivalent to enacting a cat Z rotation by ϕ .

One can similarly implement cat-controlled cat Z rotations via a generalized CR_X gate. Consider two cats dispersively coupled to the same transmon. Displacing both cats and applying selective pulses enables transmon rotations conditioned on both cats being in their logical $|1\rangle$ states. Tuning the selective pulses so that the transmon accumulates a ϕ phase is equivalent to the multicat rotation $\text{diag}(1, 1, 1, e^{i\phi})$, equivalent to a CZ gate for the choice of $\phi = \pi$. This approach easily generalizes to $n > 2$ cats.

What is remarkable about the ability of the CR_X gate to implement cat Z -type rotations is that this implementation is exponentially noise biased (since the CR_X gate is itself exponentially noise biased), yet the infidelity is not subject to the unfavorable $\sqrt{\kappa_1/\kappa_2}$ scaling of prior bias-preserving approaches to Z rotations based on the Zeno effect [5,35]. High-fidelity and high-bias Z rotations are thus possible with this approach. Gautier *et al.* [97] similarly proposed ways to circumvent the unfavorable infidelity scaling. In comparison to the proposals in that work, this CR_X -based approach is practically appealing in the context of Z rotations, but it is not obviously generalizable to enable a high-fidelity cat-cat CX gate.

Finally, the CR_X gate can also be used to enable single-shot Z -basis readout of cat qubits with readout error that decreases exponentially with $|\alpha|^2$. Readout is performed by initializing the transmon in $|g\rangle$, applying the CR_X gate, then reading out the state of the transmon. An outcome of $|g\rangle$ indicates the cat logical $|0\rangle (= |+\alpha\rangle)$, and an outcome of $|e\rangle$ indicates cat logical $|1\rangle (= |-\alpha\rangle)$. A simplified version of this readout was employed in Ref. [41]. This readout procedure does not corrupt the cat's logical Z information, and because of this nondemolition property the readout can be repeated to improve performance. (Note that the dissipative cat stabilization can be turned on during the transmon readout, i.e., between CR_X gates, so that leakage out of the code space does not accumulate with increasing repetitions.) In particular, repeating the procedure and taking the majority vote allows one to exponentially suppress contributions to the overall cat readout error associated with transmon ancilla errors [21], at the cost of a longer effective readout duration. With sufficiently many repetitions, the overall readout error is thus only limited by the cat's exponentially small (in $|\alpha|^2$) intrinsic bit-flip probability.

APPENDIX B: PULSED STABILIZATION AND LEAKAGE

Our circuit-level simulations of QEC performance use standard Pauli error channels (i.e., without leakage) to model noise in the gates. These Pauli error channels are extracted from Hamiltonian-level simulations of the gates followed by a fictitious recovery operation that maps any leaked population back into the code space in accordance with the cat's engineered dissipation. This approximate, implicit treatment of leakage is justified under two conditions: (i) engineered dissipation is applied frequently enough during syndrome extraction cycles so that leakage out of the cat code space does not accumulate from one cycle to the next; and (ii) the effects of finite leakage population on syndrome measurement are negligible. In this appendix, we show that both assumptions hold in the parameter regimes of interest in this work.

1. Pulsed stabilization

In this section we quantify the minimum amount of pulsed stabilization that is sufficient to prevent accumulating leakage out of the code space for an isolated cat. (Later, in Appendix B 2, we also quantify when pulsed stabilization is sufficient to prevent leakage from degrading logical performance during QEC.) Figure 11 shows results from master equation simulations of a cat subject to single-photon loss, where stabilization is repeatedly pulsed on and off. In particular, the cat mode is simultaneously subject to loss (dissipator $\kappa_1 D[a]$) and stabilization (dissipator $\kappa_2 D[a^2 - \alpha^2]$) for duration T_{on} , then only loss for a duration T_{off} . This on-off pulsed stabilized cycle, with total duration $T = T_{\text{on}} + T_{\text{off}}$, is repeated multiple times. The signature of whether leakage is accumulating is whether the cat has a well-defined bit-flip rate, determined as follows. After each cycle, we compute the probability $p_X(t)$ of a cat bit flip by evaluating the conserved quantities defined in Ref. [5] (we emphasize that this evaluation is not a part of our master equation simulations and does not affect the dynamics of the system; it is just a means of analyzing the simulation data). We then fit $p_X(t)$ to a power law as a function of dimensionless time, $t/T = \text{number of cycles}$. A well-defined bit-flip rate—equivalent to a constant bit-flip probability per cycle—corresponds to a power-law fit exponent of approximately 1 [i.e., $p_X(t) \propto t$].

Example fits are shown in Fig. 11(a), and fitted exponents are shown as a function of the duty cycle (T_{on}/T) and ratio of the confinement and loss rates ($\kappa_2|\alpha|^2/\kappa_1$) in Fig. 11(b). The results show that if the stabilization is too weak or the duty cycle too small, the bit-flip probability grows nonlinearly in time, indicating accumulating leakage. In contrast, when both the stabilization rate and duty cycle are sufficiently high, then the cat exhibits a well-defined (i.e., constant) bit-flip rate, indicating no leakage




FIG. 11. Efficacy of pulsed stabilization. (a) Cumulative bit-flip probability scaling. The cumulative probability of a bit flip, $p_X(t)$, is plotted (symbols) as a function of time and fit (dashed lines) to a power law. The different colors correspond to different choices of T_{on}/T and $\kappa_2|\alpha|^2/\kappa_1$. (b) Fitted power-law exponents. The data in (a) correspond to circular and square markers. The black dashed line is $(\kappa_2|\alpha|^2/\kappa_1)(T_{\text{on}}/T) \geq 25$; the good agreement of this curve with the contour of the plot indicates that the effect of pulsing can be well approximated via a reduced stabilization rate: $\kappa_2 \rightarrow \kappa_2(T_{\text{on}}/T)$. In the blue region of parameter space, the bit-flip probability increases approximately linearly with time, indicating well-stabilized cats.

accumulation. To identify the regime of well-stabilized cats, we compare the single-photon loss rate κ_1 against the effective confinement rate, $\propto \kappa_2|\alpha|^2 T_{\text{on}}/T$, where this effective confinement differs from the usual confinement rate, $\propto \kappa_2|\alpha|^2$ [17], by a factor of the duty cycle in order to account for the pulsing. We find that $(\kappa_2|\alpha|^2/\kappa_1)(T_{\text{on}}/T) \gtrsim 25$ is sufficient to guarantee a well-stabilized cat.

For context, consider the example of $\chi/2\pi = 2$ MHz, corresponding to CX and CR_X gate durations of $\pi/\chi = 250$ ns. Assuming the same duration for transmon readout and reset, the pulsed stabilization duty cycle is $T_{\text{on}}/T = 1/5$ during surface code error correction (longer transmon readout and reset duration would only yield a larger duty cycle). Assuming loss rates in the vicinity of the surface code “threshold” identified in Sec. V A, $\kappa_1 = (100 \mu\text{s})^{-1}$, and $|\alpha|^2 = 4$, we find that $\kappa_2/2\pi \gtrsim 50$ kHz is required for a well-stabilized cat. This requirement has already been exceeded by current experiments [20,22,37], and the requirement only becomes more lenient as κ_1 decreases or $|\alpha|^2$ increases. Stabilizing the cats only during the ancilla

transmons’ readout and reset is thus a realistic way to prevent accumulating leakage, and we affirm this result through simulations of repeated QEC rounds in the next subsection.

That said, additional stabilization steps could always be inserted between cat-transmon entangling gates if needed, at the price of longer syndrome extraction rounds. The impact of these additional stabilization steps on logical memory overhead can be estimated as follows. Suppose that additional stabilization steps of total duration T_s are inserted at some point during the entangling gates into a syndrome extraction cycle, increasing the cycle length to $T'_{\text{cycle}} = T_{\text{cycle}} + T_s$. This addition increases the cat Z error probability per round p_Z by a factor of approximately $T'_{\text{cycle}}/T_{\text{cycle}}$, since the cat Z errors are dominated by single-photon loss. At the same time, one can assume that the bias η does not increase, since some cat X errors are caused by mechanisms that also increase in proportion to $T'_{\text{cycle}}/T_{\text{cycle}}$, such as transmon heating. (This assumption is likely overly pessimistic: on the one hand, other error mechanisms that can propagate cat X errors, e.g., CR_X selective pulse errors, would not increase with T_s ; on the other hand, the only mechanism that would increase superlinearly, transmon double-decay errors, contributes negligibly relative to transmon-transmon heating in the high-coherence regime.) The net effect of this additional stabilization can thus be approximately captured by increasing p_Z by a factor $T'_{\text{cycle}}/T_{\text{cycle}}$ and leaving η fixed in Figs. 7 and 8.

2. Leakage

The above analysis shows that leakage does not accumulate under sufficiently strong and frequent pulsed stabilization, but even in such cases there will generally be a nonzero steady-state leakage population. That is, though the leaked population will fluctuate during a syndrome measurement round—growing while the cat stabilization is turned off and shrinking when it is turned on—the average leaked population will stabilize to some nonzero value after sufficiently many rounds. In principle, such leakage could degrade our ability to accurately measure error syndromes even without degrading the logical information stored in the cat qubit. In the following, we thus quantify the impacts of leakage on syndrome measurement performance, showing in fact that both the CX and CR_X gates are robust to leakage, such that impacts of leakage on syndrome measurements are minimal.

A useful tool for conceptualizing leakage and its impacts on cats is the shifted Fock basis [13]. The shifted Fock basis states are defined as

$$|\phi_{n,\pm}\rangle = \frac{1}{\sqrt{2}} [D(\alpha) \pm (-1)^n D(-\alpha)] |n\rangle, \quad (\text{B1})$$

where $|n\rangle$ denotes a Fock state and $D(\alpha)$ is the displacement operator. The states $|\phi_{n,+}\rangle$ and $|\phi_{n,-}\rangle$ have even and




FIG. 12. Impacts of cat leakage on syndrome measurement. (a) Shifted Fock basis. In this basis, the Hilbert space of a bosonic mode is represented as the tensor product of two components: a two-level system (encoding the parity) and a gauge mode. In the case of cat qubits, the two-level system hosts the logical information, and excitations in the gauge mode correspond to leakage. (b) Circuits simulated with cat leakage. Cat modes (black rails) are initialized in states with nonzero leakage, then entangled with the transmon (green), which is subsequently measured. To make simulation of the stabilizer circuits numerically tractable, each cat mode is initialized (traced out) immediately before (after) the corresponding entangling gate; this does not affect measurement outcomes. (c) Measurement error probabilities of the four circuits in (c) as a function of leakage. In all cases, the measurement error probabilities are far lower than the worst-case scenario of random measurement outcomes for leakage population (i.e., $n_{th}/2$ for CX and CR_X circuits, $4n_{th}/2$ for stabilizer circuits), indicating resilience to leakage.

odd photon-number parity, respectively, with $|\phi_{0,+}\rangle$ and $|\phi_{0,-}\rangle$ equivalent to the $|+\rangle_C$ and $|-\rangle_C$ cat code states. (Note that the $|\phi_{n,\pm}\rangle$ are not orthonormal; the Gram-Schmidt procedure can be applied to these states separately within each parity sector to construct a proper orthonormal basis; see Appendix C of Ref. [13] for details.) Viewed through this basis, the original Hilbert space of the bosonic mode effectively factorizes into two components: a two-level system encoding the parity ($\pm$), and a separate “gauge mode,” as depicted in Fig. 12(a). In this picture, the cat code space corresponds to the vacuum state of the gauge mode, and leakage corresponds to excitations in the gauge mode. Additionally, the engineered dissipation manifests as the loss on the gauge mode, effectively cooling it back towards the vacuum state.

To quantify the impacts of leakage on syndrome measurement, we first create a family of states with fixed logical information but varying degrees of leakage. This is done using the shifted Fock basis; we fix the logical state $|\psi_L\rangle$ while preparing the gauge mode in a thermal state $\rho(n_{th})$ with variable average occupation number n_{th} . This occupation number reflects both the amount and extent of leakage. The choice of a thermal state is conceptually motivated by the competition between the “heating” and “cooling” of the gauge mode due to leakage errors and engineered dissipation, respectively, but the qualitative conclusions below are not specific to this choice.

Taking $|\psi_L\rangle \langle\psi_L| \otimes \rho(n_{th})$ as initial states for the cat mode, we then simulate cat-transmon gates to assess the

extent to which these gates can successfully map the logical information ψ_L onto the transmon, despite the presence of leakage $n_{th} > 0$. Specifically, for the CX gate, we take $|\psi_L\rangle \in \{|+\rangle, |-\rangle\}$ and initialize the transmon in $|+\rangle = (|g\rangle + |f\rangle)/\sqrt{2}$ so that the transmon state should ideally be flipped from $|+\rangle \rightarrow |-\rangle = (|g\rangle - |f\rangle)/\sqrt{2}$ depending on which $|\psi_L\rangle$ is chosen. Analogously, for the CR_X gate, we take $|\psi_L\rangle \in \{|0\rangle, |1\rangle\}$ and initialize the transmon in $|g\rangle$, so that it should be flipped from $|g\rangle \rightarrow |e\rangle$ depending on the chosen $|\psi_L\rangle$. We quantify the performance via the probability of obtaining the ideal outcome when subsequently measuring the transmon state. These simulations are performed either with just a single cat mode and gate, or with a sequence of four modes and gates, exactly as in a syndrome measurement; see Fig. 12(b). In the latter case, the error directly corresponds to the probability of a syndrome measurement error induced by cat leakage. Decoherence is not included in these simulations in order to isolate the impact of leakage.

Figure 12(c) shows the measurement error as a function of the leakage. These are compared against the worst-case scenario where any leaked population results in a random measurement outcome. In all cases, the measurement errors are significantly lower than this worst-case error, indicating intrinsic robustness of the cat-transmon gates to leakage. In the case of the CX gate, we see no appreciable error with leakage. This robustness is naturally expected: the cat-transmon CX gate in fact enacts a parity measurement of the cat mode, so that regardless of any leakage this




FIG. 13. Steady-state leakage. (a) Repeated QEC rounds from the perspective of the data cat qubit. We perform master equation simulations of the circuit shown, where the data cat (black rail) is involved in four cat-transmon entangling gates with four different ancilla transmons (green rails), then subject to a dissipative stabilization step, as in QEC. To make the simulations numerically tractable, each ancilla transmon is initialized (traced out) immediately before (after) the corresponding gate. (b) Maximum gauge mode excitation. The maximum gauge mode excitation at any point during $n = 10$ QEC rounds is plotted as a function of $|\alpha|^2$ and the decoherence parameter q . For all points, the relative change in the maximum gauge mode excitation from round 9 to round 10 is small ($< 1\%$), indicating that $n = 10$ rounds is sufficient to capture steady-state leakage. (c) Gauge mode population over time. For the example point indicated in red in (b), we plot the gauge mode excitation as a function of time, where the excitation is evaluated after each of the five operations in a QEC round, and the dashed line indicates the maximum value.

gate successfully maps the parity-encoded logical information onto the transmon. In contrast, the CR_X gate does exhibit some increased error with leakage, but still far less than the worst-case scenario. The robustness in this case can be attributed to the finite bandwidth of the selective pulses. As shown in Fig. 4, the CR_X gate's selective pulses have significant spectral content not just at the bare transmon frequency but also at nearby frequencies detuned by a few multiples of χ . As a result, these pulses can still largely enact the desired rotations even if the displaced cat mode is not exactly in its vacuum state (due to leakage).

Beyond the qualitative observation that the cat-transmon gates are robust to leakage, we also note that the leakage-induced infidelities are unlikely to limit the performance of syndrome measurements in practice. In particular, even for the extremely large leakage populations considered here of up to 10%, the associated syndrome measurement error is only $\lesssim 2\%$, comparable to the 1% probability of transmon state-preparation and measurement (SPAM) error assumed in the main text. In practice, we expect much lower steady-state leakage in the subthreshold regime, as shown in Fig. 13. We quantify the steady-state leakage by performing master equation simulations of repeated QEC rounds, subject to the noise model of Table I (with $\kappa_2 = \chi/10$ and stabilization duration $\pi/|\chi|$, as in the main text). Each QEC round includes four cat-transmon gates (that increase leakage), followed by a period of stabilization (that decreases leakage). We plot the maximum gauge mode excitation across repeated QEC rounds as a function of $|\alpha|^2$ and the decoherence parameter q . For context, the threshold from Fig. 6 ($q = 5 \times 10^{-4}$ at $|\alpha|^2 = 6$) is indicated by the red box. At this threshold,

the maximum gauge mode excitation is already $< 10^{-2}$, and the gauge mode excitation only decreases as the value of q is decreased below threshold. In the subthreshold regime that we focus on in this work, with high-fidelity gates and leakage populations $\ll 10\%$, the contribution of leakage to syndrome measurement error is thus expected to be negligible relative to that from transmon SPAM error.

As a final note, it is interesting to compare the results of Figs. 11 and 13. The former quantifies the minimum amount of pulsed stabilization required to prevent accumulating leakage in an isolated cat, while the latter quantifies leakage in the context of QEC, where cat-transmon gates act as an additional source of leakage beyond the cat's intrinsic decoherence mechanisms. As expected, the maximum gauge mode populations are large when the requirement for leakage nonaccumulation of an isolated cat, $(\kappa_2|\alpha|^2/\kappa_1)(T_{\text{on}}/T) \geq 25$ (see Fig. 11), is not satisfied. In particular, for the parameters of Fig. 13, we have $(\kappa_2|\alpha|^2/\kappa_1)(T_{\text{on}}/T) = |\alpha|^2/50q$. The requirement is not satisfied for $q = 10^{-2}$ across the full range of $|\alpha|^2$ shown, so, as expected, large gauge mode excitation is observed at this value of q . In contrast, for $q = 5 \times 10^{-4}$, the requirement is satisfied across the full range of $|\alpha|^2$ shown. The fact that large gauge mode excitation is nevertheless observed at small $|\alpha|^2$ for $q = 5 \times 10^{-4}$ is due to additional leakage induced by the cat-transmon gates (especially the CR_X gate, where cat leakage can be induced by selective pulse errors that are larger at smaller $|\alpha|^2$; see Fig. 4). As $|\alpha|^2$ grows, this additional leakage is suppressed, and the confinement rate $\kappa_2|\alpha|^2$ grows, resulting in lower gauge mode excitation.

TABLE II. Different single-parameter noise models used to obtain the results in Fig. 14. Model 1 is the same as in Table I of the main text. Relative to this, Model 2 reduces all transmon decoherence rates by a factor of 3, and Model 3 holds the transmon decoherence rates constant (i.e., independent of q). The ratios between transmon loss and dephasing and heating rates are maintained, so all models assume an effective thermal population of $\gamma_\uparrow/\gamma_\downarrow = 0.5\%$ and $T_1 = T_2$ for the transmon.

Decoherence source	Model 1	Model 2	Model 3
Storage loss (κ_1/χ)	q	q	q
Transmon loss ($\gamma_\downarrow/\chi$)	$3q$	q	5×10^{-4}
Storage dephasing (κ_ϕ/χ)	$0.01q$	$0.01q$	$0.01q$
Transmon dephasing (γ_ϕ/χ)	$1.5q$	$0.5q$	2.5×10^{-4}
Transmon heating ($\gamma_\uparrow/\chi$)	$0.015q$	$0.005q$	2.5×10^{-6}

APPENDIX C: ALTERNATIVE NOISE MODELS

In order to investigate the dependence of our results on the choice of noise model, in this appendix we consider variants of the noise model of Table I of the main text. The three noise models we consider are given in Table II. Model 1 is the same as that in Table I and is repeated here for ease of comparison. Recall that this model assumes higher coherence for the storage mode than the transmon, which is roughly representative of 3D devices, as well as the specific 2D device of Ref. [41]. In contrast, Model 2 assumes that the storage mode and transmon have comparable coherence (storage $T_1 =$ transmon T_1), which may be a more reasonable assumption for 2D devices generally. Separately, in order to provide insights into the limitations imposed by transmon decoherence, in Model 3 the transmon decoherence rates are held fixed, unlike in Models 1 and 2 where they are varied together with the storage decoherence rates.

In Fig. 14, we plot the achievable noise bias (η) and Z-type error probability (p_Z) for these three different noise models, as in Fig. 7 of the main text. Starting with Model 1, the corresponding plot in Fig. 14 is the same as Fig. 7, except that we have grayed out the region of parameter space that lies beyond the maximum beneficial bias. This is done both to emphasize that operating in this regime of parameter space, while technically possible, only serves to degrade logical performance (see Appendix A 2), and also to visually highlight the contrast with the other noise models. Indeed, for Model 2, the main difference relative to Model 1 is that larger biases are accessible. The increase in the accessible bias follows simply from the lower bound on the cat bit-flip probability of the CX gate [Eq. (15)]. Since transmon errors during the CX gate can induce cat bit flips, higher transmon coherence relative to storage mode coherence allows for a larger accessible bias. For example, with transmon and storage T_1 both of the scale of milliseconds and χ of the scale of megahertz (corresponding to q in the range 10^{-5} – 10^{-4}), biases of 10^4 are achievable. This is a

roughly a factor of 3 larger than the maximum achievable bias for Model 1 at the same q , consistent with the fact that the transmon heating rate is 3 times lower in Model 2.

Model 3 holds the transmon decoherence rates fixed, so the corresponding plot in Fig. 14 highlights the limitations on p_Z and η imposed by transmon decoherence. Specifically, the fixed values of $\gamma_\downarrow/\chi = 5 \times 10^{-4}$ and $\gamma_\phi/\chi = 2.5 \times 10^{-4}$ correspond to per-gate transmon decay and dephasing probabilities of the order of 10^{-3} . Because transmon errors can propagate through the CR_X gate to cat Z errors, the achievable p_Z is limited to a similar value. This limitation is indicated by the yellow dash-dot line in the plot, making clear that after a point improving storage coherence (decreasing q) yields no improvement in p_Z , unless transmon coherence is also improved. At the same time, with fixed transmon decoherence, the maximum beneficial bias scales linearly with p_Z (black dashed line). This scaling results from the fact that the fixed transmon decoherence fixes the minimum achievable cat bit-flip probability of the CX gate [Eq. (15)], and with fixed bit-flip probability we have $\eta \propto p_Z$. Both this scaling, and a visual comparison of the achievable bias with Models 1 and 2 to that with Model 3, make clear that it is necessary to proportionately increase the transmon coherence in order to maintain a large noise bias as the storage mode coherence is increased.

APPENDIX D: EXTRACTING GATE ERROR CHANNELS

We numerically extract Pauli error channels describing CX and CR_X gates as follows. The gate error $\mathcal{E}$ channel is defined such that the application of a noisy gate is equivalent to the ideal gate followed by application of the error channel $\mathcal{E}$,

$$\mathcal{U}_{\text{noisy}}(\rho) = \mathcal{E}(U_{\text{ideal}}\rho U_{\text{ideal}}^\dagger). \quad (\text{D1})$$

We can characterize the error channel $\mathcal{E}$ via its Pauli transfer matrix, R . We refer the reader to Ref. [98] for a description of this and the other representations of quantum channels discussed below. The Pauli transfer matrix elements are given by

$$R_{ij} = \frac{1}{d} \text{tr}[P_i \mathcal{E}(P_j)], \quad (\text{D2})$$

where $d = 4$ is the Hilbert space dimension of two qubits, and the P_i denote (two-qubit) Pauli operators. This representation is convenient to extract from simulations, as one calculates R_{ij} by initializing the system in $U_{\text{ideal}}^\dagger P_j U_{\text{ideal}}$, applying the noisy gate operation $\mathcal{U}_{\text{noisy}}$, and computing the overlap with P_i ,

$$R_{ij} = \frac{1}{d} \text{tr}[P_i \mathcal{U}_{\text{noisy}}(U_{\text{ideal}}^\dagger P_j U_{\text{ideal}})]. \quad (\text{D3})$$




FIG. 14. Achievable noise bias, η , and Z-type error probability, p_Z , for the different noise models in Table II. Gray areas indicate regions of parameter space that are either inaccessible or where operation is not beneficial to logical memory performance. Black dashed lines indicate the maximum beneficial bias (see Appendix A 2). The yellow dash-dot line in the Model 3 plot indicates the limit imposed on p_Z by the fixed transmon decoherence. This line is obtained by computing p_Z in the limit of zero storage decoherence (i.e., $q \rightarrow 0$ in Model 3).

For error-correction simulations, a more convenient representation is the chi-matrix (or process matrix) representation,

$$\mathcal{E}(\rho) = \sum_{ij} \chi_{ij} P_i \rho P_j^\dagger. \quad (\text{D4})$$

To enable efficient error-correction simulations, we make the Pauli twirling approximation and neglect the off-diagonal terms in χ . The remaining diagonal terms encode the probabilities of different Pauli errors, and they are related to the Pauli transfer matrix via a Walsh Hadamard transform,

$$R_{ii} = \frac{1}{d} \sum_{j,k} \chi_{j,k} \text{tr}[P_i P_j P_i P_k] = \sum_j \chi_{jj} (-1)^{\langle ij \rangle}, \quad (\text{D5})$$

where $\langle i, j \rangle$ is defined via $P_i P_j = (-1)^{\langle ij \rangle} P_j P_i$. We can thus compute the diagonal chi-matrix elements (Pauli error probabilities) by applying an inverse Walsh Hadamard transform to the diagonal Pauli transfer matrix elements obtained from gate simulations.

In noisy implementations of the cat-transmon gates, errors can prevent the system from returning to the code space. That is, the storage mode state could lie outside the cat code space (spanned by $|0/1\rangle_C$), and, in the case of the CX gate, the final transmon state could lie outside the gate's $|g\rangle, |f\rangle$ code space. This possible leakage must be accounted for to obtain a normalized Pauli error channel for the gate, and we do so by augmenting the noisy gate channel with a recovery operation, i.e., $\mathcal{U}_{\text{noisy}} \rightarrow \mathcal{R} \circ \mathcal{U}_{\text{noisy}}$, where the recovery operation $\mathcal{R}$ returns the system to the code space as described below.

For the storage mode, $\mathcal{R}$ implements the ideal dissipative recovery map, which is equivalent to applying the

dissipative stabilization $\mathcal{D}[a^2 - \alpha^2]$ for infinite time. We implement this map by computing the conserved quantities described in Ref. [5]. This choice of recovery operation is justified by the periodic application of the dissipative stabilization during error-correction cycles (see Fig. 1), which suffices to prevent the accumulation of leakage outside the code space, as shown in Appendix B.

For the transmon, $\mathcal{R}$ acts differently depending on whether the transmon is operated in its $|g\rangle, |e\rangle$ subspace (CR_X gate) or its $|g\rangle, |f\rangle$ subspace (CX gate). In the case of the CR_X gate, $\mathcal{R}$ acts trivially as there is no leakage outside the code space. In the case of the CX gate, $\mathcal{R}$ returns any leaked population in $|e\rangle$ to $|f\rangle$. This is the only choice for $\mathcal{R}$ that accurately captures the resilience of the CX gate to single transmon T_1 decay events; were any $|e\rangle$ population mapped to $|g\rangle$, single decay events ($|f\rangle \rightarrow |e\rangle$) would effectively be mapped to double-decay events ($|f\rangle \rightarrow |e\rangle \rightarrow |g\rangle$), and the moderately noise-biased nature of the gate would not manifest in the Pauli error channel.

It is worth highlighting that this approach—computing gates' Pauli error channels by applying a recovery $\mathcal{R}$ after the simulation of a single noisy gate—does not capture the possible impacts of leakage propagating between subsequent gates. While it is true that the cats' dissipative stabilization prevents the accumulation of leakage, multiple gates are applied in the stretches between dissipative stabilization (see Fig. 1), so leakage in an earlier gate can in principle impact the performance of a later gate before it is corrected by the stabilization. For the cats themselves, neglecting this effect is unlikely to have significant impact; see Appendix B. In particular, small distortions in the input cat states should not propagate through CX or CR_X gates to large enough distortions to induce cat X errors. Indeed, no such damaging propagation was observed in the experiment [41] (see also the fault-tolerance arguments in

Ref. [50]). Were propagation of leakage an issue, the dissipative stabilization could be applied more frequently [40], at the price of a longer error-correction cycle duration; see Appendix B.

Separately, for the transmons, neglecting the propagation of leakage does omit the possibility that an $|f\rangle \rightarrow |e\rangle$ decay in one CX gate could combine with an $|e\rangle \rightarrow |g\rangle$ decay in a later CX gate to induce a cat X error. Treating leakage as such thus modestly underestimates the probability of double-decay-induced cat X errors. This underestimate will not affect results qualitatively, as double-decay-induced cat X errors are accounted for within individual gates. It is also unlikely to have a significant quantitative impact, since, for sufficiently high coherence, the contribution from transmon double decay (scales quadratically in the noise parameter q) will be negligible in comparison to that from transmon heating (scales linearly in q), and the impacts of heating are accurately captured by this approach.

APPENDIX E: SIMULATION ASSUMPTIONS

In this appendix, we enumerate various assumptions that are made implicitly in our simulations of the cat-transmon architecture. For each, we discuss when the assumption is justified and its potential impacts on performance.

Perfect χ matching. Our simulations of the CX gate assume that $\chi_e = \chi_f$. In practice, small mismatches $|\chi_e/\chi_f - 1| \neq 0$ can be tolerated without meaningfully increasing the cat X error rate. This tolerance results from the periodic application of engineered dissipation, which is capable of correcting the small storage mode over- or under-rotations associated with the mismatch, as demonstrated in Ref. [41]. What constitutes a tolerable mismatch generally depends on the engineered dissipation strength and how frequently it is applied. For context, Putterman *et al.* [41] found that mismatches as large as approximately 10% did not meaningfully increase the cat X error rate with stabilization applied only once per error-correction cycle. So long as the mismatch is not too large to significantly degrade the performance of a single CX gate, engineered dissipation can in principle be applied more frequently (e.g., after each CX gate) to ensure that the mismatch does not result in a significantly increased cat X error rate. This added protection, however, would come at the price of a longer error-correction cycle.

No tunable coupler heating. In implementations of the cat-transmon architecture that rely on tunable couplers to implement the tunable cat-transmon dispersive coupling [41], heating of the couplers can potentially induce cat X errors just as can heating of the transmon errors. This coupler-heating-induced cat X error probability is negligible if the coupler's dispersive coupling to the storage modes is sufficiently small in comparison to the transmon's. However, if the coupler's dispersive coupling is

comparable or larger than the transmon's, coupler heating can contribute appreciably to the cat X error probability. The impact of this additional contribution can be effectively assessed from our results simply by augmenting the transmon heating rate, indicating that coupler heating should not qualitatively impact results.

No storage mode self-Kerr nonlinearity. Typically, dispersively coupling a storage mode to a nonlinear transmon qubit causes the storage mode to acquire a self-Kerr nonlinearity. The magnitude of this self-Kerr nonlinearity will depend on how the coupling is engineered (e.g., the proposal of Ref. [80] realizes a dispersive coupling without any storage self-Kerr nonlinearity), but small self-Kerr nonlinearity can be tolerated by the cat-transmon architecture without meaningfully degrading performance. This tolerance again stems predominantly from the periodic application of engineered dissipation, which is capable of correcting small distortions to the cats caused by self-Kerr nonlinearity. For example, Refs. [37,41] found that self-Kerr nonlinearities of the order of a few kilohertz were tolerable. We also note that the CR_X gate's storage echo effectively mitigates the prospect of cat Z errors induced by the storage self-Kerr nonlinearity.

Two-level transmon approximation for the CR_X gate. We only include the lowest two transmon levels in our simulations of the CR_X gate. This approximation is valid, i.e., leakage to higher states like $|f\rangle$ is likely to be negligible, when the CR_X gate's pulses are sufficiently selective. In particular, if the bandwidth of the selective pulse is much smaller than the transmon anharmonicity, the selective pulse should not induce leakage. As a rule of thumb, the bandwidth of the selective pulses must be smaller than $4|\alpha|^2\chi$ for good CR_X performance, so leakage to higher states should be negligible when the anharmonicity is larger than this quantity. For megahertz-scale χ , this assumption is reasonable, as typical transmon anharmonicities are of the order of 200 MHz. The assumption may be violated, however, if larger χ is used in an effort to speed up gates. In this situation, DRAG can be used to mitigate transmon leakage during the CR_X gate's selective pulses.

Perfect coupler toggling. Our simulations assume that the tunable dispersive couplings can be instantly toggled from on to off, with zero dispersive coupling at the off position, whereas in actuality toggling the coupling will require some finite duration and there will be a finite on:off ratio. The primary impact of the finite switching time is to effectively lengthen the duration of the cat-transmon gates. This is especially relevant for the CR_X gate, which ideally involves toggling the coupler on and off four times during the gate sequence. This additional time cost could be offset by operating the CR_X gate at larger $|\alpha|^2$, allowing for faster selective pulses at the price of increased cat Z error. Alternatively, one could leave the coupler on for the entire gate, relying on wide-bandwidth pulses to




FIG. 15. Logical error projections for biased- and unbiased-noise models. (a) Total logical error and fits for square ($d_X = d_Z \equiv d$) surface codes for the unbiased error model. (b) Logical Z error for the biased error model, for fixed d_X and η . (c) Logical X error for the biased error model. Error bars in (a) and (b) indicate 95% confidence intervals; error bars are omitted in (c) for visual clarity.

implement the unselective transmon pulses and storage displacements accurately even in the presence of non-negligible dispersive coupling. The primary impact of the finite on:off ratio is to induce small distortions of the cat, which are again correctable via periodic application of the engineered dissipation, as demonstrated experimentally in Refs. [37,41].

APPENDIX F: FITTING AND EXTRAPOLATING LOGICAL ERRORS

Our procedure for fitting and extrapolating logical errors is illustrated in Fig. 15. For unbiased noise [Fig. 15(a)], we fit the total logical error data (aggregated across all d) to the ansatz $p_X^{(L)} + p_Z^{(L)} = Ad^2(Bp)^{Cd}$, where $d_Z = d_X = d$, and capital letters are fit parameters, with the resulting fit shown as solid lines. The d^2 prefactor accounts for the fact that the number of fault locations that contribute to logical X (Z) errors increases both with the number of rounds (d), and the distance d_Z (d_X). This is only an approximation (see a more detailed discussion in, e.g., Ref. [3]), but the ansatz nevertheless fits the data well.

For biased noise, we separately fit $p_Z^{(L)}$ and $p_X^{(L)}$. The $p_Z^{(L)}$ data [Fig. 15(b)] are fit to the ansatz $Ad_Z(Bp_Z)^{Cd_Z}$, with fitting performed separately for each different choice of d_X , η . The prefactor of d_Z accounts for the fact that we simulate d_Z rounds of error correction, and we opt not to include an additional factor of d_X in the prefactor, since d_X is held fixed for each fitting. To fit $p_X^{(L)}$ [Fig. 15(c)], data are generated from logical memory simulations with a variety of choices for p_Z , η , and d_Z . For fixed d_X , we find that $p_X^{(L)}/d_Z^2$ is well described by a power-law fit of the form $A(p_Z/\eta)^B$. The d_Z^2 accounts for the increasing number of fault locations, and p_Z/η is a proxy for the physical X error probability.

In Fig. 16, we use the fits of $p_X^{(L)}/d_Z^2$ to verify the claim made in the caption of Fig. 6 that $d_X = 5$ and $|\alpha|^2 = 6$ is sufficient to achieve $p_X^{(L)} \lesssim 10^{-10}$ in the deep subthreshold regime. Specifically, we simulate the CX and CR_X gates at fixed $|\alpha|^2 = 6$ and varied q , then extract p_Z and η from

the gate error channels. We then use the fit of $p_X^{(L)}/d_Z^2$ to estimate $p_X^{(L)}$ at various d_Z . Given the “threshold” of $q = 5 \times 10^{-4}$ in Fig. 6, in the deep subthreshold regime ($q \leq 5 \times 10^{-5}$) we find that indeed $p_X^{(L)} < 10^{-10}$ for the d_Z values shown.

APPENDIX G: VALIDATING THE SIMPLIFIED BIASED ERROR MODEL

In this appendix, we compare the logical memory error of rectangular surface codes under two different circuit-level noise models: the “full model,” where Pauli error channels are directly extracted from master equation simulations of the cat-transmon gates, and the “simplified model” that replaces these error channels with simply parameterized biased error channels.

The two parameters describing the simplified error channels are the Z error probability, p_Z , and the bias η . We compute these parameters from the full CX and CR_X gate




FIG. 16. Extrapolated $p_X^{(L)}$ for fixed $d_X = 5$ and $|\alpha|^2 = 6$. Gate simulations are used to compute p_Z and η over a range of q values (using the noise model in Table I), then $p_X^{(L)}$ is computed using the $d_X = 5$ fit from Fig. 15(c).




FIG. 17. Comparing logical error rates for the full and simplified error models. (a) Comparison of $p_Z^{(L)}$. (b) Comparison of $p_X^{(L)}$. We compute $p_X^{(L)}$ either directly using the full error model (points), or from extrapolating a fit (solid line; see Appendix F) to the simplified error model data. The full error model data points comprise various different choices of the underlying parameters α , κ_1/χ , and d_Z , and the corresponding p_Z, η are computed from the resultant CX and CR_X error channels. All error bars indicate 95% confidence intervals.

Pauli error channels as follows. We define

$$p_Z = \frac{3}{2} \times \frac{1}{2} [p_{ZI}^{cX} + p_{ZZ}^{cX} + p_{ZI}^{cRX} + p_{ZZ}^{cRX}], \quad (G1)$$

where the factor of $3/2$ is a phenomenological factor that we find is necessary to obtain good agreement, and the remainder of the expression is the average per-gate probability of a Z error on a given cat qubit, averaged over the CX and CR_X gates. We then define the bias as

$$\eta = p_Z / (p_X + p_Y). \quad (G2)$$

Here, $(p_X + p_Y)$ denotes the sum of all probabilities for X or Y errors on the cat qubit, averaged over the CX and CR_X gates. The simplified error channels are then defined (for both single- and two-qubit gates) such that all Z-type errors have equal probability, with the total probability of a Z-type error equal to p_Z , and the total probability of any other error equal to p_Z/η .

With this mapping, we compare the logical memory performance with the two different circuit-level noise models in Fig. 17. We find that the logical error obtained with the simplified model is reasonably predictive of that obtained with the full model. Specifically, in Fig. 17(a), we compare the logical Z errors, $p_Z^{(L)}$. We find good agreement between the $p_Z^{(L)}$, with the agreement shown for the specific case of $d_X = 5, |\alpha|^2 = 6$ being representative of the agreement for other choices of these parameters. In Fig. 17(b), we compare the logical X errors, $p_X^{(L)}$. For all points, we find reasonable, order-of-magnitude agreement between the two, though, for some points, the simplified model fit falls outside the full model data error bars. The impact of a possible discrepancy in $p_X^{(L)}$ can be approximated as an effective rescaling of the noise bias. In particular, we expect $p_X^{(L)} \propto (1/\eta)^{(d_X+1)/2}$, so that an underestimate of $p_X^{(L)}$ by a factor x is roughly equivalent to an overestimate of the bias by a factor $x^{-2/(d_X+1)}$. For example, at $d_X = 5$,

an underestimate of $p_X^{(L)}$ by a factor of 2 ($x = 1/2$) would be roughly equivalent to an overestimate in the bias by a factor of approximately 1.25. Figure 8 provides a sense of how different levels of noise bias affect the logical memory overhead.

APPENDIX H: COMPENSATING THE CR_X PHASE DURING ERROR CORRECTION

As mentioned in the main text, the CR_X unitary

$$CR_X = |\alpha\rangle\langle\alpha| \otimes I + |-\alpha\rangle\langle-\alpha| \otimes R_X(\pi) \quad (H1)$$

is equivalent to a cat-controlled CX gate followed by an $S^\dagger = \text{diag}(1, -i)$ gate acting on the data cat qubit. In Fig. 18, we show that this additional phase gate does not present a problem for error correction. The basic idea is as follows. By changing the sign of the selective pulses in the CR_X pulse sequence, one can effectively replace $R_X(\pi) \rightarrow R_X(-\pi)$ in the above unitary, which is then equivalent to a CX gate followed by S . We can then arrange for one CR_X gate's S to be canceled by another CR_X gate's $S^\dagger$ in the next time step of an error-correction cycle, since the phase gates on the control commute with the CX gate. The net result is that we can effectively replace time-adjacent pairs of CR_X gates with pairs of CX gates, and we recover the usual surface code syndrome extraction circuits. Importantly, this approach does not increase the length of an error-correction cycle.

Special care needs to be taken when applying this approach to data qubits at the boundary, since in this case there is not guaranteed to be a CR_X gate at an adjacent time step. In this case (red box in Fig. 18), an additional ancillary transmon qubit can be added outside of the surface code patch, and this qubit can be used together with a CR_X gate to apply an S or $S^\dagger$ gate to the data cat qubit (see Appendix A 5). In the context of fault-tolerant computation




FIG. 18. Cancellation of the CR_X gate's accompanying phase. The left panel shows a surface code patch with data cat qubits as white circles and ancilla transmon qubits as green circles. The CX (green control, white target) and CR_X (white target) gates used for syndrome extraction are shown, with the time step during the error correction when each gate is applied indicated by the number above the gate. In the blue and red boxes on the right, the CR_X gates are represented by purple rectangles containing a standard CX gate followed by an S or $S^\dagger$ acting on the data cat qubit. These phase gates in adjacent time steps commute through the CX gates and cancel as described in the text.

with lattice surgery, such additional ancillary qubits would already be present in the associated routing space, so this strategy would not incur additional overhead.

APPENDIX I: COMPARISON OF CSS AND XZZX SURFACE CODES

In this appendix, we consider a variant of the cat-transmon architecture where error correction is performed with the XZZX surface code [33], instead of with the standard “CSS” surface code described in the main text. As discussed in Ref. [33] the XZZX code exhibits a higher threshold against highly biased noise in comparison to the CSS surface code. While the higher threshold allows for reduced logical memory overheads near threshold, we show that the CSS code nevertheless offers lower overheads in the practically relevant deep subthreshold regime, and hence why we focus on the CSS code in the main text.

The rectangular version of the XZZX code [11,33] is shown in Fig. 19(a). In contrast to the square version of the XZZX code, this rectangular version allows one to independently tune d_X and d_Z by adjusting the respective edge lengths of the code. This tunability allows for better hardware efficiency at large bias. For example, in the limit of infinite bias $d_X = 1$ suffices, and the rectangular XZZX code reduces to a simple repetition code at $d_X = 1$.

However, this rectangular XZZX code requires an “unrotated” lattice geometry, in contrast to the “rotated” lattice geometry of the rectangular CSS surface code in Fig. 1(a). (The rotated lattice geometry is not suitable for rectangular XZZX codes: both d_X and d_Z would be limited by the length of the shorter side, due to the diagonal nature of the logical operators in the rotated geometry.) Owing to this rotation difference, the rectangular XZZX code requires more qubits to achieve the same code distance as the rectangular CSS code. In particular, the rectangular XZZX code comprises $(2d_X - 1)(2d_Z - 1)$ total qubits (data and ancilla), which is roughly a factor of 2 larger than the overhead of the rectangular CSS code, $2d_X d_Z - 1$. Thus, in comparison to the rectangular CSS code, the higher threshold of the XZZX code comes at the price of a larger overhead at fixed code distance. We analyze this trade-off between higher threshold and larger overhead at fixed code distance below.

In the cat-transmon architecture, XZZX error syndromes can be measured using CX and CR_X gates with the circuit shown in Fig. 19(b). In this circuit, unitaries $U, U^\dagger$ are applied to the transmon between applications of the CX and CR_X gates, where

$$U = \pi_{ef} H_{g,e} \quad (11)$$

with $H_{g,e}$ denoting the Hadamard gate in the transmon's $|g\rangle, |e\rangle$ subspace, and $\pi_{ef} = |f\rangle\langle e| + \text{H.c.}$ denoting an $|e\rangle \leftrightarrow |f\rangle$ π pulse. This unitary serves two purposes: (1) it switches the transmon to its complementary basis (via the Hadamard) so that it is sensitive to the appropriate type of error, and (2) it maps the transmon from its $|g\rangle, |e\rangle$ subspace, where the CR_X gate is operated, to its $|g\rangle, |f\rangle$ subspace, where the CX gate is operated. Importantly, prior use of the $|g\rangle, |e\rangle$ subspace does not undermine the noise bias of subsequent CX gates. In particular, transmon errors before the first CX gate (caused, e.g., by a single decay event from $|e\rangle \rightarrow |g\rangle$) can propagate to an IXXZ error on the four connected data qubits. Because this error is stabilizer equivalent to ZIII , in effect this transmon error does not propagate to cat X errors. Additionally, we note that the additional phase associated with the CR_X unitary (not shown in Fig. 19) can be compensated as described in Appendix H by canceling the phase of a CR_X unitary at the end of one round with that of a CR_X unitary at the beginning of the next round.

In Fig. 19(c), we numerically compare the error-correction performance of the CSS and XZZX surface codes. We use the “simplified” biased-noise error model of Sec. V, with p_Z varied and the bias fixed to $\eta = 10^3$. This choice of η is motivated by the fact that (i) it is achievable with the cat-transmon architecture, (ii) though the thresholds of both codes improve with bias, the thresholds saturate once $\eta \gtrsim 10^3$ [33]. As expected, the XZZX code exhibits higher thresholds under this biased-noise model.




FIG. 19. The XZZX code in the cat-transmon architecture. (a) Rectangular XZZX code. A $d_x = 3$, $d_z = 5$ XZZX code is shown, comprising $(2d_x - 1)(2d_z - 1) = 45$ total qubits. The logical operators are shown in blue and orange. An example stabilizer generator is shown in purple, with arrows indicating the order in which entangling gates are applied between the data qubits (white circles) and ancilla qubit (green circle). (b) Syndrome extraction circuit. CR_X and CX gates are used to enable measurement of Z -type and X -type Pauli operators, respectively. In between CR_X and CX gates, a transmon operation U is applied to convert from one transmon basis to another (see the main text). (c) Comparison of rectangular XZZX (dashed lines) and CSS (solid lines) codes. We consider the simplified biased-noise model described in Sec. V of the main text, with bias $\eta = 10^3$. Under this model, we plot the total logical error $p^{(L)} = p_x^{(L)} + p_z^{(L)}$ after d_z rounds of noisy syndrome extraction for various code distances d_x, d_z . For visual clarity, we plot only $d_x = 3$ ($d_x = 5$) codes in the left (right) plots. Lines are colored according to their respective logical memory overheads: $2d_x d_z - 1$ for the CSS code, and $(2d_x - 1)(2d_z - 1)$ for the XZZX code. For visual clarity, we omit plot markers and error bars: we simulate 30 uniformly spaced p_z values over the range $p_z \in [10^{-4}, 10^{-1}]$, and we simulate repeated shots until 1000 logical errors are observed, up to 10^7 total shots.

As can be seen by comparing XZZX and CSS curves of similar color in Fig. 19, the XZZX code offers lower logical error at fixed qubit overhead for p_z values near the threshold as a result. However, again comparing similar overheads (similar colors), we observe steeper slopes for the CSS code. As a result, for p_z far below threshold, the CSS code exhibits lower logical error at fixed qubit overhead. Since in this work we are primarily interested in the deep subthreshold regime, we thus focus on the CSS code.

To provide intuition for this result, we note that the choice between the XZZX and CSS codes at fixed overhead can also be understood as a choice between increased threshold or increased code distance. That is, the XZZX code offers a higher threshold, but the CSS code offers a higher distance at fixed qubit overhead. At sufficiently low error, a higher code distance is always preferable to a higher threshold, as can be seen from the following scaling argument. Roughly speaking, logical error approximately scales as $(p/p_{\text{th}})^d$. Increasing the threshold $p_{\text{th}} \rightarrow p'_{\text{th}} > p_{\text{th}}$ suppresses logical error by a factor of $(p_{\text{th}}/p'_{\text{th}})^d$, while increasing the distance $d \rightarrow d' > d$ suppresses logical error by a factor of $(p/p_{\text{th}})^{d'-d}$. The latter is p dependent, while the former is not. Thus, at some sufficiently small p , increasing the distance will always yield lower logical error than increasing the threshold, and hence why the CSS code outperforms the XZZX code at sufficiently low error.

- [1] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, [arXiv:quant-ph/9811052](#).
- [2] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
- [3] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A—At. Mol. Opt. Phys.* **86**, 032324 (2012).
- [4] D. Gottesman, A. Kitaev, and J. Preskill, Encoding a qubit in an oscillator, *Phys. Rev. A* **64**, 012310 (2001).
- [5] 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).
- [6] M. H. Michael, M. Silveri, R. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin, New class of quantum error-correcting codes for a bosonic mode, *Phys. Rev. X* **6**, 031006 (2016).
- [7] 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).
- [8] W. Cai, Y. Ma, W. Wang, C.-L. Zou, and L. Sun, Bosonic quantum error correction codes in superconducting quantum circuits, *Fundam. Res.* **1**, 50 (2021).
- [9] A. Joshi, K. Noh, and Y. Y. Gao, Quantum information processing with bosonic qubits in circuit QED, *Quantum Sci. Technol.* **6**, 033001 (2021).

- [10] K. Fukui, A. Tomita, A. Okamoto, and K. Fujii, High-threshold fault-tolerant quantum computation with analog quantum error correction, *Phys. Rev. X* **8**, 021054 (2018).
- [11] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, Practical quantum error correction with the XZZX code and Kerr-cat qubits, *PRX Quantum* **2**, 030345 (2021).
- [12] A. L. Grimsmo and S. Puri, Quantum error correction with the Gottesman-Kitaev-Preskill code, *PRX Quantum* **2**, 020101 (2021).
- [13] 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, *PRX Quantum* **3**, 010329 (2022).
- [14] 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).
- [15] 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).
- [16] 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).
- [17] R. Lescanne, M. Villiers, T. Peronnin, 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).
- [18] 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).
- [19] N. E. Frattini, R. G. Cortiñas, J. Venkatraman, X. Xiao, Q. Su, C. U. Lei, B. J. Chapman, V. R. Joshi, S. Girvin, and R. J. Schoelkopf *et al.*, Observation of pairwise level degeneracies and the quantum regime of the Arrhenius law in a double-well parametric oscillator, *Phys. Rev. X* **14**, 031040 (2024).
- [20] U. Réglade, A. Bocquet, R. Gautier, J. Cohen, A. Marquet, E. Albertinale, N. Pankratova, M. Hallén, F. Rautschke, and L.-A. Sellem *et al.*, Quantum control of a cat qubit with bit-flip times exceeding ten seconds, *Nature* **629**, 778 (2024).
- [21] A. Hajr, B. Qing, K. Wang, G. Koolstra, Z. Pedramrazi, Z. Kang, L. Chen, L. B. Nguyen, C. Jünger, and N. Goss *et al.*, High-coherence Kerr-cat qubit in 2D architecture, *Phys. Rev. X* **14**, 041049 (2024).
- [22] A. Marquet, A. Essig, J. Cohen, N. Cottet, A. Murani, E. Albertinale, S. Dupouy, A. Bienfait, T. Peronnin, and S. Jezouin *et al.*, Autoparametric resonance extending the bit-flip time of a cat qubit up to 0.3 s, *Phys. Rev. X* **14**, 021019 (2024).
- [23] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. Devoret, L. Jiang, and R. Schoelkopf, Deterministic teleportation of a quantum gate between two logical qubits, *Nature* **561**, 368 (2018).
- [24] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. 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).
- [25] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zolys-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).
- [26] B. De Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, Error correction of a logical grid state qubit by dissipative pumping, *Nat. Phys.* **18**, 296 (2022).
- [27] V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. Brock, A. Ding, and L. Frunzio *et al.*, Real-time quantum error correction beyond break-even, *Nature* **616**, 50 (2023).
- [28] D. Lachance-Quirion, M.-A. Lemonde, J. O. Simoneau, L. St-Jean, P. Lemieux, S. Turcotte, W. Wright, A. Lacroix, J. Fréchette-Viens, and R. Shillito *et al.*, Autonomous quantum error correction of Gottesman-Kitaev-Preskill states, *Phys. Rev. Lett.* **132**, 150607 (2024).
- [29] 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).
- [30] B. L. Brock, S. Singh, A. Eickbusch, V. V. Sivak, A. Z. Ding, L. Frunzio, S. M. Girvin, and M. H. Devoret, Quantum error correction of qudits beyond break-even, *Nature* **641**, 612 (2025).
- [31] P. Aliferis and J. Preskill, Fault-tolerant quantum computation against biased noise, *Phys. Rev. A—At. Mol. Opt. Phys.* **78**, 052331 (2008).
- [32] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultrahigh error threshold for surface codes with biased noise, *Phys. Rev. Lett.* **120**, 050505 (2018).
- [33] J. P. Bonilla Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, The XZZX surface code, *Nat. Commun.* **12**, 1 (2021).
- [34] J. Roffe, L. Z. Cohen, A. O. Quintavalle, D. Chandra, and E. T. Campbell, Bias-tailored quantum LDPC codes, *Quantum* **7**, 1005 (2023).
- [35] J. Guillaud and M. Mirrahimi, Repetition cat qubits for fault-tolerant quantum computation, *Phys. Rev. X* **9**, 041053 (2019).
- [36] J. Guillaud and M. Mirrahimi, Error rates and resource overheads of repetition cat qubits, *Phys. Rev. A* **103**, 042413 (2021).
- [37] H. Putterman, K. Noh, R. N. Patel, G. A. Peairs, G. S. McCabe, M. Lee, S. Aghaeimeibodi, C. T. Hann, I. Jarrige, and G. Marcard *et al.*, Preserving phase coherence and linearity in cat qubits with exponential bit-flip suppression, *Phys. Rev. X* **15**, 011070 (2025).
- [38] 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).
- [39] É. Gouzien, D. Ruiz, F.-M. Le Régent, J. Guillaud, and N. Sangouard, Performance analysis of a repetition cat code architecture: Computing 256-bit elliptic curve logarithm

- in 9 hours with 126 133 cat qubits, *Phys. Rev. Lett.* **131**, 040602 (2023).
- [40] F.-M. Le Régent, C. Berdou, Z. Leghtas, J. Guillaud, and M. Mirrahimi, High-performance repetition cat code using fast noisy operations, *Quantum* **7**, 1198 (2023).
- [41] H. Putterman, K. Noh, C. T. Hann, G. S. MacCabe, S. Aghaieimebodi, R. N. Patel, M. Lee, W. M. Jones, H. Moradinejad, and R. Rodriguez *et al.*, Hardware-efficient quantum error correction via concatenated bosonic qubits, *Nature* **638**, 927 (2025).
- [42] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, Fault-tolerant detection of a quantum error, *Science* **361**, 266 (2018).
- [43] P. Reinhold, S. Rosenblum, W.-L. Ma, L. Frunzio, L. Jiang, and R. J. Schoelkopf, Error-corrected gates on an encoded qubit, *Nat. Phys.* **16**, 822 (2020).
- [44] W.-L. Ma, M. Zhang, Y. Wong, K. Noh, S. Rosenblum, P. Reinhold, R. J. Schoelkopf, and L. Jiang, Path-independent quantum gates with noisy ancilla, *Phys. Rev. Lett.* **125**, 110503 (2020).
- [45] Y. Liu, S. Singh, K. C. Smith, E. Crane, J. M. Martyn, A. Eickbusch, A. Schuckert, R. D. Li, J. Sinanan-Singh, and M. B. Soley *et al.*, Hybrid oscillator-qubit quantum processors: Instruction set architectures, abstract machine models, and applications, *arXiv:2407.10381*.
- [46] K. S. Dubovitskii, Bit-flip errors in dissipative cat qubits: Second-order perturbation theory, *Phys. Rev. A* **111**, 012617 (2025).
- [47] C. Berdou, A. Murani, U. Reglade, W. C. Smith, M. Villiers, J. Palomo, M. Rosticher, A. Denis, P. Morfin, and M. Delbecq *et al.*, One hundred second bit-flip time in a two-photon dissipative oscillator, *PRX Quantum* **4**, 020350 (2023).
- [48] A. Blais, A. L. Grimsom, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [49] W.-L. Ma, S.-S. Li, and L. Jiang, Algebraic structure of path-independent quantum control, *Phys. Rev. Res.* **4**, 023102 (2022).
- [50] Q. Xu, P. Zeng, D. Xu, and L. Jiang, Fault-tolerant operation of bosonic qubits with discrete-variable ancillae, *Phys. Rev. X* **14**, 031016 (2024).
- [51] R. Gautier, A. Sarlette, and M. Mirrahimi, Combined dissipative and Hamiltonian confinement of cat qubits, *PRX Quantum* **3**, 020339 (2022).
- [52] Z. Leghtas, G. Kirchmair, B. Vlastakis, R. J. Schoelkopf, M. H. Devoret, and M. Mirrahimi, Hardware-efficient autonomous quantum memory protection, *Phys. Rev. Lett.* **111**, 120501 (2013).
- [53] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. Girvin, R. J. Schoelkopf, and M. H. Devoret, Fast universal control of an oscillator with weak dispersive coupling to a qubit, *Nat. Phys.* **18**, 1464 (2022).
- [54] A. A. Diringer, E. Blumenthal, A. Grinberg, L. Jiang, and S. Hacoen-Gourgy, Conditional-not displacement: Fast multioscillator control with a single qubit, *Phys. Rev. X* **14**, 011055 (2024).
- [55] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, Noise spectroscopy through dynamical decoupling with a superconducting flux qubit, *Nat. Phys.* **7**, 565 (2011).
- [56] W. Magnus, On the exponential solution of differential equations for a linear operator, *Commun. Pure. Appl. Math.* **7**, 649 (1954).
- [57] L. Cywiński, R. M. Lutchyn, C. P. Nave, and S. D. Sarma, How to enhance dephasing time in superconducting qubits, *Phys. Rev. B* **77**, 174509 (2008).
- [58] N. Didier, E. A. Sete, J. Combes, and M. P. da Silva, ac flux sweet spots in parametrically modulated superconducting qubits, *Phys. Rev. Appl.* **12**, 054015 (2019).
- [59] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly nonlinear qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [60] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer's guide to superconducting qubits, *Appl. Phys. Rev.* **6**, 021318 (2019).
- [61] F. Motzoi and F. K. Wilhelm, Improving frequency selection of driven pulses using derivative-based transition suppression, *Phys. Rev. A—At. Mol. Opt. Phys.* **88**, 062318 (2013).
- [62] Q. Xu, J. K. Iverson, F. G. Brandão, and L. Jiang, Engineering fast bias-preserving gates on stabilized cat qubits, *Phys. Rev. Res.* **4**, 013082 (2022).
- [63] K. S. Chou, T. Shemma, H. McCarrick, T.-C. Chien, J. D. Teoh, P. Winkel, A. Anderson, J. Chen, J. Curtis, and S. J. de Graaf *et al.*, Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements, *arXiv:2307.03169*.
- [64] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, and K. Chou *et al.*, Quantum memory with millisecond coherence in circuit QED, *Phys. Rev. B* **94**, 014506 (2016).
- [65] O. Milul, B. Guttel, U. Goldblatt, S. Hazanov, L. M. Joshi, D. Chausovsky, N. Kahn, E. Çiftçürek, F. Lafont, and S. Rosenblum, Superconducting cavity qubit with tens of milliseconds single-photon coherence time, *PRX Quantum* **4**, 030336 (2023).
- [66] K. Murch, S. Weber, E. Levenson-Falk, R. Vijay, and I. Siddiqi, $1/f$ noise of Josephson-junction-embedded microwave resonators at single photon energies and millikelvin temperatures, *Appl. Phys. Lett.* **100**, 142601 (2012).
- [67] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [68] C. Gidney, Stim: A fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [69] O. Higgott, Pymatching: A python package for decoding quantum codes with minimum-weight perfect matching, *ACM Trans. Quantum Comput.* **3**, 1 (2022).
- [70] The simulations assume 0.5% error each for transmon readout and reset, but the combined error probability (1%) is the only physically relevant parameter, since reset errors can be propagated to readout errors in the next round.
- [71] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2010).
- [72] 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).
- [73] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, and C. Yang *et al.*, Towards practical quantum computers: Transmon qubit with a lifetime approaching 0.5 milliseconds, *Npj Quantum Inf.* **8**, 3 (2022).
- [74] M. Tuokkola, Y. Sunada, H. Kivijärvi, L. Grönberg, J.-P. Kaikkonen, V. Vesterinen, J. Govenius, and M. Möttönen, Methods to achieve near-millisecond energy relaxation and dephasing times for a superconducting transmon qubit, [arXiv:2407.18778](https://arxiv.org/abs/2407.18778).
- [75] J. Ghosh and A. G. Fowler, Leakage-resilient approach to fault-tolerant quantum computing with superconducting elements, *Phys. Rev. A* **91**, 020302 (2015).
- [76] K. C. Miao, M. McEwen, J. Atalaya, D. Kafri, L. P. Pryadko, A. Bengtsson, A. Opremcak, K. J. Satzinger, Z. Chen, and P. V. Klimov *et al.*, Overcoming leakage in quantum error correction, *Nat. Phys.* **19**, 1780 (2023).
- [77] X. Yang, J. Chu, Z. Guo, W. Huang, Y. Liang, J. Liu, J. Qiu, X. Sun, Z. Tao, and J. Zhang *et al.*, Coupler-assisted leakage reduction for scalable quantum error correction with superconducting qubits, *Phys. Rev. Lett.* **133**, 170601 (2024).
- [78] G. S. MacCabe, H. Ren, J. Luo, J. D. Cohen, H. Zhou, A. Sipahigil, M. Mirhosseini, and O. Painter, Nano-acoustic resonator with ultralong phonon lifetime, *Science* **370**, 840 (2020).
- [79] P. Arrangoiz-Arriola, E. A. Wollack, Z. Wang, M. Pechal, W. Jiang, T. P. McKenna, J. D. Witmer, R. Van Laer, and A. H. Safavi-Naeini, Resolving the energy levels of a nanomechanical oscillator, *Nature* **571**, 537 (2019).
- [80] Y. Ye, K. Peng, M. Naghiloo, G. Cunningham, and K. P. O’Brien, Engineering purely nonlinear coupling between superconducting qubits using a qurion, *Phys. Rev. Lett.* **127**, 050502 (2021).
- [81] L. B. Nguyen, Y.-H. Lin, A. Somoroff, R. Mencia, N. Grabon, and V. E. Manucharyan, High-coherence fluxonium qubit, *Phys. Rev. X* **9**, 041041 (2019).
- [82] F. Bao, H. Deng, D. Ding, R. Gao, X. Gao, C. Huang, X. Jiang, H.-S. Ku, Z. Li, and X. Ma *et al.*, Fluxonium: An alternative qubit platform for high-fidelity operations, *Phys. Rev. Lett.* **129**, 010502 (2022).
- [83] H. Zhang, S. Chakram, T. Roy, N. Earnest, Y. Lu, Z. Huang, D. Weiss, J. Koch, and D. I. Schuster, Universal fast-flux control of a coherent, low-frequency qubit, *Phys. Rev. X* **11**, 011010 (2021).
- [84] Y.-H. Lin, L. B. Nguyen, N. Grabon, J. San Miguel, N. Pankratova, and V. E. Manucharyan, Demonstration of protection of a superconducting qubit from energy decay, *Phys. Rev. Lett.* **120**, 150503 (2018).
- [85] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).
- [86] S. Singh, A. S. Darmawan, B. J. Brown, and S. Puri, High-fidelity magic-state preparation with a biased-noise architecture, *Phys. Rev. A* **105**, 052410 (2022).
- [87] C. A. Pattison, A. Krishna, and J. Preskill, Hierarchical memories: Simulating quantum LDPC codes with local gates, *Quantum* **9**, 1728 (2025).
- [88] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, *PRX Quantum* **2**, 040101 (2021).
- [89] P. Panteleev and G. Kalachev, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2022), p. 375.
- [90] A. Leverrier and G. Zémor, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Denver, CO, USA, 2022), p. 872.
- [91] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
- [92] D. Ruiz, J. Guillaud, A. Leverrier, M. Mirrahimi, and C. Vuillot, LDPC-cat codes for low-overhead quantum computing in 2D, [arXiv:2401.09541](https://arxiv.org/abs/2401.09541).
- [93] J. M. Gambetta, F. Motzoi, S. Merkel, and F. K. Wilhelm, Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator, *Phys. Rev. A—At. Mol. Opt. Phys.* **83**, 012308 (2011).
- [94] L. B. Nguyen, G. Koolstra, Y. Kim, A. Morvan, T. Chistolini, S. Singh, K. N. Nesterov, C. Jünger, L. Chen, and Z. Pedramrazi *et al.*, Blueprint for a high-performance fluxonium quantum processor, *PRX Quantum* **3**, 037001 (2022).
- [95] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, Universal control of an oscillator with dispersive coupling to a qubit, *Phys. Rev. A* **92**, 040303 (2015).
- [96] K. Li, P. Zhao, and D. Tong, Approach to realizing nonadiabatic geometric gates with prescribed evolution paths, *Phys. Rev. Res.* **2**, 023295 (2020).
- [97] R. Gautier, M. Mirrahimi, and A. Sarlette, Designing high-fidelity Zeno gates for dissipative cat qubits, *PRX Quantum* **4**, 040316 (2023).
- [98] C. J. Wood, J. D. Biamonte, and D. G. Cory, Tensor networks and graphical calculus for open quantum systems, [arXiv:1111.6950](https://arxiv.org/abs/1111.6950).