

Preserving Phase Coherence and Linearity in Cat Qubits with Exponential Bit-Flip Suppression

Harald Putterman^{1,*}, Kyungjoo Noh^{1,*}, Rishi N. Patel¹, Gregory A. Peairs¹, Gregory S. MacCabe¹, Menyong Lee¹, Shahriar Aghaeimeibodi¹, Connor T. Hann¹, Ignace Jarrige¹, Guillaume Marcaud¹, Yuan He¹, Hesam Moradinejad¹, John Clai Owens¹, Thomas Scaffidi^{1,‡}, Patricio Arrangoiz-Arriola¹, Joe Iverson¹, Harry Levine¹, Fernando G. S. L. Brandão^{1,2}, Matthew H. Matheny¹, and Oskar Painter^{1,2,3,§}

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 12 October 2024; accepted 12 February 2025; published 25 March 2025)

Cat qubits, a type of bosonic qubit encoded in a harmonic oscillator, can exhibit an exponential noise bias against bit-flip errors with increasing mean photon number. Here, we focus on cat qubits stabilized by two-photon dissipation, where pairs of photons are added and removed from a harmonic oscillator by an auxiliary, lossy buffer mode. This process requires a large loss rate and strong nonlinearities of the buffer mode that must not degrade the coherence and linearity of the oscillator. In this work, we show how to overcome this challenge by coloring the loss environment of the buffer mode with a multipole filter and optimizing the circuit to take into account additional inductances in the buffer mode. Using these techniques, we achieve near-ideal enhancement of cat-qubit bit-flip times with increasing photon number, reaching over 0.1 s with a mean photon number of only 4. Concurrently, our cat qubit remains highly phase coherent, with phase-flip times corresponding to an effective lifetime of $T_{1,\text{eff}} \simeq 70 \mu\text{s}$, comparable with the bare oscillator lifetime. We achieve this performance even in the presence of an ancilla transmon, used for reading out the cat-qubit states, by engineering a tunable oscillator-ancilla dispersive coupling. Furthermore, the low nonlinearity of the harmonic oscillator mode allows us to perform pulsed cat-qubit stabilization, an important control primitive, where the stabilization can remain off for a significant fraction (e.g., two-thirds) of a 3 μs cycle without degrading bit-flip times. These advances are important for the realization of scalable error correction with cat qubits, where large noise bias and low phase-flip error rate enable the use of hardware-efficient outer error-correcting codes.

DOI: [10.1103/PhysRevX.15.011070](https://doi.org/10.1103/PhysRevX.15.011070)

Subject Areas: Quantum Physics,
Quantum Information,
Superconductivity

I. INTRODUCTION

The expected system size required for realizing fault-tolerant quantum computers far exceeds the scale of current experiments [1–3]. One approach to reduce the resource overhead for fault tolerance is to use qubits which have a

biased-noise structure (e.g., bit-flip error rates are much smaller than the phase-flip error rates). This noise bias property allows the quantum error correcting code to focus predominantly on correcting only one type of error, for example, phase-flip errors, leading to reduced overhead for error correction [4–14].

Cat qubits [4,15] are a specific realization of biased-noise qubits. In cat qubits, the quantum information is encoded into the infinite-dimensional Hilbert space of a quantum harmonic oscillator. The computational basis states of a cat qubit are approximately equivalent to the coherent states, $|0\rangle \simeq |+\alpha\rangle$ and $|1\rangle \simeq |-\alpha\rangle$. The complementary basis states are exactly given by the even- and odd-parity cat states, i.e., $|\pm\rangle = |C_{\alpha}^{\pm}\rangle \propto |+\alpha\rangle \pm |-\alpha\rangle$ (see Fig. 1). Cat qubits can have exponentially suppressed bit-flip (X) error rates with increasing mean photon number $|\alpha|^2$ [16–19]. This is due to the phase-space separation

*These authors contributed equally to this work.

†Contact author: putterma@amazon.com

‡Present address: Department of Physics and Astronomy, University of California, Irvine, California 92697, USA.

§Contact author: ojpb@amazon.com

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




FIG. 1. Bloch sphere of a cat qubit. Wigner functions of the computational and complementary basis states are represented next to the corresponding points on the Bloch sphere.

between the two computational basis states which gets larger as $|\alpha|$ increases. In contrast, the phase-flip (Z) error rates of a cat qubit increase only linearly with the mean photon number $|\alpha|^2$.

To realize this exponential noise bias against bit-flip errors, it is necessary to stabilize the oscillator mode in the cat-qubit state manifold. In this work, we focus on cat qubits stabilized with two-photon dissipation, so-called “dissipative cat qubits” [20–22]. Dissipative cat qubits are implemented by engineering a nonlinear coupling between an oscillator mode (also referred to as a storage mode) and a lossy buffer mode. The nonlinear coupling mediates an exchange of two photons from the storage to one photon in the buffer mode via a three-wave mixing (3WM) nonlinearity. By adding a linear buffer drive and adiabatically eliminating the lossy buffer mode, we realize a two-photon dissipation described by the master equation

$$\frac{d\hat{\rho}}{dt} = \kappa_2 D[\hat{a}^2 - \alpha^2]\hat{\rho}, \quad (1)$$

which stabilizes the oscillator mode \hat{a} in a two-dimensional steady-state subspace spanned by the coherent states $|\pm\alpha\rangle$ [20].

For cat qubits to enable reduced error correction overhead, it is essential to achieve an exponential noise bias without sacrificing the cat qubit’s phase coherence. Specifically, a large noise bias must be achieved at small values of $|\alpha|^2$ where the phase-flip error rate remains low enough to be correctable by an outer error-correcting code such as a repetition code [7,8,11,12]. Achieving this requires preserving the coherence and linearity of the storage mode under the two-photon dissipation. This is challenging due to the strong nonlinear 3WM coupling of the storage mode to the highly lossy buffer mode used for dissipative stabilization. Additionally, in some architectures, interactions needed to characterize the storage mode can also create undesired nonlinearities. While there has been encouraging recent progress addressing these challenges, prior experiments have so far demonstrated either high cat phase coherence [22] or exponential noise bias [16,18,19] but not both simultaneously.

In this paper, we present a cat-qubit system where the bit-flip error rates are exponentially suppressed and the cat qubit remains highly phase coherent. In particular, we achieve high effective storage-mode lifetimes of $T_1 \simeq 70 \mu\text{s}$ by using a multipole metamaterial filter to protect the storage from the strong buffer mode loss channel. Simultaneously, our device employs two novel strategies to preserve the storage linearity. First, a tunable coupler with high on-off ratio dispersive coupling enables characterization of the storage mode with an ancillary transmon without introducing deleterious nonlinearities. Second, we engineer an advantageous cancellation of buffer-induced nonlinearities by carefully accounting for serial inductances present in our buffer mode circuit. Together, the storage mode’s high coherence and low nonlinearity enable us to demonstrate large noise bias at low photon numbers, as is crucial for hardware-efficient error correction. We measure exponential bit-flip suppression, with near-ideal scaling [20], up to 0.1 s bit-flip times at $|\alpha|^2 = 4$. Furthermore, when pulsing the two-photon dissipation on and off with 3 μs cycles, we are able to achieve bit-flip times comparable to the continuous case despite the fact that the stabilization is on only one-third of the time during a cycle. These advances are key enablers for performing more complex operations on cat qubits such as multiqubit gates and error-correction circuits.

This paper is organized as follows. After providing an overview of our device in Sec. II, we present the key technical advances necessary for realizing a high-coherence and low-nonlinearity storage mode. First, in Sec. III, we introduce a tunable dispersive coupling with a high on-off ratio for reading out the storage-mode state, where the tunability protects the storage mode from errors due to the readout components. Next, in Sec. IV, we address the trade-off between the required strong buffer loss and storage coherence by coloring the buffer environment with a multipole metamaterial filter. Finally, in Sec. V, we discuss the importance of accounting for serial inductances within the buffer to realize storage modes with low nonlinearity. Equipped with these technical advances, we then benchmark the performance of our cat qubit. We demonstrate the cat qubit’s exponential noise bias in Sec. VI and the persistence of that noise bias under pulsed stabilization in Sec. VII.

II. DEVICE OVERVIEW

We consider a system containing a storage mode \hat{a} , a buffer mode \hat{b} , and components for reading out the state of the storage mode as shown in Fig. 2. The central element of our system is the storage mode (blue in Fig. 2), where the cat-qubit states are encoded. The storage mode is implemented as a $\lambda/2$ resonator and has frequency $\omega_s/2\pi = 5.35 \text{ GHz}$.

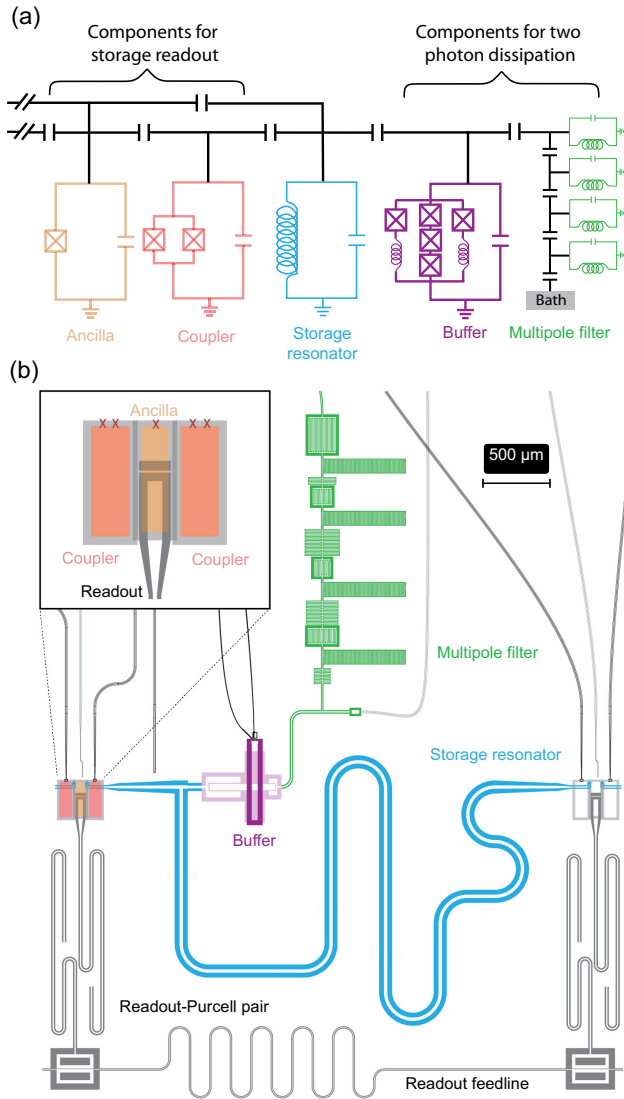


FIG. 2. Device overview. (a) Circuit representation of the key components of a dissipative cat qubit with ancilla transmon. The storage mode is read out by an ancilla transmon using a tunable-transmon coupler. The storage mode is also coupled to a buffer mode. The buffer is coupled to a four-pole metamaterial filter consisting of an array of coupled LC oscillators. The buffer and filter modes are used for implementing two-photon dissipation to stabilize a cat qubit in the storage mode. (b) Representative physical layout of the device. The storage mode is a $\lambda/2$ resonator. The filter is implemented as a series of lumped LC oscillators. The inset shows ancilla and couplers, and the X symbols represent junctions.

The buffer mode (purple in Fig. 2) is implemented using a variant of the asymmetrically threaded superconducting quantum interference device (SQUID) (ATS) element [16], consisting of a middle junction array and two side junctions. Two noteworthy aspects of our buffer design are the shallowness of the middle junction array (in the sense that it has only three junctions in the array) and the fact that each side junction has a small but non-negligible serial

inductance. The significance of these design aspects are discussed in detail in Sec. V. The buffer has two flux lines controlling the fluxes through the two loops of the ATS, one symmetrically (φ_Σ) and the other differentially (φ_Δ). It is also coupled to a lossy environment through an output line colored by a four-pole metamaterial bandpass filter [23,24] (green in Fig. 2), with the passband ranging from ~ 2.6 to ~ 3.6 GHz. At the chosen operating point, the frequency and the loss rate of the buffer mode are given by $\omega_b/2\pi = 3.01$ GHz and $\kappa_b/2\pi = 10.7$ MHz, respectively.

The storage mode is read out by an ancilla transmon (yellow in Fig. 2) with a dedicated readout resonator. The storage mode is dispersively coupled to the ancilla transmon through a tunable-transmon coupler (red in Fig. 2). The frequency of the coupler can be tuned by varying the external flux of the coupler’s SQUID loop. This tunability allows us to selectively turn on the dispersive coupling between the storage mode and the ancilla qubit only when the storage mode needs to be read out (see Sec. III).

The device consists of two silicon chips which are flip-chip bonded together [25,26]. The bottom chip hosts the linear elements including the storage mode and uses a tantalum superconducting thin-film layer [27]. The top chip has the nonlinear elements containing Josephson junctions such as the buffer mode, the ancilla transmon, and the tunable-transmon coupler and uses a series of aluminum superconducting thin-film layers. The schematic in Fig. 2(a) shows the circuit diagram of our device, with the circuit layout shown in Fig. 2(b). As indicated in the circuit layout, the cat-qubit device studied here is part of a larger integrated circuit involving an array of cat qubits, with each storage mode coupled two ancilla qubits, one on the left and one on the right.

III. TUNABLE DISPERSIVE COUPLING FOR READING OUT A STORAGE MODE

A commonly used strategy for reading out quantum states of a linear oscillator mode is to couple it dispersively to an ancillary qubit. Through the dispersive coupling, the ancilla qubit experiences frequency shifts depending on the number of excitations in the oscillator mode. Such frequency shifts can then be used to characterize the state of the oscillator mode [28–31]. However, this same dispersive coupling can also introduce unwanted frequency shifts of the oscillator mode when the number of excitations in the ancilla qubit changes due to, for instance, decay or heating. Such unwanted frequency shifts have been one of the main performance bottlenecks for various error-correction protocols with oscillator modes [16,32,33]. In this section, we provide a solution to this problem by engineering tunable dispersive coupling between a storage mode and an ancilla transmon via a tunable-transmon coupler.

Here, we focus on a subsystem comprising a storage mode, a tunable-transmon coupler, and a fixed-frequency ancilla transmon which are capacitively coupled (see

Appendix H 1). The effective Hamiltonian of this subsystem takes the form

$$\hat{H} = \omega_s \hat{a}^\dagger \hat{a} + \omega_c |e_c\rangle \langle e_c| + \omega_a |e_a\rangle \langle e_a| + \chi_{sa} \hat{a}^\dagger \hat{a} |e_a\rangle \langle e_a|, \quad (2)$$

in the dressed eigenbasis where all the parameters ω_s , ω_c , ω_a , and χ_{sa} depend on the external flux of the coupler, $\Phi_{x,c}$. Here, $|e_c\rangle$ and $|e_a\rangle$ refer to the first excited state of the coupler and the ancilla, respectively. Thus, ω_s , ω_c , and ω_a are the frequencies of the storage, coupler, and ancilla, respectively, in the lowest energy manifold consisting of the ground state and the first excited state. Most importantly, χ_{sa} is the strength of the dispersive coupling between the storage and the ancilla, i.e., the frequency shift of the storage mode when the ancilla is excited from the ground state to its first excited state $|e_a\rangle$.

The external flux of the coupler, $\Phi_{x,c}$, tunes the frequency of the coupler mode. In Fig. 3(a), we show the frequencies of the storage, ancilla, and coupler as a function of $\Phi_{x,c}$. The frequency of the coupler tunes from $\omega_{c,\max}/2\pi = 8.40$ GHz to $\omega_{c,\min}/2\pi = 5.77$ GHz as the coupler external flux is varied from 0 to the half flux quantum $\Phi_0/2$. Note that the storage and ancilla frequencies are also lowered as the coupler approaches its minimum frequency. This is due to the increased mode hybridization with the coupler and the associated repulsion between the energy levels.

We select the idle (off) position to be when the coupler is at its maximum frequency with $\Phi_{x,c} = 0$. To resolve the small dispersive coupling between the storage and ancilla with the coupler in the off position, we perform Ramsey measurements of the ancilla transmon with a variable mean photon number coherent state in the storage mode (see Appendix K). In particular, a mean photon number up to 20 is used to amplify the impact of the dispersive coupling on the ancilla. With this method, we find that the strength of the dispersive coupling is only $\chi_{sa}/2\pi = -1.4 \pm 0.1$ kHz in the off position.

We can dynamically turn on the dispersive coupling by applying a flux pulse on the tunable coupler [see ‘‘coupler fast flux’’ in Fig. 3(d)]. By pulsing the coupler to lower frequencies, we increase the hybridization between modes and, consequently, increase the dispersive coupling between the storage and ancilla. In Fig. 3(b), we characterize the dispersive coupling at each flux amplitude by measuring the difference in phase accumulation rate on the storage mode conditioned on the state of the ancilla (see Appendix J for more details). The maximum dispersive coupling is achieved with the coupler at its minimum frequency, where $|\chi_{sa}|/2\pi > 5$ MHz.

In practice, we do not use the coupler’s minimum frequency position ($\Phi_{x,c} = \Phi_0/2$), as the dispersive nature of the interaction breaks down in the large storage photon number regime due to strong hybridization (see

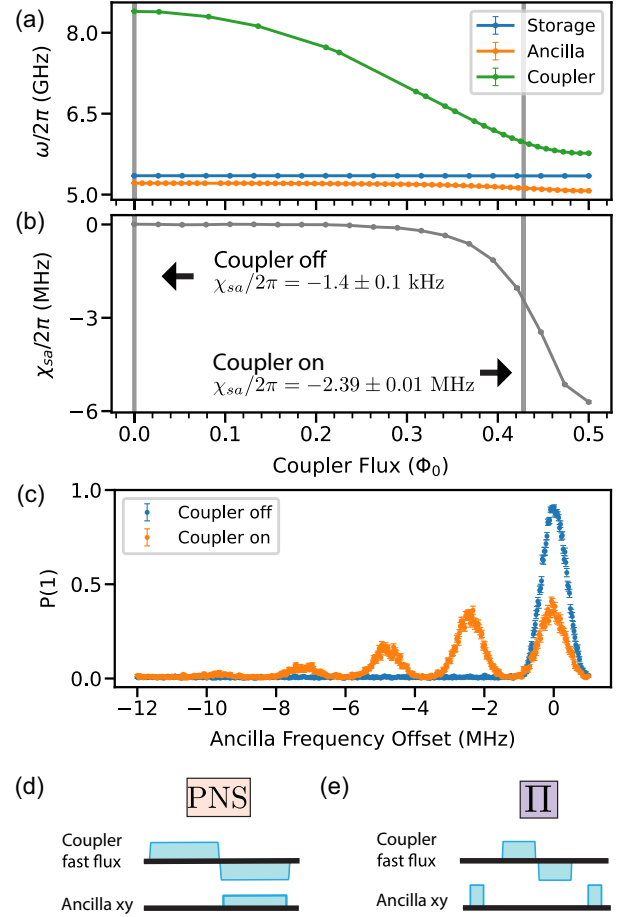


FIG. 3. Tunable dispersive coupling for reading out a storage mode. (a) Frequencies of the coupler, ancilla, and storage mode as a function of the external flux of the coupler. (b) Storage-ancilla dispersive coupling as a function of the coupler flux measured using a storage conditional phase experiment. When the tunable coupler is pulsed to the on position with $\Phi_{x,c} = 0.43\Phi_0$, the storage and ancilla transmon develop a dispersive coupling $|\chi_{sa}|/2\pi > 2$ MHz. Using an alternative measurement method (based on Ramsey experiments on the ancilla), we measure the $\chi_{sa}/2\pi = -1.4 \pm 1$ kHz in the off position with $\Phi_{x,c} = 0$. (c) Number-splitting spectroscopy on the ancilla transmon with the coupler in either the on or off position and the storage mode in a coherent state. (d) Photon-number-selective (PNS) pulse sequence on the ancilla used in (c) to measure the storage mode. (e) Storage parity mapping (Π) pulse sequence used in, for example, Fig. 4(c). Both storage-mode measurement sequences in (d) and (e) use net-zero waveforms for the tunable coupler flux pulse.

Appendix H 1). Instead, we select an on position to be $\Phi_{x,c} = 0.43\Phi_0$, where we observe a dispersive coupling of $\chi_{sa}/2\pi = -2.39 \pm 0.01$ MHz. Relative to the off position $\chi_{sa}/2\pi$ of -1.4 kHz, this constitutes a dynamic range of >1000 for the storage-ancilla dispersive coupling. The high on-off ratio is made possible by carefully optimizing the direct capacitive coupling between the storage and ancilla so that the net dispersive interaction

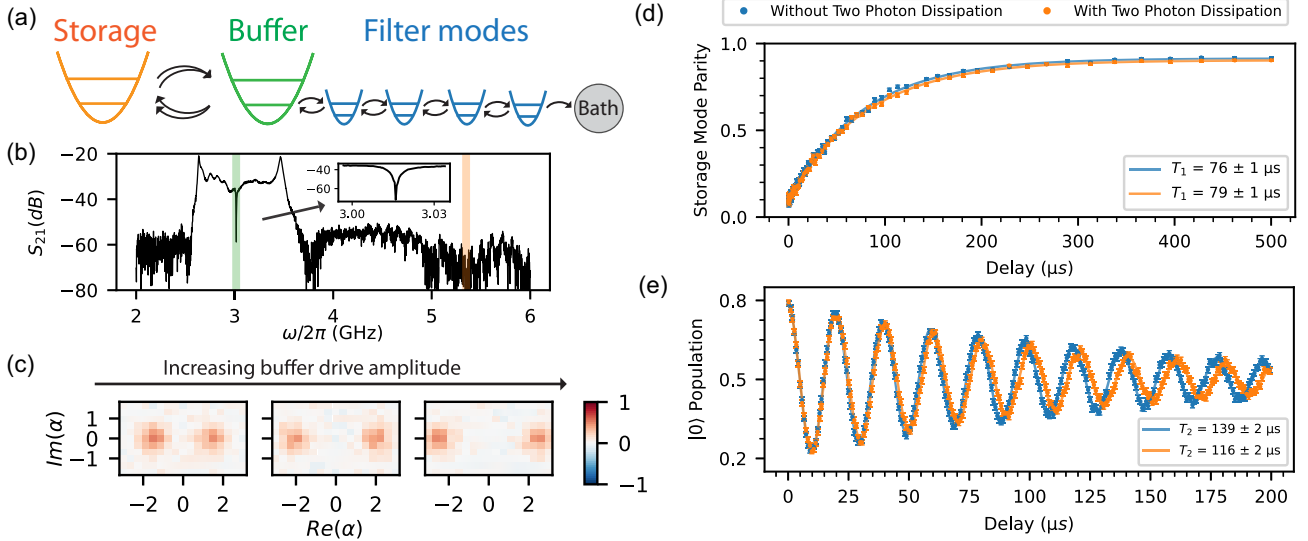


FIG. 4. Engineering two-photon dissipation while preserving storage-mode coherence. (a) Schematic diagram of a storage mode which is coupled to a lossy buffer mode by a three-wave mixing process exchanging two photons from the storage mode with one photon in the buffer mode. The strong loss on the buffer mode is engineered by coupling the buffer mode to a lossy transmission line through a four-pole metamaterial filter. (b) The transmission spectrum through the four-pole filter. The filter colors the environment seen by the buffer mode. The buffer frequency (green) lies within the passband resulting in strong damping of the buffer mode. The storage-mode frequency (orange) lies outside of the filter passband resulting in suppression of storage-mode loss to the filtered environment. Note that these S_{21} data are convolved with the uncalibrated response of a 2.5–4 GHz isolator so the filter spectrum may be different than that reported outside the 2.5–4 GHz range and, thus, at the storage-mode frequency. Nonetheless, the storage loss is strongly filtered, with electromagnetic simulations predicting a storage loss rate due to the filter of approximately 0.2 kHz. (c) Wigner tomograms of the steady-state cat-qubit manifolds of the two-photon dissipation with increasing buffer drive amplitude. The coherent state amplitude of a cat qubit in the storage mode increases as the amplitude of the linear buffer drive increases. (d) Measured T_1 of the storage mode with and without the pure two-photon dissipation ($\kappa_2 \mathcal{D}[\hat{a}^2]$, i.e., without the buffer drive) being applied during delay. (e) Measured T_2 of the storage mode with and without the pure two-photon dissipation being applied during delay.

between them is minimized around the idle off position of the tunable coupler [34,35].

The tunability of the dispersive interaction is further demonstrated in Fig. 3(c), through number-splitting spectroscopy of the ancilla transmon. See Fig. 3(d). Here, the storage mode is driven into a coherent state, and a weak π pulse (length 1.4 μs) is applied to the ancilla transmon. By varying the drive frequency of the ancilla π pulse, we measure the spectrum of the ancilla transmon. When the coupler is in the off position (i.e., no coupler flux pulse applied), we cannot resolve the splitting of the ancilla frequency, as expected and desired. When the coupler is pulsed to the on position, we see distinct peaks in the ancilla spectrum, corresponding to different Fock states of the storage mode. The spacing between the neighboring peaks is consistent with the dispersive shift χ_{sa} , measured in Fig. 3(b).

As shown in Fig. 3(e), the tunable dispersive coupling also enables parity measurement of the storage mode by combining a Ramsey measurement on the ancilla with a flux pulse ($\Phi_{x,c}/\Phi_0 = 0.38$; see Appendix C) on the coupler. In the parity measurement, one needs to apply unselective $\pi/2$ pulses on the ancilla, as well as a conditional (on the ancilla being excited) π rotation of the storage

mode through the storage-ancilla dispersive coupling. Because of the large on-off ratio, the unselective $\pi/2$ pulses are easy to achieve, since they are applied when the storage-ancilla dispersive coupling is turned off and the resonant frequency of the ancilla is unperturbed by the storage-mode state. The parity measurement allows us to perform Wigner tomography [28] of the storage mode [see, e.g., Fig. 4(c)]. Further advantages of the tunable dispersive coupling in the specific context of a cat qubit are demonstrated in Sec. VI.

IV. ENGINEERING TWO-PHOTON DISSIPATION WHILE PRESERVING STORAGE-MODE COHERENCE

Equipped with the ability to read out the storage-mode state, we now describe the engineering of two-photon dissipation for stabilizing a cat qubit. Two-photon dissipation can be realized by engineering a 3WM Hamiltonian $\hat{H}_{3WM} = g_2 \hat{a}^2 \hat{b}^\dagger + \text{H.c.}$, a linear drive $\hat{H}_d = -g_2 \alpha^2 \hat{b}^\dagger + \text{H.c.}$, and a loss $\kappa_b \mathcal{D}[\hat{b}]$ on the buffer mode. After adiabatically eliminating the buffer mode, the dynamics of this system obey Eq. (1) with $\kappa_2 \simeq 4g_2^2/\kappa_b$. Here, the key technical challenge is the implementation of the 3WM

Hamiltonian, which requires a strong flux pump on the buffer mode and large buffer loss rate. The strong driven-dissipative processes on the buffer mode should not degrade the coherence of the storage mode.

In order to protect the storage mode from the strong loss channel of the buffer mode, we color the loss environment of the buffer mode with a multipole meta-material filter, so called because it is constructed from many subwavelength components. [11,23,24,36]. Unlike commonly used Purcell filters composed of only one resonator mode, our filter consists of four lumped-element resonator modes which are capacitively coupled. It can also be understood as a series of coupled oscillators as illustrated in Fig. 4(a).

The use of a four-pole filter allows us to realize a passband with a bandwidth of about 1 GHz with sharp band edges [see Fig. 4(b)]. The broad bandwidth makes the system resilient against mistargeting of the buffer frequency. Frequency offsets on the order of 100 MHz can be tolerated, since the buffer loss rate is maintained at around 10 MHz throughout much of the passband. The storage mode, however, is far outside of the filter passband and, therefore, is protected against loss through the filter. From electromagnetic simulations, we estimate that the induced loss rate on the storage mode due to the buffer loss channel is approximately 0.2 kHz, over 4 orders of magnitude smaller than the buffer loss rate of $\kappa_b/2\pi = 10.7$ MHz. This corresponds to a T_1 ceiling of > 750 μ s on the storage mode.

In Fig. 4(c), we demonstrate stabilization of the cat-qubit manifold by integrating the 3WM pump, linear drive, and loss on the buffer mode. In particular, we turn on the 3WM interaction with strength $g_2/2\pi = 578 \pm 4$ kHz (corresponding to $\kappa_2/2\pi = 124 \pm 2$ kHz; see Appendix I) and a linear buffer drive with a variable amplitude. As the strength of the linear buffer drive increases, the distance between the two steady states (i.e., the “size”) of the cat qubit increases. When the buffer drive is turned off while the 3WM pump is kept on, a pure two-photon dissipation $\kappa_2\mathcal{D}[\hat{a}^2]$ is implemented which confines the storage mode into the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ Fock-state manifold. Note that we use the term pure two-photon dissipation to refer to the dynamics due to $\kappa_2\mathcal{D}[\hat{a}^2]$ with $|\alpha| = 0$ to distinguish it from the case with $|\alpha| > 0$.

We study the lifetime (T_1) and coherence time (T_2) of the storage mode in this integrated system using the following procedure. We prepare the state $(|\hat{n} = 0\rangle + |\hat{n} = 1\rangle)/\sqrt{2}$ (up to state-preparation errors and phase) in the confined Fock-state manifold of the storage mode by initializing a large coherent state and applying a pure two-photon dissipation to the storage mode [18]. We then let the storage mode evolve, with the pure two-photon dissipation either turned off or kept on, for a variable delay. Ideally, the pure two-photon dissipation acts trivially within the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ manifold, so it should not affect the

storage-mode coherence. However, strong off-resonant driving of various unintended transitions can degrade the storage-mode coherence.

To characterize the storage T_1 time, after the variable delay time, we measure the photon-number parity of the storage mode and extract the decay rate of the parity observable. As shown in Fig. 4(d), the storage T_1 with the pure two-photon dissipation turned off is measured to be 76 ± 1 μ s. This relatively long T_1 value is achieved due to (i) the optimized deposition and etching of the tantalum superconducting thin film used in the fabrication of the chip which hosts the storage-mode coplanar waveguide resonators [27,37] and (ii) the effective protection of the storage mode by the four-pole filter. Furthermore, we observe that the storage T_1 is not degraded when the pure two-photon dissipation is kept on during the delay. This can be attributed to careful allocation of the storage and buffer mode frequencies to avoid undesired resonances under a strong 3WM flux pump. In particular, in the pumped buffer system other undesired sideband interactions can be driven by the flux pump, for example, terms such as $\hat{a}^\dagger\hat{b}^{\dagger 3} + \text{H.c.}$ Some of these interactions can lead to an increased parity decay rate of the storage mode, since they may create (or annihilate) an odd number of photons in the storage mode. In our system, the storage-buffer frequency configuration is carefully optimized to make these undesired sideband transitions sufficiently off resonant, hence avoiding degradation in T_1 under the 3WM flux pump (see Appendix H 3).

The storage T_2 is extracted through a Ramsey-style measurement where, after the preparation into $(|\hat{n} = 0\rangle + |\hat{n} = 1\rangle)/\sqrt{2}$ and variable delay, we apply a storage displacement and a final measurement of the storage mode’s vacuum state probability using a photon-number-selective pulse on the ancilla. As shown in Fig. 4(e), the storage-mode coherence time T_2 is measured to be over 100 μ s with the pure two-photon dissipation turned off. With the dissipation turned on, T_2 is only marginally degraded. These results demonstrate a highly coherent storage mode which can host a stabilized, phase-coherent cat qubit. See Appendix L for more details on the T_1 and T_2 measurements.

V. BUFFER-INDUCED STORAGE-MODE NONLINEARITIES

Realizing a cat qubit with an exponential noise bias and a low phase-flip error rate requires more than high storage-mode coherence times. It also requires minimizing undesired, parasitic nonlinearities in the storage mode. In particular, we identify the buffer-induced nonlinearities (e.g., the storage self-Kerr and the storage-buffer cross-Kerr) as the most crucial parameters. We then discuss how the serial inductances in the buffer mode can be utilized to minimize these parasitic nonlinearities.




FIG. 5. Buffer-induced storage-mode nonlinearities. (a) Circuit diagram of the buffer mode, including additional serial inductances (colored in red) in series with the side junctions. In the central path, we use a shallow three-element junction array. (b) Frequency of the buffer mode as a function of φ_Δ with $\varphi_\Sigma = \pi/2$. Adding serial inductances to the model is necessary to accurately reproduce the experimental data. (c) Numerical predictions based on a model where the underlying circuit parameters are tuned to match (b). We let E_J of the model vary and plot the predictions of storage self-Kerr (top) and the storage-buffer cross-Kerr (bottom) as a function of the average buffer side junction energy $E_J = (E_{J,1} + E_{J,2})/2$. The value of E_J extracted from fitting (b) (black dashed line) is close to the optimal value of E_J (blue dashed line). To verify the model, we measure the storage self-Kerr to be 0.97 ± 0.02 kHz (green horizontal line), which is close to the model predicted value of $K_s/2\pi = 1.1$ kHz at the independently extracted device E_J value.

In the presence of the storage-buffer cross-Kerr and storage self-Kerr, the Hamiltonian of the storage-buffer subsystem is given by

$$\hat{H} = \left(g_2 (\hat{a}^2 - \alpha^2) \hat{b}^\dagger + \text{H.c.} \right) + \chi_{sb} \hat{a}^\dagger \hat{a} \hat{b}^\dagger \hat{b} + \frac{K_s}{2} \hat{a}^{\dagger 2} \hat{a}^2. \quad (3)$$

The first two terms are the resonant 3WM term and linear drive on the buffer mode, which together stabilize a cat qubit. The last two terms are the undesired storage-buffer cross-Kerr and storage self-Kerr. By adding the buffer loss $\kappa_b \mathcal{D}[\hat{b}]$ and adiabatically eliminating the buffer mode, the effective master equation within the storage mode is approximately given by

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}_{\text{eff}}, \hat{\rho}] + \frac{4g_2^2}{\kappa_b} \mathcal{D} \left[\left(1 - \frac{2i\chi_{sb}}{\kappa_b} \hat{a}^\dagger \hat{a} \right) (\hat{a}^2 - \alpha^2) \right] \hat{\rho}, \quad (4)$$

with

$$\hat{H}_{\text{eff}} = -\frac{2g_2^2 \chi_{sb}}{\kappa_b^2} (\hat{a}^{\dagger 2} - \alpha^2) \hat{a}^\dagger \hat{a} (\hat{a}^2 - \alpha^2) + \frac{K_s}{2} \hat{a}^{\dagger 2} \hat{a}^2, \quad (5)$$

in the regime $\chi_{sb} \langle \hat{a}^\dagger \hat{a} \rangle / \kappa_b \ll 1$ (see Appendix H 4).

The effective master equation above clearly shows the detrimental effects of the storage-buffer cross-Kerr χ_{sb} . Specifically, the two-photon dissipation operator gets distorted by a factor of $(1 - 2i\chi_{sb} \hat{a}^\dagger \hat{a} / \kappa_b)$, and an additional term is added to the Hamiltonian \hat{H}_{eff} which acts non-trivially outside of the cat-qubit manifold. An intuitive reason for the particularly damaging effect of these undesired terms is as follows. Noise makes the storage mode

deviate from the cat-qubit manifold which then excites the buffer mode. This buffer excitation should be accompanied only by an ideal two-photon dissipation operator $\hat{a}^2 - \alpha^2$ on the storage mode. However, the buffer excitation additionally dephases the storage mode due to the storage-buffer cross-Kerr. This extra dephasing during the correction, captured by the distortion factor $(1 - 2i\chi_{sb} \hat{a}^\dagger \hat{a} / \kappa_b)$ and the added term in the Hamiltonian, can then significantly degrade the bit-flip performance of the cat qubit.

In practice, the storage self-Kerr and storage-buffer cross-Kerr originate from the self-Kerr of the buffer mode and the storage-buffer hybridization. In an ideal realization of the ATS buffer element, the self-Kerr of the buffer mode vanishes at the saddle points (i.e., first-order flux-insensitive points). However, as noted in previous work [16], an ATS in practice typically has nonzero self-Kerr on the saddle points due to various imperfections, such as the finite number of junctions in the junction array and the asymmetry in the side junction energies. Here, we demonstrate that as in other superconducting systems [38,39] it is important to consider additional inductances in series with the junctions of the circuit, as they can contribute significantly to the frequency and the self-Kerr of the buffer mode. Specifically, we consider inductances in series with the side junctions. These additional inductances contribute significantly to the frequency and the self-Kerr of the buffer mode. The additional nonlinearity is particularly important, because the other leading-order contributors are canceled out.

As shown in Fig. 5(a), the serial inductances introduce two extra flux nodes $\hat{\phi}_{P_1}$ and $\hat{\phi}_{P_2}$ to the system, in addition to the flux node $\hat{\phi}$ of the buffer mode. We focus on the regime where these serial inductances are much smaller than the inductances due to the side junctions, i.e., $E_{L,P_i} = \hbar^2 / (4e^2 L_{P_i}) \gg E_{J,i}$ for $i = 1, 2$. In this regime, the extra

nodes $\hat{\phi}_{P_1}$ and $\hat{\phi}_{P_2}$ can be treated classically and set to be the values that minimize the inductive potential function [38,40]. Then, we obtain an effective Hamiltonian (see Appendix H 2):

$$\begin{aligned} \hat{H} = & 4E_C \hat{N}^2 - NE_{J,\text{array}} \cos\left(\frac{\hat{\phi}}{N}\right) \\ & - E_{J,1} \cos(\hat{\phi} + \varphi_{x,1}) - E_{J,2} \cos(\hat{\phi} - \varphi_{x,2}) \\ & + \frac{E_{J,1}^2}{4E_{L,P_1}} \cos[2(\hat{\phi} + \varphi_{x,1})] + \frac{E_{J,2}^2}{4E_{L,P_2}} \cos[2(\hat{\phi} - \varphi_{x,2})], \end{aligned} \quad (6)$$

where the first two lines represent the Hamiltonian of an ATS as described in Ref. [16] and N is the number of junctions in the junction array. The last line represents the additional contributions from the serial inductances. Notably, these additional terms have a different flux periodicity from that of the usual ATS potential due to the extra factor of 2 in the arguments of the cosine terms.

Figure 5(b) provides direct experimental evidence for the contributions of the serial inductances in the buffer mode of our device. We show the frequency of the buffer mode (black curve) as a function of $\varphi_\Delta \equiv (\varphi_{x,1} - \varphi_{x,2})/2$ with $\varphi_\Sigma \equiv (\varphi_{x,1} + \varphi_{x,2})/2 = \pi/2$. Notably, a fitted model without the serial inductances (blue curve) fails to capture the correction that is periodic in φ_Δ with period π (this is in comparison to the 2π periodicity in an ideal ATS). On the other hand, the fitted model with the serial inductances (orange curve) captures this doubly periodic behavior. From the latter model, which assumes $L_{P_1} = L_{P_2} = L_P$, we extract a serial inductance of $L_P = 27\text{pH}$, corresponding to $E_{J,i}/E_{L,P_i} \lesssim 0.01$ in our device for $i = 1, 2$.

The effect of the serial inductances on the buffer frequency is useful for characterizing the inductances but is itself not critical to the cat-qubit performance. However, the impact of serial inductances on the self-Kerr of the buffer mode is crucial. For example, assuming symmetry between the two side junctions, i.e., $E_{J,1} = E_{J,2} = E_J$ and $E_{L,P_1} = E_{L,P_2} = E_{L,P}$, we find that the buffer self-Kerr is perturbatively given by

$$K_b = -\frac{E_C}{N^2} + \frac{8E_C E_J^2}{E_{L,P} E_L}, \quad (7)$$

on the saddle points [e.g., $(\varphi_\Sigma, \varphi_\Delta) = (\pi/2, \pi/2)$ and $(\varphi_\Sigma, \varphi_\Delta) = (\pi/2, 3\pi/2)$] and where $E_L \equiv E_{J,\text{array}}/N$ is the effective inductive energy of the junction array (see Appendix H 2). This expression shows that the serial inductances introduce positive self-Kerr on the buffer mode.

Previous works have used a large number of junctions in the junction array (e.g., $N = 5$ [16,17] or $N = 13$ [19]) to minimize the negative self-Kerr contribution $-E_C/N^2$ from

the junction array. On the other hand, our analysis suggests a different strategy. Given the serial inductance, we can optimize the circuit parameters of the buffer such that the negative self-Kerr from the junction array is canceled out by the positive self-Kerr from the serial inductances. In particular, this strategy allows us to use a shallow junction array with just $N = 3$, reducing the required junction energy $E_{J,\text{array}}$ to achieve a given target E_L .

In Fig. 5(c), we illustrate how nonlinearities can be minimized by carefully tuning the design value of E_J . We provide predictions of the storage self-Kerr K_s and the storage-buffer cross-Kerr χ_{sb} as a function of E_J using a model where the other underlying circuit parameters are tuned to match the buffer frequency data in Fig. 5(b). This model predicts that these buffer-induced storage-mode nonlinearities would vanish at an optimal average side junction energy of $E_J/h = 40$ GHz (blue dashed line) due to an ideal balancing of the positive and negative self-Kerr contributions on the buffer mode. For the device of this work, we extract an E_J/h of 57 GHz (black dashed line). Despite the offset from the optimal value, the predicted storage self-Kerr of $K_s/2\pi = 1.1$ kHz and the storage-buffer cross-Kerr of $\chi_{sb}/2\pi = 156$ kHz are still tolerable given the strength of the 3WM interaction $g_2/2\pi = 578 \pm 4$ kHz and the buffer decay rate $\kappa_b/2\pi = 10.7$ MHz. We independently measure the self-Kerr of the storage mode and find it to be $K_s/2\pi = 0.97 \pm 0.02$ kHz [green horizontal strip in Fig. 5(c); see also Appendix N], in good agreement with the model prediction.

VI. TWO-PHOTON DISSIPATION PERFORMANCE

Having presented the high coherence and low nonlinearity of the storage mode in Secs. III–V, we now proceed to characterize the performance of a cat qubit encoded in this mode. As shown in the diagram in Fig. 6(a), we characterize the bit-flip time of a cat qubit by initializing a coherent state $|+\alpha\rangle \simeq |0\rangle$ and then stabilizing the cat-qubit manifold with a 3WM pump and a drive on the buffer mode. Finally, using the tunable coupler, we perform parity measurements after $\pm\alpha$ displacements to read out the cat qubit in the $|0\rangle/|1\rangle$ basis (Z basis). The bit-flip time $T_{\text{bit-flip}} = 2T_Z$ is characterized by fitting these Z -basis populations to an exponential curve and then extracting the decay exponent T_Z (see Appendix F).

The dark blue data points in Fig. 6(a) show the bit-flip time of our cat qubit increases exponentially in the mean photon number $|\alpha|^2$ and begins to saturate at a maximum value of 0.3–0.6 s when $|\alpha|^2 \gtrsim 4$. We fit the data for $|\alpha|^2 \leq 4$ to the relation $T_{\text{bit-flip}} \propto e^{\gamma|\alpha|^2}/|\alpha|^2$ and find that the extracted exponent $\gamma = 2.24 \pm 0.09$ is comparable with the optimal value of $\gamma = 2$ expected under a white-noise dephasing model [20] (larger values of γ up to $\gamma = 4$ are possible with only single-photon loss in the absence of dephasing [11,41]). A fit to a model without the factor of

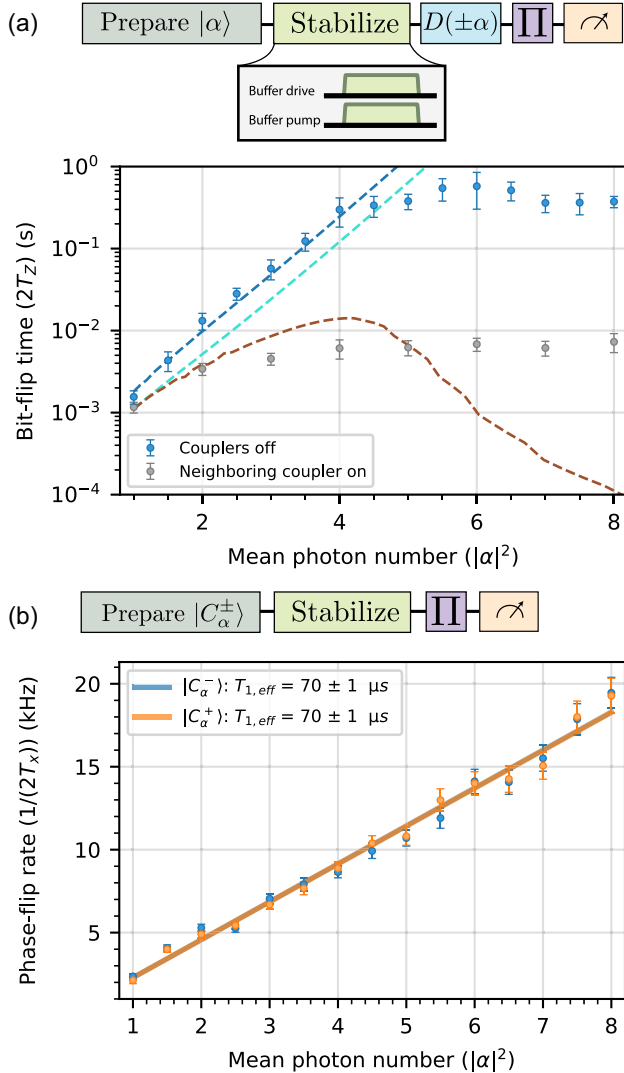


FIG. 6. Two-photon dissipation performance. (a) Bit-flip times of our cat qubit as a function of the mean photon number. The blue points are the bit-flip times of the cat qubit when all the couplers are in the off position. In the gray curve, we show the bit-flip times when the coupler in the neighboring unit cell is near its minimum frequency position ($\Phi/\Phi_0 = 0.4$). The dashed turquoise curve shows simulated bit-flip times based on the device parameters with couplers off. The dashed blue curve shows simulated bit-flip times where the white-noise dephasing parameter is phenomenologically tuned to match the experimental data. In the dashed brown curve, we show a simulation with twice larger E_J and L_P of the buffer where significant degradation is observed due to buffer-induced nonlinearities. (b) Phase-flip rates of the cat qubit. Lines correspond to a linear fit to the phase-flip rate, $\gamma_Z = \kappa_{1,eff}|\alpha|^2$, where $T_{1,eff} = 1/\kappa_{1,eff}$.

$|\alpha|^2$, $T_{\text{bit-flip}} \propto e^{\gamma'|\alpha|^2}$, which provides more direct intuition on the rate of scaling with photon number, yields $\gamma' = 1.75 \pm 0.06$. This corresponds to a factor of 5.8 increase in the bit-flip time for every added photon. Owing to the large scaling exponent and long storage-mode coherence times, we achieve a bit-flip time over 1 ms

with only $|\alpha|^2 = 1$, which increases to above 0.1 s with just $|\alpha|^2 = 4$.

The dashed turquoise curve in Fig. 6(a) shows the simulated bit-flip time based on the measured storage-mode coherence times in Sec. IV and the predicted (and measured) storage-mode nonlinearities in Sec. V. The simulation agrees approximately with the experimental data (up to a multiplicative offset) in the $|\alpha|^2 \leq 4$ regime, where the exponential bit-flip suppression is observed. The simulation with tuned white-noise dephasing parameter (dashed blue curve; see Appendix H 4) more closely captures the experimental data without the multiplicative offset. Finally, to show the importance of carefully selecting buffer parameters, we perform an additional simulation with the buffer mode having twice as large side junction energy E_J and serial inductance L_P (dashed brown curve). The simulation predicts poor bit-flip performance which is primarily due to the intolerably large storage self-Kerr and storage-buffer cross-Kerr (approximately 8 times larger than the device parameters). This demonstrates the importance of understanding the contributions of the buffer serial inductances to the nonlinearities of the storage-buffer system.

To confirm the effectiveness of the extinction of the coupling provided by the tunable coupler, we perform an experiment where we apply dissipative stabilization while we purposefully turn on the coupler between the storage mode and the transmon ancilla in a neighboring unit cell of our device. When the external flux of this one coupler is set to $0.4\Phi_0$ (i.e., close to the on position), the bit-flip time of the cat qubit under dissipative stabilization saturates at a value below 10 ms, as indicated by the gray data points in Fig. 6(a). This saturation in bit-flip time is due to heating of the coupler and the ancilla in the neighboring unit cell. With the coupler in the on position, a heating event in either the coupler or the ancilla gives rise to a dispersive shift in the storage-mode frequency larger than that of the (1-MHz-level) confinement rate of the dissipative stabilization over the range of mean photon number studied here. This leads to a loss of confinement of the storage mode and a high probability of a cat-qubit bit flip with each heating event [16]. With all of the couplers at their off position, the bit-flip time far exceeds the ancilla and coupler heating rate limit of 10 ms but eventually saturates at a maximum value of 0.3–0.6 s. We have not identified a definitive mechanism explaining this saturation behavior, and it remains an area of continued study (see Appendix O for more details).

While the bit-flip rates of a cat qubit are (ideally) suppressed exponentially with $|\alpha|^2$, the phase-flip rates of a cat qubit are expected to increase linearly, i.e., $\gamma_Z = \kappa_{1,eff}|\alpha|^2$, where $\kappa_{1,eff}$ is the effective single-photon loss rate. As shown in the diagram in Fig. 6(b), we characterize phase-flip rates of our cat qubit by initializing the storage mode into an even or odd cat state by displacing

the storage mode and performing a parity measurement. Then, we stabilize the cat qubit with the two-photon dissipation and finally perform a parity measurement to detect parity flips. The phase-flip rate is then extracted through an exponential fit of the parity decay curve. The results in Fig. 6(b) show that the phase-flip rate increases linearly in the mean-photon number $|\alpha|^2$ as expected. From linear fits to this curve, we find that the effective phase-flip lifetime is $T_{1,\text{eff}} = 1/\kappa_{1,\text{eff}} = 70 \pm 1 \mu\text{s}$. This value is consistent with an independent measurement of the storage-mode energy decay time of $T_1 = 74 \pm 1 \mu\text{s}$ (see Appendix L).

The results in this section demonstrate that our cat qubit is capable of achieving nearly ideal exponential bit-flip suppression in the regime of $|\alpha|^2 \lesssim 4$, where the phase-flip rates remain relatively low. In particular, we estimate that, for a typical error-correction cycle length of approximately $2 \mu\text{s}$, the phase-flip probabilities of our cat qubit are approximately 3% at $|\alpha|^2 = 1$ and 11% at $|\alpha|^2 = 4$. Since the error threshold of a repetition code is approximately 11% [42] (assuming a phenomenological noise model), the phase-flip error rates achieved are sufficiently low to be corrected by a repetition code in the regime where $|\alpha|^2 \lesssim 4$. At the same time, we observe a large noise bias ranging from 10 to over 1000 across this same range of average photon numbers. Thus, the observed phase-flip and bit-flip times meet the requirements for more complex operations with cat qubits, such as multiqubit gates and error-correction experiments with an outer repetition code.

VII. PULSED STABILIZATION

In error-correction architectures involving cat qubits, it is often required to turn off two-photon dissipation on a cat qubit for certain syndrome-extraction operations. For example, implementation of a CNOT gate with a cat qubit as the target typically relies on rotating the target cat qubit conditioned on the state of a control qubit [7,11,43–45]. For this conditional rotation to work well for a dissipative cat qubit, one needs to either engineer a complex two-photon dissipation scheme with its stabilized manifold rotating conditioned on the state of a control qubit [7,11] or turn off the two-photon dissipation such that the target cat qubit can freely rotate under a separate mechanism [44,45]. In the latter case, it is essential to ensure that the bit-flip performance of the cat qubit is not significantly degraded when the two-photon dissipation is turned off for some fraction of the time during a cycle, a situation we refer to as pulsed stabilization. This is a challenge, because storage-mode errors due to mechanisms such as undesired nonlinearities and dephasing are allowed to freely accumulate when two-photon dissipation is turned off. While these errors are unlikely to lead to phase flips because these errors and the subsequent two-photon dissipation recovery all preserve storage-mode parity, significant excursions in

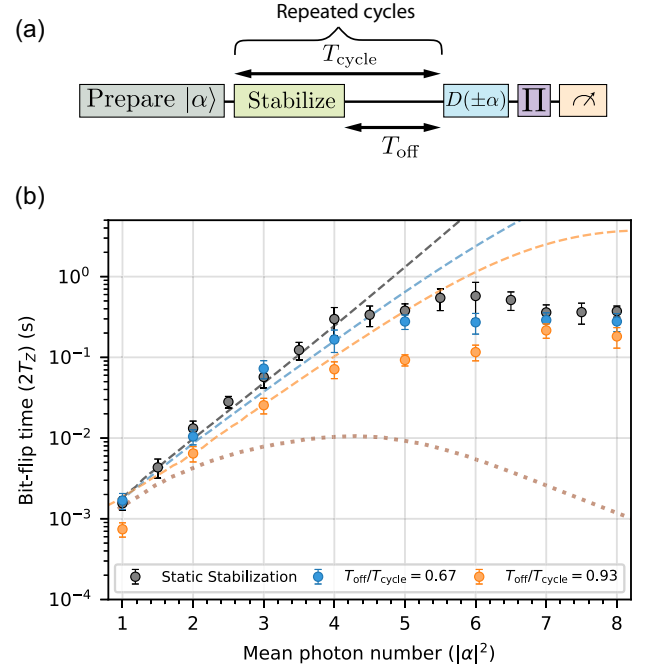


FIG. 7. Pulsed stabilization. (a) A high-level depiction of the measurement sequence for pulsed cat-qubit stabilization. Unlike previous experiments where we have the two-photon dissipation continually turned on, here we repeatedly apply pulsed stabilization cycles. In each cycle of length T_{cycle} , the stabilization is turned off for a duration T_{off} . (b) Bit-flip times of our cat qubit under pulsed stabilization as a function of the mean photon number $|\alpha|^2$ for various duty cycle ratios of $T_{\text{on}}/T_{\text{cycle}}$ for $T_{\text{cycle}} = 3 \mu\text{s}$. The markers correspond to the measured data of a baseline static stabilization (black), pulsed stabilization with $T_{\text{on}} = 1 \mu\text{s}$ (blue), and pulsed stabilization with $T_{\text{on}} = 200 \text{ ns}$ (orange). The dashed lines correspond to simulations for each of the three cases. The dotted brown line corresponds to a simulation with $T_{\text{on}} = 1 \mu\text{s}$ where the storage self-Kerr is increased to $K_s/2\pi = 5 \text{ kHz}$, corresponding to a value 5 times that of the self-Kerr predicted and measured for the device.

phase space can be recovered to the incorrect coherent state leading to bit-flip errors.

In this section, we study the bit-flip performance of our cat qubit under pulsed stabilization. That is, as shown in Fig. 7(a), we initialize a coherent state and then repeatedly apply a cycle of pulsed stabilization with cycle time T_{cycle} in which the two-photon dissipation is pulsed on and off for a duration of T_{on} and $T_{\text{off}} = T_{\text{cycle}} - T_{\text{on}}$, respectively. Then, we characterize the bit-flip rate of our cat qubit by measuring the displaced parities similarly as in Sec. VI. When calibrating the pulsed stabilization, we ensure that the phase of the storage mode is carefully tracked so that two-photon dissipation is aligned with the storage-mode coherent states when it turns back on.

In Fig. 7(b), we show the bit-flip time of our cat qubit as a function of the mean photon number $|\alpha|^2$ for various duty cycle ratios, $T_{\text{on}}/T_{\text{cycle}}$, with a fixed cycle time of

$T_{\text{cycle}} = 3 \mu\text{s}$. We include the baseline case of static stabilization where the two-photon dissipation is continually turned on (black). Relative to the baseline, we observe no degradation in the bit-flip time even when the dissipation is turned on for only $1 \mu\text{s}$, or one-third of the time during the pulsed stabilization cycle (blue curve). Note that, unsurprisingly, the bit-flip times do not exceed the saturated value observed in the static stabilization case. Significant degradation of the cat-qubit bit-flip times under pulsed stabilization is observed only in an extreme limit of $T_{\text{on}}/T_{\text{cycle}} = 7\%$, when the pulsed stabilization is turned on for only $T_{\text{on}} = 200 \text{ ns}$ of the $3 \mu\text{s}$ cycle length.

Before performing a more quantitative comparison to the data, we provide an intuitive picture for the preservation of bit-flip times under pulsed stabilization. The protection against bit flips afforded by the cat-qubit encoding owes to the exponential separation between the $|\pm\alpha\rangle$ coherent states. Two-photon dissipation should be viewed as a mechanism to preserve this separation. When the stabilization is turned off, leakage out of the cat-qubit manifold due to nonlinearities and decoherence are allowed to accumulate. In contrast, if the stabilization is continually on, leakage due to these mechanisms is brought back to the cat-qubit manifold quickly, preventing errors from accumulating. In this way of viewing it, the preservation of the cat-qubit bit-flip times even with the stabilization off for significant periods should be attributed to the low decoherence and nonlinearities of the storage mode. Even though leakage is allowed to accumulate, on timescales of $2 \mu\text{s}$, the accumulation is not sufficiently high that upon performing dissipative recovery bit flips occur.

The dashed curves correspond to simulations of the pulsed stabilization procedure for different duty cycles of stabilization, $T_{\text{on}}/T_{\text{cycle}}$. These simulations use the same refined value for the white-noise dephasing parameter as in the dashed blue curve in Fig. 6(a). For all but the dotted brown curve, we also assume in these simulations the device estimated self-Kerr of $K_s/2\pi = 1.1 \text{ kHz}$ and storage-buffer cross-Kerr of $\chi_{sb}/2\pi = 156 \text{ kHz}$. We find that, for a duty cycle of $T_{\text{on}}/T_{\text{cycle}} = 1/3$ (blue), the degradation in the bit-flip times is relatively marginal, consistent with our measurements. When the pulsed stabilization is reduced to a 7% duty cycle (orange), the simulations predict a more significant reduction in the bit-flip time relative to the static stabilization case (black), especially at large values of mean photon number $|\alpha|^2 \gtrsim 6$. Although this is consistent with our measured observations, the simulation model does not capture the previously discussed, and currently unexplained, measured saturation in bit-flip time at 0.3–0.6 s. The dotted brown curve corresponds to an additional simulation of the case $T_{\text{on}}/T_{\text{cycle}} = 1/3$, where we increase the storage self-Kerr nonlinearity to $K_s/2\pi = 5 \text{ kHz}$ (the cross-Kerr is left at $\chi_{sb}/2\pi = 156 \text{ kHz}$). This self-Kerr is 5 times larger than

what we predict and measure in our device (see Sec. V) but representative if the buffer-mode nonlinearities are not carefully minimized. In this case, we compute that the bit-flip times are significantly degraded and remain below approximately 10 ms due to the larger distortion caused by the increase in storage self-Kerr. This reinforces the importance of the accurate buffer modeling in Sec. V and the inclusion of serial inductances in the ATS, which ensure that the buffer-induced storage-mode nonlinearities stay in check.

VIII. CONCLUSION

The work presented here sets the stage for performing complex multiqubit gates and error-correction experiments involving dissipatively stabilized cat qubits. The low buffer-mode nonlinearities and high storage-mode coherence allow us to achieve long cat-qubit bit-flip times with small mean photon numbers. For typical error-correction cycle times, the phase-flip error probability of our cat qubit can stay below the error threshold of an outer repetition code up to the mean photon number of $|\alpha|^2 \simeq 4$, where bit-flip times exceed 0.1 s. We can also retain the benefits of exponentially suppressed bit-flip error rates even when the stabilization is turned off for a significant fraction of a pulsed-stabilization cycle. Additionally, the cat qubits of this work are implemented using planar superconducting quantum circuits and are amenable to scalable micro-fabrication processes. This scalability naturally allows for concatenating many such cat qubits into an outer error-correcting code, a compelling architecture for reducing the resource overhead associated with quantum error correction [7–9, 11–14].

Going forward, it will be important to drive the noise bias of cat qubits many orders of magnitude higher, from the 10^3 level of this work to levels approaching 10^7 , in order to fully realize the potential for hardware-efficient quantum error correction at algorithmically relevant error rates [45]. Crucially, the cat-qubit phase-flip rates still need to remain sufficiently low even at larger cat-qubit mean photon numbers needed for realizing such a large noise bias. Thus, it will also be important to push storage-mode T_1 lifetimes to 1 ms and beyond to accommodate larger values of cat-qubit mean photon number. Critical to this pursuit, the strategies presented in this work can be used to accommodate higher storage coherence times and achieve stronger two-photon dissipation rates without introducing parasitic effects. Future work to uncover and mitigate the mechanism(s) responsible for the observed saturation of bit-flip times at the 1-s level in our current devices will also be important. More generally, pushing the limits of cat-qubit bit-flip times and noise bias will open up new opportunities to study error mechanisms in superconducting quantum circuits that are rare events and occur on long timescales [46, 47].

ACKNOWLEDGMENTS

We thank Alex Retzker and Yufeng Ye for helpful comments on the manuscript and the staff from across the AWS Center for Quantum Computing that enabled this project. We also thank Fiona Harrison, Harry Atwater, David Tirrell, and Tom Rosenbaum at Caltech and 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: DEVICE FABRICATION

The cat-qubit system described in the main text is implemented as one unit cell of a larger integrated device. The device consists of two chip dies which are fabricated separately on high-resistivity silicon and then flip-chip bonded together [25,26]. The first (“qubit”) chip die contains aluminum metallization layers, including Al/AIOx/Al Josephson junctions which form the nonlinear circuit elements in the transmon ancilla and couplers, as well as the ATS buffer. The Josephson junctions leads, deposited with a double-angle evaporation process, are respectively shorted to the aluminum ground plane and island capacitors with an added “bandage” contact [48,49]. The second chip die is metallized with thin-film alpha phase tantalum (thickness approximately 100 nm), which has lower microwave losses than aluminum and is capable of high-coherence storage modes in coplanar waveguide resonators [27,50]. The patterned circuit consists of storage resonators and other linear elements such as the metamaterial filter and readout resonators. Decoupling fabrication of the constituent dies with separate process flows allows us to take advantage of tantalum for achieving high-coherence storage modes without the constraints imposed by integrating Al/AIOx/Al Josephson junctions on the same chip die.

APPENDIX B: WIRING DIAGRAM AND EXPERIMENTAL SETUP

A diagram including the fridge wiring and room-temperature setup for one unit cell is shown in Fig. 8.

APPENDIX C: TABLE OF DEVICE PROPERTIES

Tables containing the device properties are given in Tables I and II.

APPENDIX D: CALIBRATION DETAILS

The buffer is flux biased to a saddle point, and the couplers neighboring the storage mode are flux biased to their maximum frequency positions. To calibrate the dispersive interaction between the ancilla and storage, we first select a flux pulse amplitude for the coupler. We initially start with a flux pulse amplitude which brings the coupler

to its minimum frequency position. We calibrate the interaction frequency of the ancilla which is shifted from the idling ancilla frequency due to the increased hybridization with the coupler. With these calibrations completed, we have the capability to excite the ancilla in a storage-photon-number-selective manner. Using the photon-number-selective ancilla pulses, we perform the initial calibration of the storage interaction frequency and flux pulse lengths using the storage conditional phase experiment described in Appendix J. For parity measurements, the ancilla phase must also be carefully tracked during the flux pulse, so we calibrate an ancilla phase correction used for this purpose.

Using the coupler’s minimum frequency position is impractical in our system, as the strong couplings make the dispersive nature of the storage-ancilla interaction break down (see Appendix H 1). We, thus, calibrate a few additional operating points of the coupler besides its minimum frequency position including the “parity measurement” and “number splitting” positions whose parameters are provided in Table I. The parity measurement position is used for Wigner tomography and the final storage-mode measurements in bit-flip and phase-flip measurements in Fig. 6. The number splitting position is used in the number splitting measurements in Fig. 3 and T_2 measurements in Fig. 4. Calibration of these additional operating points follows a similar procedure used for calibrating the minimum frequency position. For the number splitting measurements, which take longer than parity measurements and are more sensitive to flux pulse nonidealities (for example, long timescale distortions), we fine-tune the ancilla interaction frequency with a spectroscopy measurement.

For both the parity measurements and number-selective measurements, we use flux pulses with net-zero [51] waveforms with Gaussian shoulders. The number-selective ancilla pulses themselves have a Gaussian shape with a pulse length of 1.4 μ s and are applied during the second half of the net-zero flux pulse. Once parity measurements are calibrated, we can calibrate the storage-mode displacement as explained in Appendix M.

For the initial two-photon dissipation calibration, we measure the buffer 3WM condition using a vector network analyzer. The two-photon dissipation is fine-tuned to match the Stark-shifted storage frequency under the 3WM buffer pump by performing a storage Ramsey experiment with a pure two-photon dissipation applied. This calibration is performed frequently to ensure that the pump and drive frequencies of the two-photon dissipation stay up to date with the Stark-shifted storage-mode frequency [i.e., $(\omega_p + \omega_d)/2 \simeq \bar{\omega}_s$].

The buffer drive frequency (ω_d) should ideally match the Stark-shifted buffer frequency ($\bar{\omega}_b$). In practice, the two-photon dissipation performs comparably with an appreciable buffer detuning ($\bar{\omega}_b - \omega_d$) on the order of 1 MHz,

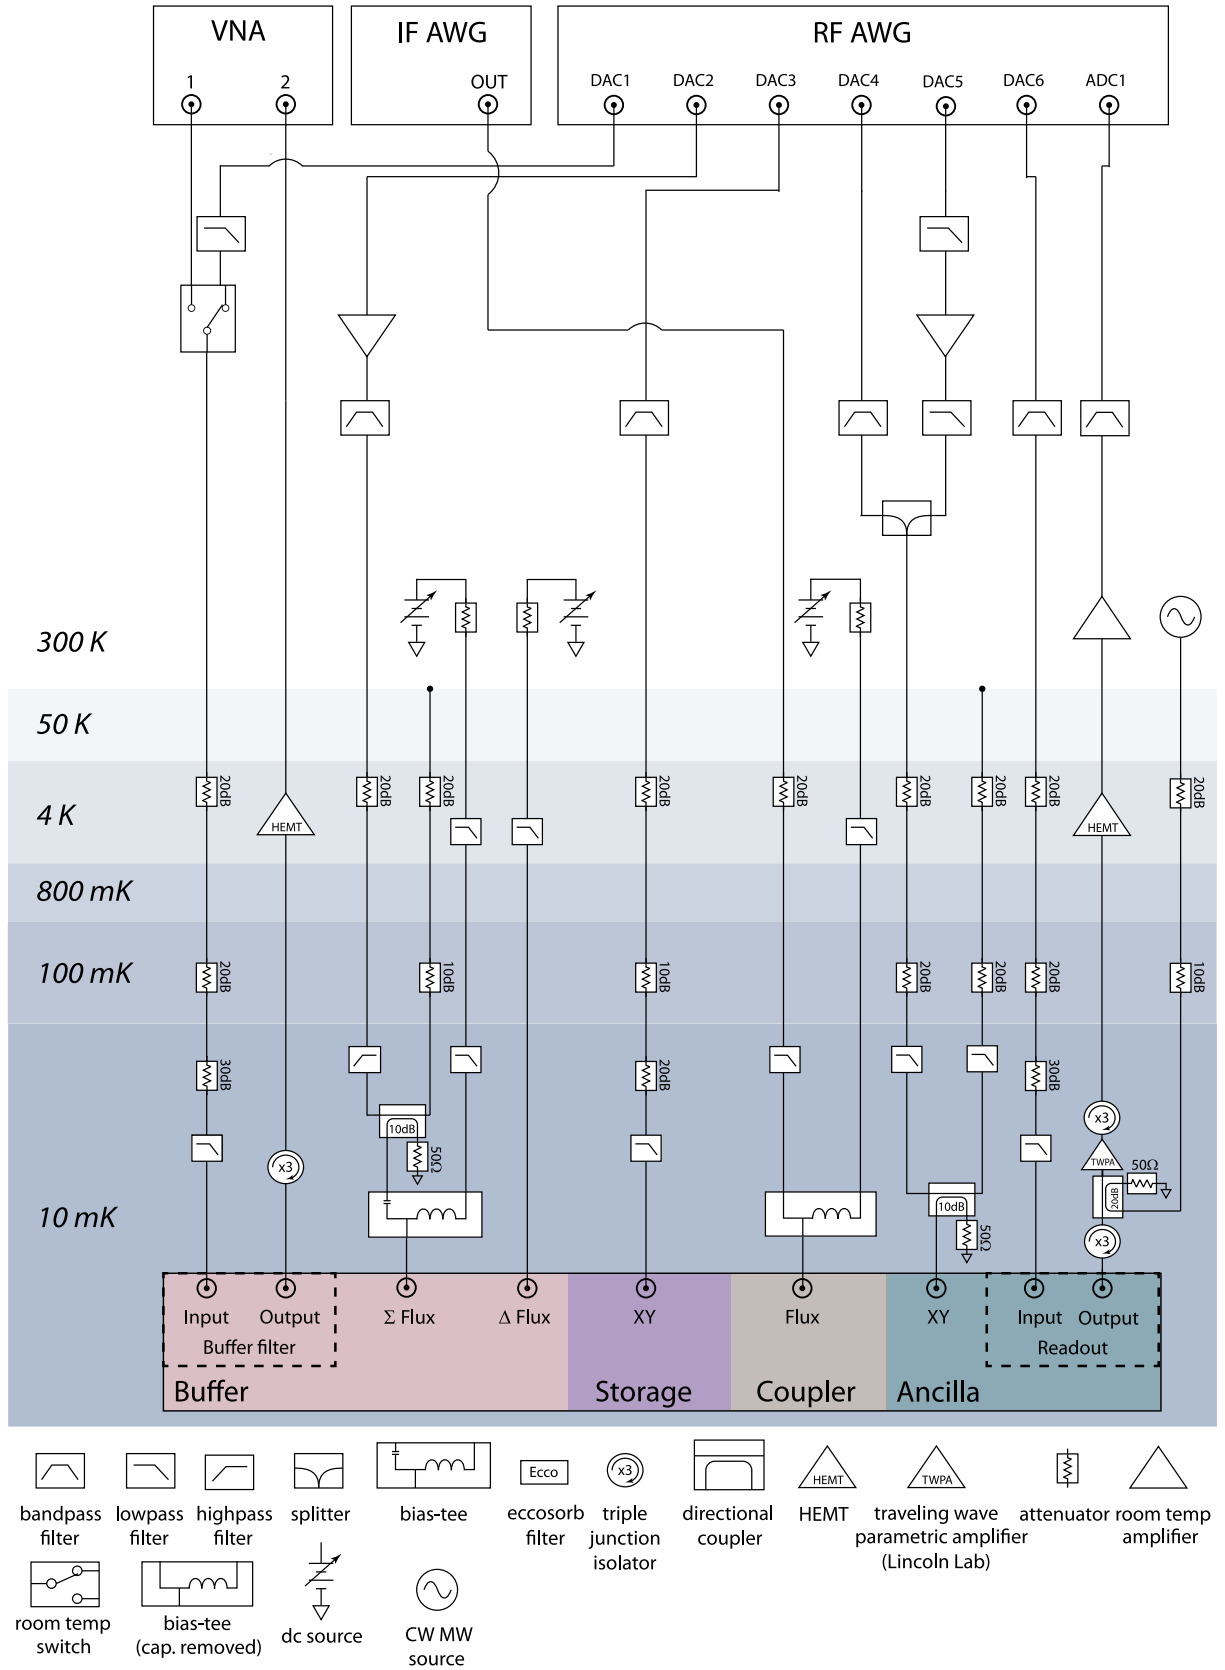


FIG. 8. Wiring diagram. Some details not included are eccosorb filters on almost all of the lines and room-temperature switches to toggle between characterization of different components.

TABLE I. Table of device properties for the unit cell used to characterize the storage mode.

Property	Descriptor (reference)	Unit	Value
Buffer properties			
$\omega_b/2\pi$	Buffer frequency	GHz	3.01
$\kappa_b/2\pi$	Buffer loss rate	MHz	10.7
$(\omega_b - \omega_d)/2\pi$	Buffer frequency—drive frequency	MHz	29.3
Coupler maximum frequency (idle) position properties			
$\Phi_{x,c}/\Phi_0$	Coupler flux value		0
$\omega_c/2\pi$	Coupler frequency	GHz	8.40
$\omega_a/2\pi$	Ancilla frequency	GHz	5.21
$\omega_s/2\pi$	Storage frequency	GHz	5.35
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	kHz	-1.4
$\omega_r/2\pi$	Ancilla readout frequency	GHz	7.23
$K_s/2\pi$	Storage Kerr	kHz	0.97
$(\omega_s - \bar{\omega}_s)/2\pi$	Storage Stark shift under 3WM	MHz	0.39
Parity measurement position properties			
Φ/Φ_0	Coupler flux value		0.38
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	MHz	-0.9
Number splitting position properties			
Φ/Φ_0	Coupler flux value		0.43
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	MHz	-2.4
Coupler minimum frequency position properties			
Φ/Φ_0	Coupler flux value		0.5
$\omega_c/2\pi$	Coupler frequency	GHz	5.77
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	MHz	-5.7

TABLE II. Abridged table of device properties for the additional unit cell containing the coupler used in Fig. 6(a).

Property	Descriptor (reference)	Unit	Value
Coupler maximum frequency (idle) position properties			
$\Phi_{x,c}/\Phi_0$	Coupler flux value		0
$\omega_c/2\pi$	Coupler frequency	GHz	8.9
$\omega_a/2\pi$	Ancilla frequency	GHz	5.2
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	kHz	-2.0
Coupler minimum frequency position properties			
Φ/Φ_0	Coupler flux value		0.5
$\omega_c/2\pi$	Coupler frequency	GHz	5.8
$\chi_{sa}/2\pi$	Storage-ancilla cross-Kerr	MHz	-4.3

which is to be compared against the buffer decay rate κ_b . As a result, we only coarsely optimize the buffer drive frequency to 1-MHz-level precision. The calibration of the linear relationship between the average photon number $|\alpha|^2$ of the cat qubit and the buffer drive amplitude is described in Appendix M.

APPENDIX E: BASIS CONVENTION

We define the complementary basis states $|\pm\rangle$ of the cat qubit as the even and odd cat states, i.e.,

$$|\pm\rangle \equiv |C_\alpha^\pm\rangle = \frac{|+\alpha\rangle \pm |-\alpha\rangle}{\sqrt{2(1 \pm e^{-2|\alpha|^2})}}. \quad (\text{E1})$$

Then, through the Hadamard transformation, the computational basis states approach the coherent states $|\pm\alpha\rangle$ in the large $|\alpha|$ limit where $e^{-2|\alpha|^2}$ becomes negligible:

$$\begin{aligned} |0\rangle &\equiv \frac{1}{\sqrt{2}}(|+\rangle + |-\rangle) \xrightarrow{e^{-2|\alpha|^2} \ll 1} |+\alpha\rangle, \\ |1\rangle &\equiv \frac{1}{\sqrt{2}}(|+\rangle - |-\rangle) \xrightarrow{e^{-2|\alpha|^2} \ll 1} |-\alpha\rangle. \end{aligned} \quad (\text{E2})$$

Note that, in the opposite limit where $|\alpha|$ approaches zero, we have

$$\begin{aligned} |+\rangle &\xrightarrow{|\alpha| \rightarrow 0} |\hat{n} = 0\rangle, \\ |-\rangle &\xrightarrow{|\alpha| \rightarrow 0} |\hat{n} = 1\rangle, \end{aligned} \quad (\text{E3})$$

where $|\hat{n} = 0\rangle$ and $|\hat{n} = 1\rangle$ are the vacuum and single-photon states, respectively (i.e., eigenstates of the number operator $\hat{n} = \hat{a}^\dagger \hat{a}$). This implies that, in the $|\alpha| \rightarrow 0$ limit, the computational basis states of a cat qubit are given by

$$\begin{aligned} |0\rangle &\equiv \frac{1}{\sqrt{2}}(|+\rangle + |-\rangle) = \frac{1}{\sqrt{2}}(|\hat{n} = 0\rangle + |\hat{n} = 1\rangle), \\ |1\rangle &\equiv \frac{1}{\sqrt{2}}(|+\rangle - |-\rangle) = \frac{1}{\sqrt{2}}(|\hat{n} = 0\rangle - |\hat{n} = 1\rangle), \end{aligned} \quad (\text{E4})$$

i.e., equal superposition of the vacuum and single-photon states with the ± 1 phases.

In this basis convention, a single-photon loss operator \hat{a} maps the $|+\rangle$ state into the $|-\rangle$ state and vice versa, leading to phase-flip (Z) errors. A Pauli- X measurement (i.e., $|+\rangle/|-\rangle$ basis measurement) of a cat qubit can be performed by reading out the photon-number parity $\exp[i\pi\hat{n}]$, since the $+1$ (or -1) eigenstate of \hat{X} has an even (or odd) number of photons. Since the computational basis states are approximately given by two coherent states $|\pm\alpha\rangle$ in the large $|\alpha|$ limit, a Pauli- Z measurement (i.e., $|0\rangle/|1\rangle$ basis measurement) can be performed by measuring a displaced photon-number parity of the storage mode [16,22].

APPENDIX F: ERROR RATE CONVENTION

To define error rates and decay times of a cat qubit, we consider a simple Lindblad master equation of the form

$$\frac{d\hat{\rho}(t)}{dt} = \gamma_X \mathcal{D}[\hat{X}]\hat{\rho}(t) + \gamma_Y \mathcal{D}[\hat{Y}]\hat{\rho}(t) + \gamma_Z \mathcal{D}[\hat{Z}]\hat{\rho}(t), \quad (\text{F1})$$

where $\hat{\rho}(t)$ is a qubit density operator and \hat{X} , \hat{Y} , and \hat{Z} are the qubit Pauli operators. Under this master equation,

an initial qubit state $\hat{\rho}(0)$ is depolarized by a Pauli error channel

$$\hat{\rho}(t) = [1 - p_X(t) - p_Y(t) + p_Z(t)]\hat{\rho}(0) + p_X(t)\hat{X}\hat{\rho}(0)\hat{X} + p_Y(t)\hat{Y}\hat{\rho}(0)\hat{Y} + p_Z(t)\hat{Z}\hat{\rho}(0)\hat{Z}, \quad (\text{F2})$$

where the Pauli error probabilities are given by

$$\begin{aligned} p_X(t) &= \frac{1}{4}(1 - e^{-2(\gamma_X + \gamma_Y)t} + e^{-2(\gamma_Y + \gamma_Z)t} - e^{-2(\gamma_X + \gamma_Z)t}), \\ p_Y(t) &= \frac{1}{4}(1 - e^{-2(\gamma_X + \gamma_Y)t} - e^{-2(\gamma_Y + \gamma_Z)t} + e^{-2(\gamma_X + \gamma_Z)t}), \\ p_Z(t) &= \frac{1}{4}(1 + e^{-2(\gamma_X + \gamma_Y)t} - e^{-2(\gamma_Y + \gamma_Z)t} - e^{-2(\gamma_X + \gamma_Z)t}). \end{aligned} \quad (\text{F3})$$

Then, it follows that the expectation values of the Pauli \hat{Z} and \hat{X} operators decay exponentially in time, i.e.,

$$\begin{aligned} \langle \hat{Z} \rangle(t) &= e^{-2(\gamma_X + \gamma_Y)t} \langle \hat{Z} \rangle(0), \\ \langle \hat{X} \rangle(t) &= e^{-2(\gamma_Z + \gamma_Y)t} \langle \hat{X} \rangle(0). \end{aligned} \quad (\text{F4})$$

These observables $\langle \hat{Z} \rangle(t)$ and $\langle \hat{X} \rangle(t)$ are directly measurable in experiments. Specifically, we measure $\langle \hat{Z} \rangle(t)$ based on the population difference in the two coherent states $|\pm a\rangle$ which are measured via displaced parity measurements. Moreover, $\langle \hat{X} \rangle(t)$ is measured through a parity measurement. Then, we fit the measured $\langle \hat{Z} \rangle(t)$ and $\langle \hat{X} \rangle(t)$ curves to an exponential decay curve and define T_Z and T_X as the decay constants. For $\langle \hat{Z} \rangle(t)$, we do not include a constant offset in the fit [i.e., $\langle \hat{Z} \rangle(t) = A \exp[-t/T_Z]$], whereas for $\langle \hat{X} \rangle(t)$ a constant offset is included [i.e., $\langle \hat{X} \rangle(t) = A \exp[-t/T_X] + B$]. A constant offset is needed in the latter case, because in the small $|\alpha|^2$ regime (e.g., $|\alpha|^2 = 1$), the steady-state parity $\langle \hat{X} \rangle(t \rightarrow \infty)$ is nonzero due to asymmetric phase-flip rates between the $|+\rangle \rightarrow |-\rangle$ and $|-\rangle \rightarrow |+\rangle$ transitions (e.g., in the extreme limit of $|\alpha|^2 = 0$, $|+\rangle \rightarrow |-\rangle$ is unlikely to happen, because $|+\rangle$ is given by the vacuum state, but the single-photon state $|-\rangle$ can easily decay to the vacuum state $|+\rangle$).

For a stabilized cat qubit with $|\alpha|^2 \gg 1$, we anticipate that the noise model is well described by the above Lindblad master equation with $\gamma_Z \gg \gamma_X \gg \gamma_Y$ due to the biased-noise structure. Thus, we approximately have

$$T_Z = \frac{1}{2(\gamma_X + \gamma_Y)} \simeq \frac{1}{2\gamma_X}, \quad T_X = \frac{1}{2(\gamma_Z + \gamma_Y)} \simeq \frac{1}{2\gamma_Z}, \quad (\text{F5})$$

and the Pauli error probabilities are given by

$$p_X(t) \simeq \gamma_X t, \quad p_Z(t) \simeq \gamma_Z t, \quad (\text{F6})$$

in the short time limit (i.e., $\gamma_Y t \ll \gamma_X t, \gamma_Z t \ll 1$). Based on the relation in Eq. (F6), we refer to γ_X and γ_Z as the bit-flip

and phase-flip rates, respectively, i.e., added bit-flip and phase-flip probability per unit time. Correspondingly, we define the bit-flip and phase-flip times to be

$$\begin{aligned} T_{\text{bit-flip}} &= \frac{1}{\gamma_X} = 2T_Z, \\ T_{\text{phase-flip}} &= \frac{1}{\gamma_Z} = 2T_X, \end{aligned} \quad (\text{F7})$$

respectively, where the last equality in each line is due to Eq. (F5). Thus, in our error rate convention, the bit-flip and phase-flip times are given by twice the decay times of the measured $\langle \hat{Z} \rangle(t)$ and $\langle \hat{X} \rangle(t)$ curves, respectively.

APPENDIX G: ERROR BAR CONVENTIONS

The mean and standard deviation of repeated measurements are computed using the posterior beta distribution. Parameters reported from fits are standard error. Fits are weighted by the inverse of the uncertainty on the points.

APPENDIX H: DEVICE HAMILTONIAN

In the rotating frame, the desired effective Hamiltonian of our system is given by

$$\hat{H} = \left(g_2(t)(\hat{a}^2 - \alpha^2)\hat{b}^\dagger + \text{H.c.} \right) + \chi_{sa}[\Phi_{x,c}(t)]\hat{a}^\dagger \hat{a} |e_a\rangle \langle e_a|, \quad (\text{H1})$$

where \hat{a} and \hat{b} are the annihilation operators of the storage and the buffer, respectively, and $|e_a\rangle$ is the first excited state of the ancilla transmon. $|\alpha|^2$ is approximately the mean photon number of a cat qubit. $g_2(t)$ is the strength of the desired 3WM mixing process which is a crucial ingredient of the two-photon dissipation process. $|g_2(t)|$ is controlled by the amplitude of the sigma flux pump on the buffer mode. The 3WM process is continuously turned on [i.e., $|g_2(t)| > 0$] in “static” stabilization of a cat qubit (Sec. VI) and is periodically turned on and off in a pulsed stabilization scenario (Sec. VII). In both cases, the coupler flux is nominally set to $\Phi_{x,c}(t) = 0$ (i.e., parking the coupler at its maximum frequency) to minimize the storage-ancilla dispersive coupling χ_{sa} while the cat qubit is being stabilized. Finally, at the end of an experimental sequence, the 3WM mixing process is turned off [i.e., $g_2(t) = 0$] and a flux pulse is applied to the coupler [i.e., $|\Phi_{x,c}(t)| > 0$] to turn on the storage-ancilla dispersive coupling χ_{sa} . This dispersive coupling enables the characterization of the storage mode through the ancilla transmon (Sec. III).

Notable imperfections that are not included in the desired effective Hamiltonian above are the storage-buffer cross-Kerr $\chi_{sb}\hat{a}^\dagger \hat{a} \hat{b}^\dagger \hat{b}$, storage self-Kerr [e.g., either $(K_s/2)\hat{a}^{\dagger 2} \hat{a}$ when the coupler is in the off position or ancilla-state-dependent storage self-Kerr when the coupler is in an on position], and the storage-coupler dispersive coupling

$\chi_{sc}[\Phi_{x,c}(t)]\hat{a}^\dagger\hat{a}|e_c\rangle\langle e_c|$. Besides these imperfections, undesired resonances can occur due to sideband transitions of a strong drive or large photon numbers in the storage mode. In the following subsections, we provide circuit-quantization-level details of the underlying coupling mechanisms and discuss these imperfections. Moreover, we provide systematic design strategies to suppress these undesired processes.

Appendix H1 provides more details on the tunable dispersive coupling, important considerations for minimizing the storage self-Kerr, and possible undesired resonances when the coupler frequency gets too close to the storage and the ancilla frequencies. Appendix H2 focuses on details related to the serial inductances of the buffer modes and how they affect the storage-buffer cross-Kerr and the storage self-Kerr. Appendix H3 discusses how the strong sigma flux pump needed for the 3WM process can lead to undesired sideband resonances if the system parameters are not carefully arranged. Lastly, in Appendix H4, we provide more details on the numerical simulation of the cat-qubit bit-flip times based on the effective Hamiltonian of our device. In Appendix H4, we also discuss the effective model for two-photon dissipation with added Kerr nonlinearities.

1. Tunable dispersive coupling between a storage mode and an ancilla transmon

The tunable dispersive interaction between a storage mode and an ancilla transmon is realized by using a tunable-transmon coupler. The Hamiltonian of the storage-coupler-ancilla subsystem is given by

$$\begin{aligned} \hat{H} = & \omega_{s,0}\hat{a}^\dagger\hat{a} + 4E_{C,c}\hat{N}_c^2 - E_{J1,c}\cos(\hat{\phi}_c + \varphi_{x,c}) \\ & - E_{J2,c}\cos\hat{\phi}_c + 4E_{C,a}\hat{N}_a^2 - E_{J,a}\cos\hat{\phi}_a \\ & + \lambda_{sc}i(\hat{a}^\dagger - \hat{a})\hat{N}_c + \lambda_{sa}\hat{N}_c\hat{N}_a + \lambda_{sa}i(\hat{a}^\dagger - \hat{a})\hat{N}_a. \end{aligned} \quad (\text{H2})$$

Here, the subscripts “s,” “c,” and “a” represent the storage, coupler, and the ancilla, respectively. Note that *a* refers to the ancilla only when it is used as a subscript and it should not be confused with the annihilation operator \hat{a} of the storage. In isolation, the storage mode is a linear quantum harmonic oscillator with a bare frequency $\omega_{s,0}$. The coupler is a tunable-frequency transmon and, thus, has a SQUID loop with two Josephson junctions with the junction energies $E_{J1,c}$ and $E_{J2,c}$. $\varphi_{x,c} \equiv 2\pi\Phi_{x,c}/\Phi_0$ is the dimensionless external flux applied to the coupler’s SQUID loop, where $\Phi_0 = h/e$ is the flux quantum. The ancilla is a fixed-frequency transmon with a Josephson junction which has the junction energy $E_{J,a}$. The charging energies of the coupler and the ancilla are given by $E_{C,c}$ and $E_{C,a}$, respectively. We assume vanishing offset charge for both the coupler and the ancilla, because they are in a “transmon regime” and insensitive to the details of the offset charge [52]. The storage, coupler, and the ancilla qubits are all

capacitively coupled with each other through their charge operators $i(\hat{a}^\dagger - \hat{a})$, \hat{N}_c , and \hat{N}_a , respectively.

When a coupler mediates tunable coupling between two qubits of the same type (e.g., as in Refs. [34,35]), the strengths of the qubit-coupler couplings are typically designed symmetrically for the two qubits neighboring the coupler. However, in our system, a coupler mediates tunable interaction between two different types of elements, i.e., a linear storage mode and a nonlinear ancilla transmon. Thus, we design the coupling strengths asymmetrically such that the coupler-storage coupling is much weaker than the coupler-ancilla coupling. Through this asymmetric coupling, we minimize the coupler-induced nonlinearity of the storage mode which is important for implementing a long-lived cat qubit in the storage mode. Note that, due to this design choice, the ancilla transmon is subject to strong coupler-induced nonlinearities. However, the impact of these induced nonlinearities on the ancilla is minimal given the ancilla’s large intrinsic nonlinearity.

The asymmetry between the coupler-storage coupling and the coupler-ancilla coupling in our system is clearly illustrated by the coupler-induced frequency shifts of the storage and the ancilla shown in Fig. 3(a). As the coupler tunes from its maximum (8.40 GHz) to minimum (5.77 GHz) frequency, the frequency of the storage is repelled down by only 3 MHz (from 5.349 to 5.346 GHz), whereas the frequency of the ancilla is pushed down by 141 MHz (from 5.211 to 5.070 GHz). Note that the storage frequency is higher than the ancilla frequency and closer to the coupler frequency. Thus, if the coupler-storage and the coupler-ancilla coupling strengths were chosen symmetrically, the storage is expected to hybridize more strongly with the coupler than the ancilla does with the coupler. However, thanks to the large asymmetry in the coupling strengths, the coupler is strongly hybridized only with the ancilla and is weakly hybridized with the storage despite such frequency arrangement.

Minimizing the induced nonlinearities of the storage mode is also crucial for ensuring the dispersive nature of the storage-ancilla interaction. As featured in the last term in Eq. (2), the desired dispersive interaction between the storage and the ancilla takes the form $\hat{H}_{sa} = \chi_{sa}\hat{a}^\dagger\hat{a}|e_a\rangle\langle e_a|$ in the dressed eigenbasis. In practice, the storage-ancilla interaction also contains ancilla-state-dependent self-Kerr terms of the storage, i.e.,

$$\begin{aligned} \hat{H}_{sa} = & \omega_s\hat{a}^\dagger\hat{a} + \omega_a|e_a\rangle\langle e_a| + \chi_{sa}\hat{a}^\dagger\hat{a}|e_a\rangle\langle e_a| \\ & + \frac{K_{s,a:|g\rangle}}{2}\hat{a}^{\dagger 2}\hat{a}^2|g_a\rangle\langle g_a| + \frac{K_{s,a:|e\rangle}}{2}\hat{a}^{\dagger 2}\hat{a}^2|e_a\rangle\langle e_a|. \end{aligned} \quad (\text{H3})$$

Here, $|g_a\rangle$ and $|e_a\rangle$ are the ground and the first excited states, respectively, of the ancilla. $K_{s,a:|g\rangle}$ and $K_{s,a:|e\rangle}$ are the storage self-Kerr when the ancilla is in the state $|g_a\rangle$ and $|e_a\rangle$,

respectively. For now, we assume that the coupler always stays in its ground state $|g_c\rangle$ and ignore its explicit presence (other than noting all the parameters in the effective Hamiltonian are dependent on the coupler flux $\varphi_{x,c}$).

Under the above effective Hamiltonian, the storage-mode frequencies are given by

$$\begin{aligned}\omega_{s,a:|g\rangle}^{(n\rightarrow n+1)} &= E(|n+1, g_a\rangle) - E(|n, g_a\rangle) \\ &= \omega_s + K_{s,a:|g\rangle}n, \\ \omega_{s,a:|e\rangle}^{(n\rightarrow n+1)} &= E(|n+1, e_a\rangle) - E(|n, e_a\rangle) \\ &= \omega_s + \chi_{sa} + K_{s,a:|e\rangle}n.\end{aligned}\quad (\text{H4})$$

In an ideal storage-ancilla dispersive coupling, the storage-mode frequency should be independent of the storage photon number n and is shifted only when the ancilla is excited from $|g_a\rangle$ to $|e_a\rangle$. However, due to the self-Kerr terms, the storage frequency is shifted not only by the ancilla excitation (by χ_{sa}), but also by its own excitations (by $K_{s,a:|g\rangle}n$ or $K_{s,a:|e\rangle}n$). These unwanted frequency shifts due to self-Kerr terms are particularly detrimental to a cat qubit with a large average photon number. Therefore, the self-Kerr nonlinearities $K_{s,a:|g\rangle}$ or $K_{s,a:|e\rangle}$ should be kept much smaller than the dispersive shift χ_{sa} .

In Figs. 9(a) and 9(b), we show the predicted storage-ancilla dispersive shift χ_{sa} and the storage self-Kerr nonlinearities $K_{s,a:|g\rangle}$ and $K_{s,a:|e\rangle}$ as a function of the coupler flux $\Phi_{x,c}/\Phi_0 = \varphi_{x,c}/2\pi$. The parameters of the model used in the prediction are tuned up to reproduce the experimentally measured mode frequencies ω_s , ω_c , and ω_a and the storage-ancilla dispersive shift χ_{sa} over the entire range of the coupler flux from $\Phi_{x,c}/\Phi_0 = 0$ to $\Phi_{x,c}/\Phi_0 = 0.5$ [shown in Figs. 3(a) and 3(b)]. At the coupler flux $\Phi_{x,c}/\Phi_0 = 0.38$ used for the parity measurement in our experiments, the model predicts $\chi_{sa}/2\pi = -0.88$ MHz (consistent with the experimental data) as well as $K_{s,a:|g\rangle}/2\pi = -1.7$ kHz and $K_{s,a:|e\rangle}/2\pi = 0.9$ kHz. Thus, in this case, $K_{s,a:|g\rangle}$ and $K_{s,a:|e\rangle}$ are approximately 3 orders of magnitude smaller than χ_{sa} , preserving the dispersive nature of the storage-ancilla interaction even when the storage has, e.g., ten photons on average, since $|K_s\alpha^2/\chi_{sa}| \ll 1$ for both $K_s = K_{s,a:|g\rangle}$ and $K_s = K_{s,a:|e\rangle}$ at $|\alpha|^2 = 10$.

At the coupler's minimum frequency position ($\Phi_{x,c}/\Phi_0 = 0.5$), the model predicts $\chi_{sa}/2\pi = -5.71$ MHz (to be compared against the experimental value of -5.73 ± 0.03 MHz) as well as $K_{s,a:|g\rangle}/2\pi = -332$ kHz and $K_{s,a:|e\rangle}/2\pi = -84$ kHz. In particular, when the ancilla is in the ground state $|g_a\rangle$, the storage self-Kerr $K_{s,a:|g\rangle}$ is only an order of magnitude smaller than the dispersive shift χ_{sa} . Thus, in this case, the dispersive nature of the storage-ancilla interaction breaks down in the higher excited state manifold of the storage (e.g., $|K_{s,a:|g\rangle}\alpha^2/\chi_{sa}|$ is already

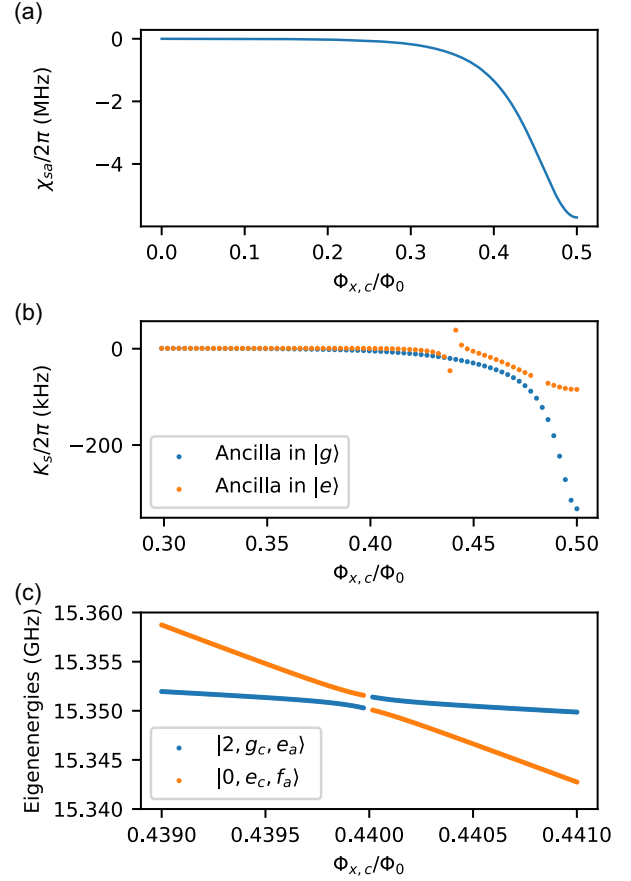


FIG. 9. Full circuit-quantization-level numerical model of a tunable dispersive interaction. (a) Storage-ancilla dispersive coupling χ_{sa} as a function of the external flux of the tunable coupler $\Phi_{x,c}$ in the unit of the flux quantum Φ_0 . (b) Self-Kerr of the storage mode K_s conditioned on the ancilla transmon being in $|g_a\rangle$ or $|e_a\rangle$ as a function of coupler flux. Note that $K_{s,a:|e\rangle}$ exhibits a resonance feature around $\Phi_{x,c}/\Phi_0 = 0.44$. (c) Avoided crossing between $|2, g_c, e_a\rangle$ and $|0, e_c, f_a\rangle$ indicating undesired conversion of two storage photons into one coupler excitation ($|g_c\rangle \rightarrow |e_c\rangle$) and one ancilla excitation ($|e_a\rangle \rightarrow |f_a\rangle$). This resonance is the root cause of the resonance feature in $K_{s,a:|e\rangle}$ observed above.

close to 1 with $|\alpha|^2 = 10$, whereas it should be much smaller than 1). This is why we avoided using the coupler's minimum frequency position as the on position in our experiments.

Besides the intolerably large K_s/χ_{sa} ratios, another reason for avoiding $\Phi_{x,c}/\Phi_0 = 0.5$ as the on position is the undesired resonances that arise as the coupler frequency is lowered and gets close to the storage and ancilla frequencies. For example, the $K_{s,a:|e\rangle}$ curve in Fig. 9(b) exhibits a resonance feature around $\Phi_{x,c}/\Phi_0 = 0.44$. The root cause of this feature is the avoided crossing between two energy levels $|2, g_c, e_a\rangle$ and $|0, e_c, f_a\rangle$ [see Fig. 9(c)]. This transition affects $K_{s,a:|e\rangle}$ because the eigenenergy of the state $|2, g_c, e_a\rangle$ is involved in the self-Kerr $K_{s,a:|e\rangle} \equiv E(|2, g_c, e_a\rangle) - 2E(|1, g_c, e_a\rangle) + E(|0, g_c, e_a\rangle)$.

In this $|2, g_c, e_a\rangle \leftrightarrow |0, e_c, f_a\rangle$ transition, which becomes resonant at around $\Phi_{x,c}/\Phi_0 = 0.44$, two storage photons are converted into one coupler excitation ($|g_c\rangle \rightarrow |e_c\rangle$) and one ancilla excitation ($|e_a\rangle \rightarrow |f_a\rangle$). As a result, if the on-position coupler flux is placed around or above $\Phi_{x,c}/\Phi_0 = 0.44$, this transition can be brought into resonance and cause undesired coupler and ancilla heating events. Thus, it is crucial to understand and avoid these resonances when realizing a tunable dispersive coupling between a linear mode and a nonlinear qubit.

2. ATS with serial inductances

Here, we provide more details of the buffer mode used for stabilizing a cat qubit. To achieve an optimal performance of the cat-qubit stabilization (e.g., in terms of bit-flip times), it is important to minimize the self-Kerr of the buffer mode such that the storage mode does not inherit undesired nonlinearities from the buffer mode (e.g., storage-buffer cross-Kerr and the storage self-Kerr). In this section, we show that inductances in series with the side junctions of the buffer can have significant impacts on the strength of the self-Kerr of the buffer mode. Moreover, we provide a design strategy to minimize the self-Kerr of the buffer mode on its saddle points by optimizing the Josephson energies of the side junctions.

As shown in Fig. 5, we consider a buffer containing a center inductor (implemented via a “shallow” Josephson junction array just with a few, e.g., three junctions in the array), two side junctions, and a shunt capacitor. Importantly, each side junction has a serial inductance. These serial inductances introduce extra flux nodes (with fluxes ϕ_{P_1} and ϕ_{P_2}) between the ground node (with zero flux by definition) and the buffer node (with flux ϕ). Thus, the inductive potential of the buffer mode is given by

$$\begin{aligned}
 V = & -NE_{J,\text{array}} \cos\left(\frac{\phi}{N}\right) \\
 & - E_{J,1} \cos(\phi + \varphi_{x,1} - \phi_{P_1}) + \frac{1}{2} E_{L,P_1} \phi_{P_1}^2 \\
 & - E_{J,2} \cos(\phi - \varphi_{x,2} - \phi_{P_2}) + \frac{1}{2} E_{L,P_2} \phi_{P_2}^2. \quad (\text{H5})
 \end{aligned}$$

Here, $E_{J,\text{array}}$ is the junction energy of each junction in the center junction array, N is the number of junctions in the junction array, $E_{J,1}$ and $E_{J,2}$ are the junction energies of the two side junctions, and E_{L,P_1} and E_{L,P_2} are the inductive energies of the serial inductances of the side junctions. Moreover, $\varphi_{x,1}$ and $\varphi_{x,2}$ are two external fluxes applied to the buffer mode.

In practice, the serial inductances L_{P_1} and L_{P_2} are much smaller than the inductances of the Josephson junctions (e.g., by many orders of magnitude), leading to $E_{L,P_1} \gg E_{J,1}$ and $E_{L,P_2} \gg E_{J,2}$. In this regime, one may attempt to eliminate the extra fluxes ϕ_{P_1} and ϕ_{P_2} by perturbatively optimizing them to minimize the potential [38,40]. Then, the inductive potential can be simplified into the form (up to an additive constant offset)

$$\begin{aligned}
 V \simeq & -NE_{J,\text{array}} \cos\left(\frac{\phi}{N}\right) \\
 & - E_{J,1} \cos(\phi + \varphi_{x,1}) + \frac{E_{J,1}^2}{4E_{L,P_1}} \cos[2(\phi + \varphi_{x,1})] \\
 & - E_{J,2} \cos(\phi - \varphi_{x,2}) + \frac{E_{J,2}^2}{4E_{L,P_2}} \cos[2(\phi - \varphi_{x,2})], \quad (\text{H6})
 \end{aligned}$$

involving only the buffer flux ϕ , where we use

$$\begin{aligned}
 & \min_{y \in \mathbb{R}} \left(A \cos(x - y) + \frac{1}{2} B y^2 \right) \\
 & = \min_{y \in \mathbb{R}} \left(A [\cos(x) \cos(y) + \sin(x) \sin(y)] + \frac{1}{2} B y^2 \right) \\
 & \simeq \min_{y \in \mathbb{R}} \left(A [\cos(x) + \sin(x)y] + \frac{1}{2} B y^2 \right) \\
 & = A \cos(x) - \frac{A^2}{2B} \sin^2(x) \\
 & = A \cos(x) + \frac{A^2}{4B} \cos(2x) + \text{const.} \quad (\text{H7})
 \end{aligned}$$

Alternatively, we can get the same relation via (see Ref. [40])

$$\begin{aligned}
 \min_{y \in \mathbb{R}} \left(A \cos(x - y) + \frac{1}{2} B y^2 \right) & = A \left(1 + \sum_{\nu \geq 1} \frac{2J_\nu(\frac{A}{B})}{\frac{A}{B} \nu^2} [\cos(\nu x) - 1] \right) \\
 & = A \left(1 + \frac{2J_1(\frac{A}{B})}{\frac{A}{B}} [\cos(x) - 1] + \frac{2J_2(2\frac{A}{B})}{4\frac{A}{B}} [\cos(2x) - 1] + \dots \right) \\
 & = A \left(1 + \frac{\frac{A}{B} + \mathcal{O}[(\frac{A}{B})^3]}{\frac{A}{B}} [\cos(x) - 1] + \frac{(\frac{A}{B})^2 + \mathcal{O}[(\frac{A}{B})^4]}{4\frac{A}{B}} [\cos(2x) - 1] + \dots \right) \\
 & \simeq A \cos(x) + \frac{A^2}{4B} \cos(2x) + \text{const.} \quad (\text{H8})
 \end{aligned}$$

To get an intuition on the impact of the serial inductances to the buffer mode, we make simplifying assumptions $E_{J,1} = E_{J,2} = E_J$ and $E_{L,P_1} = E_{L,P_2} = E_{L,P}$ (although the method should be more generally applicable). Then, we find

$$\begin{aligned} V &\simeq -NE_{J,\text{array}} \cos\left(\frac{\phi}{N}\right) \\ &\quad - E_J(\cos(\phi + \varphi_{x,1}) + \cos(\phi - \varphi_{x,2})) \\ &\quad + \frac{E_J^2}{4E_{L,P}}(\cos[2(\phi + \varphi_{x,1})] + \cos[2(\phi - \varphi_{x,2})]) \\ &= -NE_{J,\text{array}} \cos\left(\frac{\phi}{N}\right) - 2E_J \cos(\varphi_\Sigma) \cos(\phi + \varphi_\Delta) \\ &\quad + \frac{E_J^2}{2E_{L,P}} \cos(2\varphi_\Sigma) \cos[2(\phi + \varphi_\Delta)], \end{aligned} \quad (\text{H9})$$

where $\varphi_\Sigma \equiv (\varphi_{x,1} + \varphi_{x,2})/2$ and $\varphi_\Delta \equiv (\varphi_{x,1} - \varphi_{x,2})/2$. When φ_Σ is an odd integer multiple of $\pi/2$, we have $\cos(2\varphi_\Sigma) = -1$ and get

$$V \simeq -NE_{J,\text{array}} \cos\left(\frac{\phi}{N}\right) - \frac{E_J^2}{2E_{L,P}} \cos[2(\phi + \varphi_\Delta)]. \quad (\text{H10})$$

If the serial inductances is zero (i.e., $E_{L,P} \rightarrow \infty$), the potential energy is simply given by the one due to the junction array, i.e., $V = -NE_{J,\text{array}} \cos(\phi/N)$, independent of the delta flux φ_Δ . Then, in the limit of $N \rightarrow \infty$ (while keeping $E_{J,\text{array}} = NE_L$), the potential converges to a harmonic potential $V = \frac{1}{2}E_L\phi^2$ and the buffer mode's self-Kerr is vanishingly small at all values of φ_Δ along the cut defined by $\varphi_\Sigma = (2m+1)\pi/2$ for some integer $m \in \mathbb{Z}$. Moreover, the buffer mode will have a constant frequency at all values of φ_Δ .

If the serial inductance is not zero (i.e., $E_{L,P}$ is finite), the potential will always have a nontrivial anharmonic contribution $V' = -[E_J^2/(2E_{L,P})]\cos[2(\phi + \varphi_\Delta)]$. To understand the impact of this serial inductance potential term, we add a capacitive charging energy and construct the following Hamiltonian:

$$\begin{aligned} \hat{H} &= 4E_C(\hat{N} - n_g)^2 - NE_{J,\text{array}} \cos\left(\frac{\hat{\phi}}{N}\right) \\ &\quad - \frac{E_J^2}{2E_{L,P}} \cos[2(\hat{\phi} + \varphi_\Delta)]. \end{aligned} \quad (\text{H11})$$

Restricting to $n_g = 0$ and viewing

$$\hat{H}_0 = 4E_C\hat{N}^2 + \frac{1}{2}\left(\frac{E_{J,\text{array}}}{N}\right)\hat{\phi}^2 \quad (\text{H12})$$

as an unperturbed Hamiltonian, we find the correction to the energy eigenvalues to the first order in $\hat{H} - \hat{H}_0$. This yields

$$\begin{aligned} E_n &\simeq \left(\sqrt{8E_CE_L} - \frac{E_C}{N^2} + \frac{2E_J^2}{E_{L,P}} \sqrt{\frac{2E_C}{E_L}} \cos(2\varphi_\Delta) \right) n \\ &\quad - \frac{1}{2}E_C \left(\frac{1}{N^2} + \frac{8E_J^2}{E_{L,P}E_L} \cos(2\varphi_\Delta) \right) n(n-1), \end{aligned} \quad (\text{H13})$$

where we define $E_L \equiv E_{J,\text{array}}/N$. Thus, the frequency and the self-Kerr of the buffer mode are perturbatively given by

$$\begin{aligned} \omega_b &\simeq \sqrt{8E_CE_L} - \frac{E_C}{N^2} + \frac{2E_J^2}{E_{L,P}} \sqrt{\frac{2E_C}{E_L}} \cos(2\varphi_\Delta), \\ K_b &\simeq -\frac{E_C}{N^2} - \frac{8E_CE_J^2}{E_{L,P}E_L} \cos(2\varphi_\Delta). \end{aligned} \quad (\text{H14})$$

To benchmark the validity of the analytic expressions in Eq. (H14), we perform numerical circuit quantization of the buffer mode in the presence of serial inductors. In the numerics, we do not eliminate the extra modes due to the serial inductors. To make the numerics feasible, it is important to include some small nonzero capacitance in parallel with each serial inductor such that the frequency of the serial inductor mode does not diverge. In practice, we set this capacitance such that the frequency of the serial inductor mode is much higher than that of the buffer mode and confirm that different choices of the capacitance values in this regime do not change the key properties of the buffer mode. As shown in Fig. 10, the analytic expressions in Eq. (H14) agree qualitatively with the exact results from numerical circuit quantization.

Note that, on the saddle points (i.e., φ_Δ is an odd integer multiple of $\pi/2$ in addition to φ_Σ being also an odd integer multiple of $\pi/2$), we have $\cos(2\varphi_\Delta) = -1$, and, thus, the self-Kerr of the buffer mode is perturbatively given by

$$K_b \simeq -\frac{E_C}{N^2} + \frac{8E_CE_J^2}{E_{L,P}E_L}. \quad (\text{H15})$$

Here, the first term is due to the contribution from the junction array, and the second term is due to the contribution from the serial inductance. That is, the junction array introduces negative self-Kerr of the buffer mode and the serial inductance introduces positive self-Kerr of the buffer mode at the saddle points. Thus, by balancing the contributions from these two terms, i.e.,

$$\frac{E_C}{N^2} = \frac{8E_CE_J^2}{E_{L,P}E_L}, \quad (\text{H16})$$

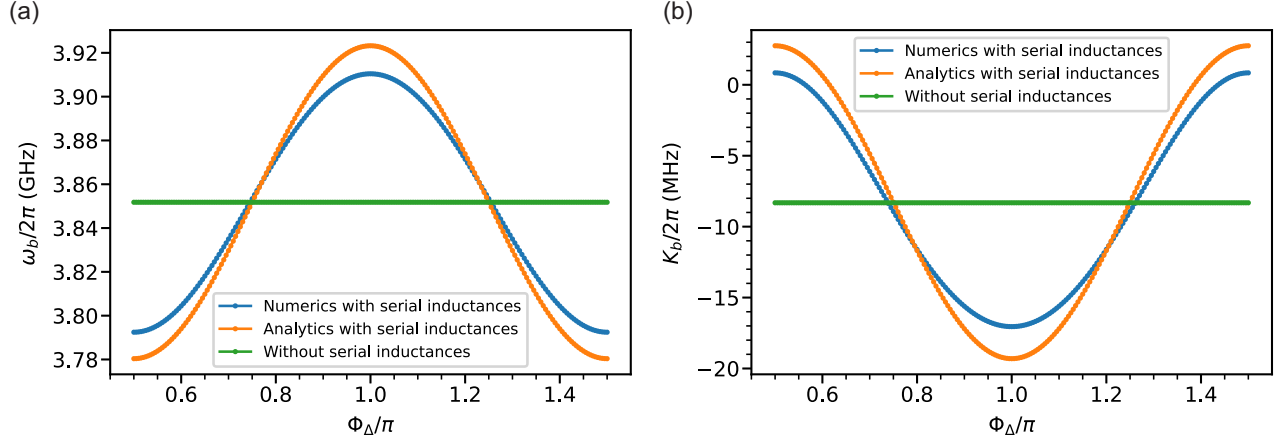


FIG. 10. Numerical model of a buffer. (a) Frequency and (b) self-Kerr (i.e., anharmonicity) of the buffer mode in the presence (blue and orange curves) or absence (green curves) of the inductances in series with the side junctions. The blue curves represent predictions via exact numerical circuit quantization, and the orange curves represent the analytic expressions in Eq. (H14). Note that the two side junction energies were assumed to be identical such that the analytic expressions are applicable, although the numerical approach can be more generally applied.

we can minimize the self-Kerr of the buffer mode even when N is finite (e.g., $N = 3$ as opposed to $N \rightarrow \infty$). In practice, the side junction energy E_J is one of the more flexible design knobs (e.g., not affecting the buffer frequency) among the parameters involved. Hence, choosing it to be $E_J = \sqrt{E_{J,\text{array}}E_{L,P}/8N^3}$ subject to given values of $E_{J,\text{array}}$ and $E_{L,P}$ is one way to satisfy the above condition for minimizing the buffer self-Kerr.

3. Buffer-pump-induced parity flip of the storage

To realize the desired 3WM interaction ($g_2\hat{a}^2\hat{b}^\dagger + \text{H.c.}$), we pump the sigma flux of the buffer mode [16]. Here, we discuss how this flux pump can introduce undesired loss mechanism of the storage mode due to driven resonances. In particular, we make use of the Floquet theory in the form of the expanded Hilbert space formalism [53,54]. There are simpler alternatives such as naively tabulating all possible frequency-collision conditions. However, with this simple approach, one may either accidentally neglect relevant higher-order transitions or become overly conservative by accounting for irrelevant resonance conditions that are prohibited by the “selection rules.” In contrast, the Floquet approach can serve as a systematic tool for capturing arbitrary higher-order processes as well as estimating the width of the resonance features.

We are ultimately interested in a composite system consisting of a storage, a flux-pumped buffer, and the additional modes associated with the buffer serial inductances. For now, we focus on an isolated flux-pumped ATS to simplify the presentation. In this case, the Hamiltonian of the buffer is given by

$$\begin{aligned} \hat{H}_b(t) = & 4E_C\hat{N}_b^2 - NE_{J,\text{array}}\cos\left(\frac{\hat{\phi}_b}{N}\right) \\ & - (E_{J,1} + E_{J,2})\cos[\varphi_\Sigma(t)]\cos(\hat{\phi}_b + \varphi_\Delta) \\ & + (E_{J,1} - E_{J,2})\sin[\varphi_\Sigma(t)]\sin(\hat{\phi}_b + \varphi_\Delta), \end{aligned} \quad (\text{H17})$$

where $\varphi_\Sigma(t)$ is the pumped sigma flux of the buffer and φ_Δ is the static delta flux of the buffer. We further assume that the buffer is statically parked at one of its saddle points, e.g., $(\varphi_\Sigma, \varphi_\Delta) = (\pi/2, \pi/2)$. Then, by adding a sigma flux pump with amplitude ϵ_p and frequency ω_p , we have $\varphi_\Sigma(t) = (\pi/2) + \epsilon_p \cos(\omega_p t)$ and $\varphi_\Delta = \pi/2$ yielding the buffer Hamiltonian

$$\begin{aligned} \hat{H}_b(t) = & 4E_C\hat{N}_b^2 - NE_{J,\text{array}}\cos\left(\frac{\hat{\phi}_b}{N}\right) \\ & - (E_{J,1} + E_{J,2})\sin[\epsilon_p \cos(\omega_p t)]\sin\hat{\phi}_b \\ & + (E_{J,1} - E_{J,2})\cos[\epsilon_p \cos(\omega_p t)]\cos\hat{\phi}_b. \end{aligned} \quad (\text{H18})$$

Note that the driven buffer Hamiltonian above is time dependent and is periodic in time t with a period of $2\pi/\omega_p$. Using the expanded Hilbert space formalism [53,54], we can construct a time-independent Hamiltonian which has the equivalent Floquet spectrum as the time-dependent periodic Hamiltonian. In particular, we introduce a pump mode with a charge operator \hat{N}_p and a cosine-phase operator $\cos\hat{\phi}_p$, which are defined as

$$\begin{aligned} \hat{N}_p = & \sum_{n_p=-\infty}^{\infty} n_p |n_p\rangle\langle n_p|, \\ \cos\hat{\phi}_p = & \frac{1}{2} \sum_{n_p=-\infty}^{\infty} (|n_p + 1\rangle\langle n_p| + \text{H.c.}). \end{aligned} \quad (\text{H19})$$

Here, one may interpret n_p as the number of pump photons in the pump mode. Then, the equivalent time-independent Floquet Hamiltonian is given by

$$\hat{H}_F = \hat{H}_b + \hat{H}_p + \hat{H}_{bp}, \quad (\text{H20})$$

where $\hat{H}_b = 4E_C \hat{N}_b^2 - NE_{J,\text{array}} \cos(\hat{\phi}_b/N) + (E_{J,1} - E_{J,2}) \times \cos \hat{\phi}_b$ is the static buffer Hamiltonian, $\hat{H}_p = \omega_p \hat{N}_p$ is the Hamiltonian of the pump mode, and

$$\begin{aligned} \hat{H}_{bp} = & -(E_{J,1} + E_{J,2}) \sin(\epsilon_p \cos \hat{\phi}_p) \sin \hat{\phi}_b \\ & - (E_{J,1} - E_{J,2}) [1 - \cos(\epsilon_p \cos \hat{\phi}_p)] \cos \hat{\phi}_b \end{aligned} \quad (\text{H21})$$

is the interaction Hamiltonian between the buffer mode and the pump mode due to the sigma flux pump.

From the pump-mode Hamiltonian $\hat{H}_p = \omega_p \hat{N}_p$, we see that the pump mode releases an excess energy of ω_p (i.e., pump frequency) when it loses one pump photon. Such an excess energy is gained by system through the interaction Hamiltonian \hat{H}_{bp} to either realize a desired driven process or induce an undesired transition. For illustration, we expand the interaction Hamiltonian to the first order in the pump amplitude ϵ_p , i.e.,

$$\hat{H}_{bp} = -(E_{J,1} + E_{J,2}) \epsilon_p \cos \hat{\phi}_p \sin \hat{\phi}_b + \mathcal{O}(\epsilon_p^2). \quad (\text{H22})$$

The term $\cos \hat{\phi}_p$ acts on the pump mode and either adds or takes away one pump photon in the pump mode [see Eq. (H19)]. The term $\sin \hat{\phi}_b$ acts on the bare buffer mode. For an isolated buffer mode, this term can induce various transitions in the buffer mode which involve an odd number of excitation gain or loss (e.g., $|0\rangle \leftrightarrow |1\rangle$ and $|0\rangle \leftrightarrow |3\rangle$) due to the parity. When the resonant frequency of one of these transitions, say, $|0\rangle \leftrightarrow |3\rangle$, coincides with the frequency of the pump mode, i.e., $\omega_p = \omega_{b,0 \leftrightarrow 3}$, an avoided crossing happens in the expanded Hilbert space between the two relevant states $|0, n_p = 0\rangle \leftrightarrow |3, n_p = -1\rangle$. This avoided crossing is then translated into a time-dependent $|0\rangle \leftrightarrow |3\rangle$ transition in the original Hilbert space of the system after the pump mode is projected out [54].

In our system, the buffer mode is not isolated and is instead coupled to the storage mode as well as the additional modes associated with the serial inductances. Even then, most of the above analysis is still applicable with the only difference that the term $\sin \hat{\phi}_b$ can now introduce more complex transitions involving the modes that are hybridized with the buffer mode. The specific transition most relevant to the cat-qubit implementation is a 3WM process $\hat{a}^2 \hat{b}^\dagger + \text{H.c.}$ on the dressed modes of the storage and the buffer. In particular, this transition is realized by the matrix elements of the form $\langle n-2, 1 | \sin \hat{\phi}_b | n, 0 \rangle$ and $\langle n, 0 | \sin \hat{\phi}_b | n-2, 1 \rangle$, where $|n, m\rangle$ is a static dressed

eigenstate containing n storage photons and m buffer photons in the absence of the pump (here, we assume that the additional modes associated with the serial inductances always stay in their ground states given the high frequency of these modes well above 10 GHz). Let $\bar{\omega}_s$ and $\bar{\omega}_b$ denote the Stark-shifted frequencies of the storage and the buffer, respectively, under the buffer pump. Then, when the resonance condition $\omega_p = 2\bar{\omega}_s - \bar{\omega}_b$ is met, an avoided crossing between the two Floquet dressed states $|n, 0, n_p = -1\rangle$ and $|n-2, 1, n_p = 0\rangle$ occurs in the expanded Hilbert space, leading to the desired 3WM process in the system when the pump mode is projected out.

The above expanded Hilbert space framework allows us to systematically search for any undesired resonances under the buffer pump. That is, one can diagonalize the time-independent Floquet Hamiltonian \hat{H}_F , including the pump mode, and then look for undesired avoided crossings between a pair of Floquet dressed states. Since the buffer mode is heavily attenuated with a large decay rate κ_b , we primarily focus on the spurious transitions involving the buffer ground state. In particular, we compute the Stark-shifted frequency of the buffer mode [i.e., $\bar{\omega}_b \equiv E(|0, 1, n_p = 0\rangle) - E(|0, 0, n_p = 0\rangle)$] which is shown in Fig. 11(a) for various pump amplitudes of $\epsilon_p \in [0, 0.02, 0.04, 0.06, 0.08, 0.1]$. The underlying parameters of the model Hamiltonian are tuned up to match the experimentally observed energy spectrum [e.g., as in Fig. 5(b)].

As the pump amplitude ϵ_p is increased from 0 (dark purple) to 0.1 (light purple) in Fig. 11(a), the frequency of the buffer mode is negatively Stark shifted. Note that there are kinks in the buffer frequencies around both $\omega_p/2\pi \sim 7.7$ GHz and $\omega_p/2\pi \sim 7.2$ GHz. We identify that the former feature around $\omega_p/2\pi \sim 7.7$ GHz is due to the desired 3WM avoided crossing between $|2, 0, n_p = -1\rangle$ and $|0, 1, n_p = 0\rangle$ (i.e., $\hat{a}^2 \hat{b}^\dagger + \text{H.c.}$) with the resonance condition of $\omega_p = 2\bar{\omega}_s - \bar{\omega}_b$. On the other hand, the latter feature around $\omega_p/2\pi \sim 7.2$ GHz is due to an undesired avoided crossing between $|1, 3, n_p = -2\rangle$ and $|0, 0, n_p = 0\rangle$ (i.e., $\hat{a}^\dagger \hat{b}^{\dagger 3} + \text{H.c.}$) with the resonance condition of $2\omega_p = \bar{\omega}_s + 3\bar{\omega}_b$. Upon adiabatic elimination of the buffer mode, this process leads to a pump-induced heating of the storage mode. Note that this undesired resonance is a second-order process, involving two pump photons, that could have been overlooked if one looked only for first-order processes manually without the help of a systematic Floquet analysis.

Figures 11(b) and 11(c) show various experimental signatures of the undesired resonant process $\hat{a}^\dagger \hat{b}^{\dagger 3} + \text{H.c.}$ observed in our device. In Fig. 11(b), we show the measured buffer frequency versus buffer-pump frequency, similar to that simulated in Fig. 11(a). The variation in the

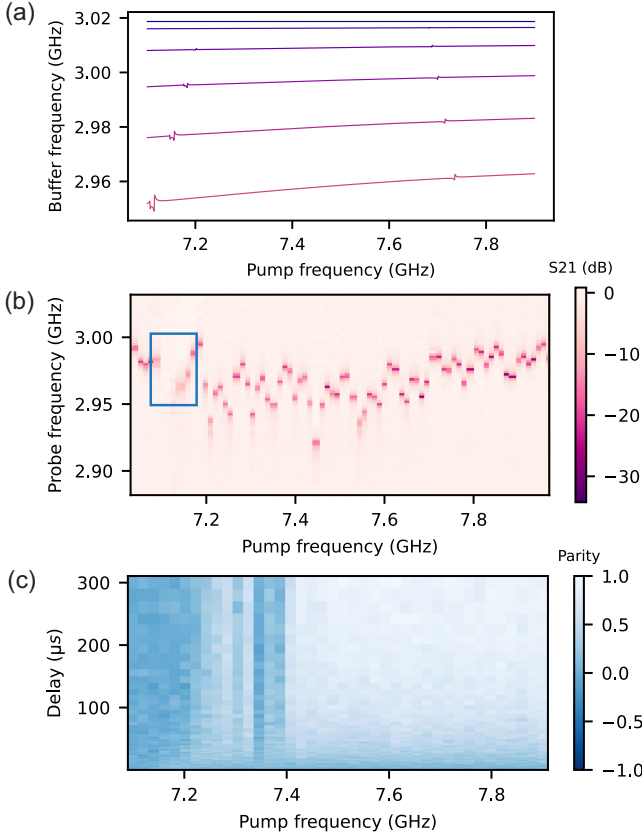


FIG. 11. Pump-induced parity changes. (a) Numerically simulated buffer frequency as a function of the pump frequency for different pump amplitudes $\epsilon_p \in [0, 0.02, 0.04, 0.06, 0.08, 0.1]$. As the pump amplitudes increase, the buffer frequency is shifted lower due to the buffer Stark shift. In addition to the desired 3WM condition near 7.7 GHz, we observe an additional resonance near 7.2 GHz corresponding to $(3\omega_b + \omega_a)/2$. (b) Experimentally measured spectrum of the buffer mode resonance as a function of pump frequency for $\epsilon_p \sim 0.08-0.1$. There is a notable loss in contrast in the buffer spectrum in a frequency band below 7.2 GHz, as indicated by the blue rectangle. Note the 3WM condition (expected at approximately 7.7 GHz) is not observable on a scan as coarse as this. (c) Parity of the storage mode as a function of time starting in a superposition of $|0\rangle/|1\rangle$ Fock states prepared with $3 \mu\text{s}$ of two-photon dissipation. The storage mode is expected to decay to the vacuum state with parity of $+1$, but we observe multiple frequency bands where the steady-state parity of the storage mode is far from $+1$.

buffer mode frequency versus the buffer-pump frequency is due to variations in flux pump power delivery with frequency which makes direct observation of resonances difficult. Notably, though, below approximately $\omega_p/2\pi \sim 7.2$ GHz, we observe a reduction in the resonant buffer transmission contrast. In Fig. 11(c), we show the corresponding measured photon number parity of the storage mode as a function of time after preparation in the $|0\rangle/|1\rangle$ Fock state, under the same buffer pump conditions. In the absence of a nonlinear resonance or for the desired 3WM process, the storage mode would decay into the ground

state of parity $+1$. Here, we see the impact of the undesired resonance in a more pronounced manner, where in a broad frequency band around $\omega_p/2\pi \sim 7.2$ GHz, the storage-mode parity does not decay to $+1$ but rather reaches a steady-state parity far from $+1$. This is consistent with the action of buffer-pump-induced heating due to the undesired $\hat{a}^\dagger \hat{b}^{\dagger 3} + \text{H.c.}$ resonant process.

If the two resonant pump frequencies $\omega_p = 2\bar{\omega}_s - \bar{\omega}_b$ (desired 3WM) and $\omega_p = (\bar{\omega}_s + 3\bar{\omega}_b)/2$ (undesired resonance) are close to each other, the phase coherence of the storage mode at the desired 3WM condition is degraded by the near-resonant pump-induced storage heating. In our system, we have carefully chosen the device parameters to ensure that the desired 3WM condition is far detuned from the undesired resonance condition by about 500 MHz in pump frequency. Thus, the storage-mode coherence times in our device are not significantly degraded by the two-photon dissipation as demonstrated in Figs. 4(d), 4(e), and 6(b).

We lastly remark that there is another resonance feature in Fig. 11(c) around $\omega_p/2\pi \sim 7.35$ GHz (i.e., degradation in the storage parity) that is not captured by our model prediction in Fig. 11(a). In this regard, we note that our model does not account for the decay of the buffer mode. If the buffer loss is accounted for in the model, we expect that each resonance feature will be broadened in the pump frequency due to the decay-induced linewidth of the buffer mode. Specifically, in our system, the buffer loss is engineered with a metamaterial bandpass filter with a wide passband and multiple filter modes. Thus, we anticipate that the pump-induced resonance features in our system are not only broadened by the loss channel, but also come with additional ripples associated with internal modes of the multipole filter. The latter could be a root cause of the unexplained parity degradation at the pump frequency of $\omega_p/2\pi \sim 7.35$ GHz. We leave a more detailed analysis of this phenomena as a future research direction for engineering driven-dissipative systems.

4. Analysis of cat-qubit bit-flip times with an effective model

In our cat-qubit experiments, the coupler is parked in the off position for the bulk of a control sequence and is pulsed to the on position only in the final measurement step (or sometimes also in the state-preparation step). Thus, to understand the bit-flip times, which are not sensitive to the state-preparation and measurement errors, it suffices to consider an effective Hamiltonian with the coupler in the off position. Hence, we consider the effective Hamiltonian in Eq. (3), which is duplicated here for the readers' convenience:

$$\hat{H} = \left(g_2(\hat{a}^2 - \alpha^2)\hat{b}^\dagger + \text{H.c.} \right) + \chi_{sb}\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} + \frac{K_s}{2}\hat{a}^{\dagger 2}\hat{a}^2. \quad (\text{H23})$$

In addition to this effective Hamiltonian, we further consider a Lindblad dissipator of the form

$$\mathcal{L}[\hat{\rho}(t)] = \kappa_b \mathcal{D}[\hat{b}]\hat{\rho}(t) + \kappa_1 \mathcal{D}[\hat{a}]\hat{\rho}(t) + \kappa_\phi \mathcal{D}[\hat{a}^\dagger \hat{a}]\hat{\rho}(t). \quad (\text{H24})$$

To understand how the storage-buffer cross-Kerr and the storage self-Kerr degrade the bit-flip performance of a cat qubit, we for now ignore the storage decay and dephasing times, i.e., $\kappa_1 = \kappa_\phi = 0$ (in the numerical simulations below, we consider nonzero decay and dephasing rates). Then, we apply the effective operator formalism [55] to understand the system analytically. We consider the subspace with the buffer mode in its ground state as the ground-state manifold and the rest as the excited-state manifold. Moreover, we ignore the second and higher excited states of the buffer mode in our analytic calculations.

Adopting the notation in Ref. [55], we identify \hat{L} and \hat{H}_e as $\hat{L} = \sqrt{\kappa_b}$ and $\hat{H}_e = \chi_{sb} \hat{a}^\dagger \hat{a}$ and find that the non-Hermitian Hamiltonian of the excited-state manifold is given by $\hat{H}_{\text{NH}} = \hat{H}_e - (i/2)\hat{L}^\dagger \hat{L} = \chi_{sb} \hat{a}^\dagger \hat{a} - i\kappa_b/2$. Then, in the limit of $\chi_{sb} \langle \hat{a}^\dagger \hat{a} \rangle / \kappa_b \ll 1$, we have

$$\hat{H}_{\text{NH}}^{-1} = \frac{2i}{\kappa_b} \left(1 - \frac{2i\chi_{sb}}{\kappa_b} \hat{a}^\dagger \hat{a} + \mathcal{O}\left(\left(\frac{\chi_{sb}}{\kappa_b} \hat{a}^\dagger \hat{a}\right)^2\right) \right). \quad (\text{H25})$$

We further identify \hat{V}_+ and \hat{H}_g as $\hat{V}_+ = g_2(\hat{a}^2 - \alpha^2) = \hat{V}_+^\dagger$ and $\hat{H}_g = (K_s/2)\hat{a}^{\dagger 2}\hat{a}^2$, respectively, and arrive at the effective master equation

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}_{\text{eff}}, \hat{\rho}] + \frac{4g_2^2}{\kappa_b} \mathcal{D}\left[\left(1 - \frac{2i\chi_{sb}}{\kappa_b} \hat{a}^\dagger \hat{a}\right)(\hat{a}^2 - \alpha^2)\right]\hat{\rho}, \quad (\text{H26})$$

with

$$\hat{H}_{\text{eff}} = -\frac{2g_2^2\chi_{sb}}{\kappa_b^2}(\hat{a}^{\dagger 2} - \alpha^2)\hat{a}^\dagger \hat{a}(\hat{a}^2 - \alpha^2) + \frac{K_s}{2}\hat{a}^{\dagger 2}\hat{a}^2, \quad (\text{H27})$$

by using $\hat{L}_{\text{eff}} = \hat{L}\hat{H}_{\text{NH}}^{-1}\hat{V}_+$ and $\hat{H}_{\text{eff}} = -\hat{V}_-[\hat{H}_{\text{NH}}^{-1} + (\hat{H}_{\text{NH}}^{-1})^\dagger]\hat{V}_+/2 + \hat{H}_g$.

To get an intuition on this effective master equation, we first consider a special case where the storage-buffer cross-Kerr is absent, i.e., $\chi_{sb} = 0$. In this case, we have two-photon dissipation $d\hat{\rho}/dt = -i[\hat{H}_{\text{eff}}, \hat{\rho}] + \kappa_2 \mathcal{D}[\hat{a}^2 - \alpha^2]\hat{\rho}$ with an added storage Kerr nonlinearity $\hat{H}_{\text{eff}} = (K_s/2)\hat{a}^{\dagger 2}\hat{a}^2$. The addition of the storage-mode Kerr nonlinearity ($K_s \ll \kappa_2$) to the dissipatively stabilized cat qubit has a relatively harmless effect in the case without the storage-buffer cross-Kerr. In particular, the two-photon drive can simply be reapportioned between dissipative and Hamiltonian terms to yield a hybrid dissipative-Kerr cat stabilization [56], i.e.,

$$\begin{aligned} \frac{d\hat{\rho}}{dt} = & -i\left[\frac{K_s}{2}(\hat{a}^{\dagger 2} - \alpha^{*2})(\hat{a}^2 - \alpha^2), \hat{\rho}\right] \\ & + \frac{4g_2^2}{\kappa_b} \mathcal{D}[\hat{a}^2 - \alpha^2]\hat{\rho}, \end{aligned} \quad (\text{H28})$$

where the steady-state coherent-state amplitude squared is now $\alpha'^2 = \alpha^2/(1 + iK_s/\kappa_2)$ with $\kappa_2 = 4g_2^2/\kappa_b$. This indicates that the storage Kerr nonlinearity is simply used for a Kerr-cat-type stabilization making use of some of the two-photon drive from the dissipative stabilization. This causes a slight change in amplitude and the phase of the stabilized cat qubit but does not degrade its bit-flip times. Thus, the storage self-Kerr alone without the storage-buffer cross-Kerr is not detrimental to a dissipative cat qubit.

As discussed in detail in the main text, a nonzero storage-buffer cross-Kerr causes dephasing of the storage mode when the buffer is excited. For example, such dephasing manifests as a distorted engineered jump operator $(1 - 2i\chi_{sb}\hat{a}^\dagger \hat{a}/\kappa_b)(\hat{a}^2 - \alpha^2)$. This additional dephasing can then significantly degrade the bit-flip times of a cat qubit when combined with various other mechanisms such as storage loss, dephasing, and self-Kerr which bring the storage mode outside of the cat-qubit manifold and, consequently, excite the buffer mode.

Note that the effective master equation in Eqs. (H26) and (H27) is derived only up to the first order in $\chi_{sb}\hat{a}^\dagger \hat{a}/\kappa_b$ using a series-expanded expression for \hat{H}_{NH}^{-1} . While these simple expressions give a useful intuition on how the storage-buffer cross-Kerr degrades the bit-flip times of a cat qubit, as shown in Fig. 12, the resulting effective dynamics (green curve) does not agree quantitatively with the exact dynamics given by the full storage-buffer model (blue curve), especially when $|\chi_{sb}\alpha^2|$ is comparable to κ_b . However, one can drastically improve the accuracy of the effective master equation by not using a series expansion for \hat{H}_{NH}^{-1} and instead using an exact expression (orange curve):

$$\hat{H}_{\text{NH}}^{-1} = \sum_{n=0}^{\infty} \frac{2i}{\kappa_b + 2i\chi_{sb}n} |n\rangle\langle n|. \quad (\text{H29})$$

The computational cost of solving the effective master equation does not increase with this modification, while the accuracy is substantially improved.

The numerical simulations in Fig. 12 clearly illustrate the damaging effects of the storage-buffer cross-Kerr. In our numerical simulations, we use the measured parameters when applicable and otherwise use the parameters predicted from circuit-quantization-level models. For the buffer decay rate, we use $\kappa_b/2\pi = 10.7$ MHz. For κ_1 and κ_ϕ , we use the relations

$$\frac{1}{T_1} = \kappa_1, \quad \frac{1}{T_2} = \frac{1}{2}\kappa_1 + \frac{1}{2}\kappa_\phi \quad (\text{H30})$$




FIG. 12. Comparison of effective models for two-photon dissipation with added Kerr nonlinearities. We simulate the case where E_J and $E_{L,P}$ are double the experimental values [as in the brown curve in Fig. 6(a)]. In this regime, the effect of the storage-buffer cross-Kerr nonlinearities is significant. We compare a full model which explicitly includes both the storage mode and the buffer mode (blue), an effective model with the buffer adiabatically eliminated with an exact expression of H_{NH}^{-1} (orange), and an effective model with the buffer adiabatically eliminated which uses only the leading-order term in H_{NH}^{-1} (green).

and estimate κ_1 and κ_ϕ based on the measured storage-mode T_1 and T_2 coherence times under the two-photon dissipation (i.e., $T_1 = 79 \mu\text{s}$ and $T_2 = 116 \mu\text{s}$).

For the storage self-Kerr K_s and the storage-buffer cross-Kerr χ_{sb} , we use the values predicted by a detailed modeling of our storage-buffer system, accounting for the presence of serial inductances in the buffer mode (see Sec. V). In particular, when we use the circuit parameters that are tuned up to reproduce the experimentally measured buffer frequency spectrum, we predict $K_s/2\pi = 1.1 \text{ kHz}$ and $\chi_{sb}/2\pi = 156 \text{ kHz}$. These K_s and χ_{sb} values are used in the dashed turquoise curve in Fig. 6(a). When we use the hypothetical circuit parameters where the average buffer side junction energy E_J and the serial inductance L_P are doubled compared to the actual device parameters, we predict $K_s/2\pi = 9.8 \text{ kHz}$ and $\chi_{sb}/2\pi = 2029 \text{ kHz}$, which are used in the dashed brown curve in Fig. 6(a).

Since the cat-qubit bit-flip times are not sensitive to state preparation and measurement errors, we do not explicitly simulate the realistic readout of the storage mode using an ancilla transmon. Instead, we assume ideal state preparation and measurement of a cat qubit in the $|0\rangle/|1\rangle$ basis (i.e., Z basis). Then, we extract the bit-flip time using the same approach used in the experiments, including the convention of the bit-flip time as defined in Appendix F.

From Fig. 6(a), it is worth noting that experimentally measured bit-flip times exceed the simulated bit-flip times (dashed turquoise curve) by a multiplicative offset in the regime where the bit-flip error rates are exponentially suppressed in $|\alpha|^2$. One of the possible reasons behind this slight discrepancy is that we consider white-noise pure

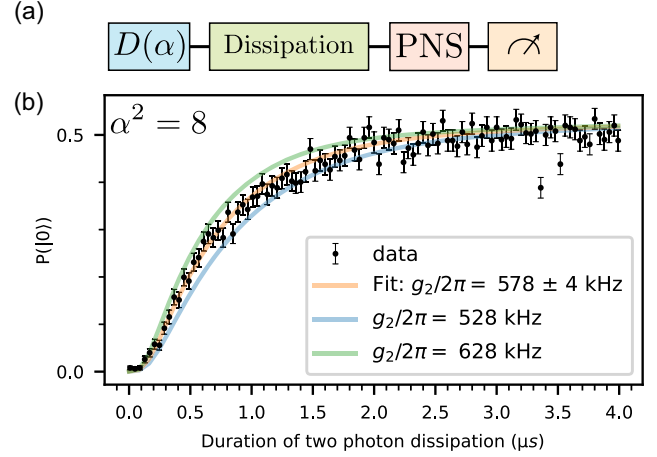


FIG. 13. g_2 fitting. (a) Pulse sequence for the experiment. We displace the storage mode into a coherent state. We apply a pure two-photon dissipation without the buffer drive to map the storage mode into the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ manifold. We measure the storage-mode vacuum population using a number-selective π pulse on the ancilla. (b) Experimental data of the vacuum population of the storage mode as a function of the two-photon dissipation duration starting from $|\alpha|^2 = 8$. Curves correspond to simulations of the experiment with an effective model containing the storage mode and buffer mode. The model fit (orange) yields $g_2/2\pi = 578 \pm 4 \text{ kHz}$, from which we infer $\kappa_2/2\pi = 124 \pm 2 \text{ kHz}$ using the independently measured value of $\kappa_b/2\pi = 10.7 \text{ MHz}$. Reference curves (blue and green) are also included to show the sensitivity of the convergence.

dephasing of the storage mode in our effective model above. On the other hand, the pure storage-mode dephasing in our device may have a colored-noise spectrum (e.g., more concentrated on the lower frequencies) as opposed to a frequency-independent white-noise spectrum. However, one can phenomenologically tune the white-noise dephasing rate κ_ϕ to attempt to explain the observed cat-qubit bit-flip times. In the dashed blue curve in Fig. 6(a) and in all pulsed cat-qubit stabilization simulations in Fig. 7(b), we have indeed empirically reduced κ_ϕ by a factor of 2 such that the simulations agree better with the static bit-flip times measured experimentally. However, we have otherwise kept all other parameters the same.

APPENDIX I: g_2 AND κ_2 FITTING

To determine g_2 and κ_2 , we fit the decay of the storage mode to the vacuum state from a coherent state when a pure two-photon dissipation $\kappa_2 \mathcal{D}[\hat{a}^2]$ is applied. The pure two-photon dissipation maps the storage mode to the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ manifold. Since we start from an initial state with vanishing average parity, this results in the vacuum population of the storage mode quickly approaching approximately 0.5 where the convergence timescale is determined by the two-photon dissipation strength. We fit this vacuum population curve to determine g_2 and κ_2 .

The sequence used for the experiment is shown in Fig. 13(a). The storage mode is displaced to a coherent state $|\alpha\rangle$ with $|\alpha|^2 = 8$, two-photon dissipation is applied for variable durations, and finally the vacuum population of the storage mode is read out using a photon-number-selective pulse on the ancilla. Figure 13(b) shows the vacuum population of the storage mode as a function of the two-photon dissipation duration. At the beginning, the vacuum population of the storage mode is negligible, because the initial coherent state has low overlap with it. On the timescale of $3\ \mu\text{s}$, the storage mode has mostly converged to the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ manifold as indicated by the vacuum population reaching approximately 0.5. The rest of the population in $|\hat{n} = 1\rangle$ will decay to the vacuum state due to the storage T_1 . By fitting the experimental data against a Lindblad master equation model, we extract $g_2/2\pi = (\kappa_2\kappa_b)^{1/2}/4\pi = 578 \pm 4\ \text{kHz}$. From this value of $g_2/2\pi$ and the independently measured value of $\kappa_b/2\pi = 10.7\ \text{MHz}$, we infer a value of $\kappa_2/2\pi = 124 \pm 2\ \text{kHz}$ ($\kappa_2 \approx 4g_2^2/\kappa_b$ when the buffer mode is adiabatically eliminated).

APPENDIX J: STORAGE CONDITIONAL PHASE MEASUREMENT

We use a storage conditional phase measurement to determine the storage-ancilla dispersive coupling and the flux pulse length needed for parity measurement. The pulse sequence for the conditional phase measurement is shown in Fig. 14(a). The conditional phase measurement involves preparing the storage into a small coherent state, preparing the ancilla state ($|i\rangle$, where $i = g, e$), applying a flux pulse with a variable pulse length t , unpreparing the ancilla state, displacing the storage mode with a variable phase (eight phases ϕ from 0 to $7\pi/8$ in increments of $\pi/8$), and finally reading out the storage-mode vacuum population. Both the initial and final displacement use an amplitude of $\alpha = 0.5$. The raw data resulting from this experiment are shown in Fig. 14(b). For each ancilla state, we observe fringes, indicating the evolving phase of the storage mode.

During the course of the flux pulse, the initial storage-mode coherent state $|\alpha\rangle$ approximately evolves into a final coherent state with the same amplitude but a different phase $|\alpha e^{i\phi_0(t,i)}\rangle$. We aim to extract $\phi_0(t, i)$ from the fringes of the conditional phase measurement for each ancilla state and flux pulse length. We determine $\phi_0(t, i)$ by fitting the measured vacuum population to the model:

$$P_0 = A e^{-\alpha^2[2-2\cos(2\pi\phi-2\pi\phi_0)]} + C. \quad (\text{J1})$$

In Fig. 14(c), we plot the difference in storage-mode phase between the ancilla being in $|g\rangle$ versus $|e\rangle$ [i.e., $\phi_0(t, e) - \phi_0(t, g)$]. By fitting the phase difference to a linear model in pulse length t , we determine the dispersive shift χ_{sa} from the ancilla on the storage mode. The flux

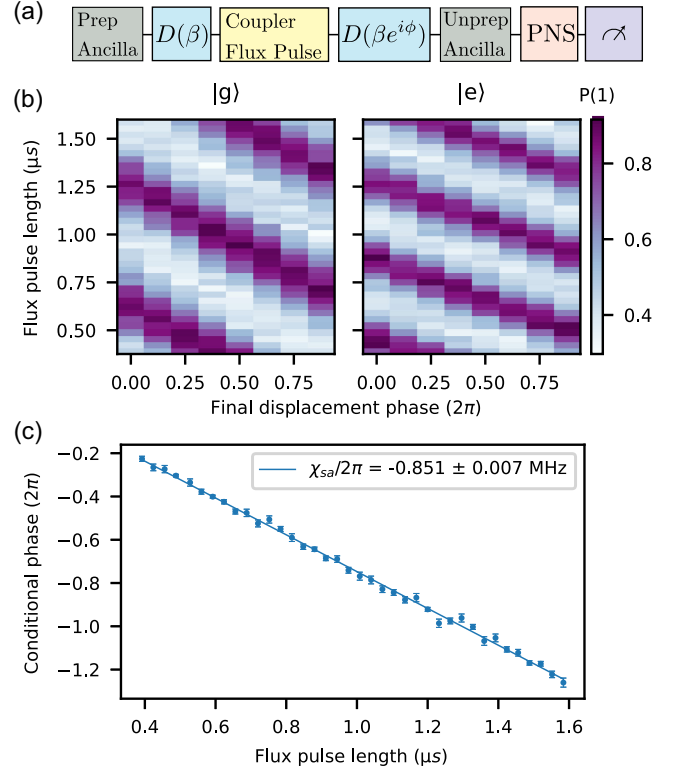


FIG. 14. Storage conditional phase measurement. (a) A high-level description of the pulse sequence for characterizing the storage-ancilla dispersive coupling strength. We prepare the ancilla into either $|g\rangle$ or $|e\rangle$, displace the storage, apply a coupler flux pulse with a variable pulse length, displace back the storage with a variable phase, unprepare the ancilla state (inverse of the state preparation), and finally measure the storage-mode vacuum population using a vacuum-selective pulse on the ancilla. (b) Fringes measured for two different ancilla states $|g\rangle$ and $|e\rangle$ as a function of the flux pulse length. From these fringes, we determine the storage conditional phase for each flux pulse length. (c) Storage conditional phase as a function of the flux pulse length. From a fitted linear relationship between the flux pulse length and the storage conditional phase, we extract $\chi_{sa}/2\pi = -0.851 \pm 0.007\ \text{MHz}$. This dataset was collected at the “parity measurement” position of the coupler.

pulse length required to achieve a π conditional phase is not exactly $\pi/|\chi_{sa}|$ because of transients of the flux pulse. To determine the parity measurement flux pulse length, we find where the storage conditional phase reaches $-\pi$. Note that, in addition to determining the dispersive coupling strength and flux pulse length for the parity measurement, we determine the frequency of the storage with the coupler in the on position by fitting the storage phase as a function of the flux pulse length with the ancilla in $|g\rangle$.

For the storage vacuum population measurements, a weak number-selective pulse is applied on the ancilla. The number-selective pulses we use have a duration of $1.4\ \mu\text{s}$ and a Gaussian pulse shape. Number-selective pulses are applied during the second half of a symmetric




FIG. 15. Measurement of storage-ancilla χ_{sa} with coupler in the off position. (a) Ramsey measurements of the ancilla for varying storage populations various values of the mean photon number in the storage mode (ranging from 0 to 20). (b) Fit of the ancilla frequency as a function of the storage-mode mean photon number to determine the storage-ancilla χ_{sa} .

flux waveform which pulses the coupler to the on position. Note that the extracted value of $\chi_{sa}/2\pi = -0.851 \pm 0.007$ MHz is reported in Table I (in the parity measurement position properties section).

APPENDIX K: MEASUREMENT OF THE OFF-POSITION STORAGE-ANCILLA DISPERSIVE INTERACTION STRENGTH

The storage conditional phase measurement is not effective at resolving the vanishingly small storage-ancilla dispersive coupling strength in the off position. Thus, for the off-position χ_{sa} measurement, we instead amplify the effects of the dispersive coupling by driving the storage into a highly excited coherent state (with a mean photon number up to 20). Specifically, we perform an ancilla Ramsey experiment with the storage in a variable amplitude coherent state as shown in Fig. 15(a). The ancilla experiences an average shift which scales as $\chi_{sa}|\alpha|^2$. The ancilla frequency shift as a function of the storage photon number is shown in Fig. 15(b). A linear fit to the frequency yields a dispersive coupling of $\chi_{sa}/2\pi = -1.4 \pm 0.1$ kHz.

APPENDIX L: STORAGE T_1 AND T_2 MEASUREMENTS

For both the T_1 and T_2 measurements, we start off by preparing the storage mode into the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$

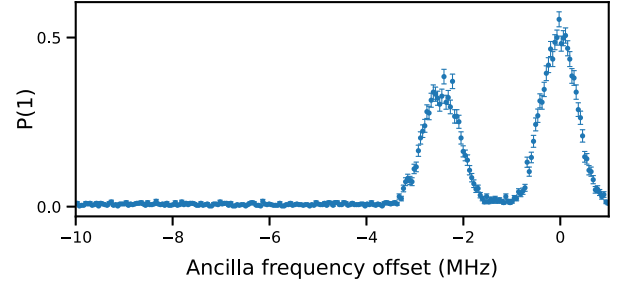


FIG. 16. Dissipative preparation of state in the $|\hat{n} = 0\rangle/|\hat{n} = 1\rangle$ manifold. Ancilla number splitting after applying two-photon dissipation for $8 \mu\text{s}$ starting from $|\alpha|^2 = 4$.

manifold using a pure two-photon dissipation without the buffer drive [18]. Specifically, we displace the storage mode to $|\alpha|^2 = 4$ and apply the pure two-photon dissipation for $8 \mu\text{s}$ (number-splitting spectroscopy of the prepared state is shown in Fig. 16). Next, we wait for a variable delay time with or without the pure two-photon dissipation being applied.

For T_1 measurements, the final step is measuring the storage-mode parity. Measurement of the storage-mode parity is performed with the coupler pulsed to the parity measurement flux position. The parity readout is symmetrized by using two phases for the final ancilla $\pi/2$ pulse of the parity measurement [31].

For T_2 measurements, in the final steps we displace the storage mode and measure the storage-mode vacuum population using a weak photon-number-selective π pulse applied to the transmon ancilla. The storage-mode displacement has an amplitude of $\alpha = 0.83 \simeq \ln \sqrt{1/2}$, which ensures that, when the storage decays to vacuum during the variable time delay, the final displacement results in approximately half of the population remaining in the vacuum state. In order to ensure that the final measurement is accurate before the steady state is reached, we symmetrize the final displacement. Specifically, we perform the final measurement with both positive and negative displacements. Without symmetrization, the use of a frame detuning during the Ramsey sequence results in a non-symmetric decay envelope.

The reported storage T_1 (T_2) values are determined by performing five (four) interleaved T_1 (T_2) measurements without and with the pure two-photon dissipation being applied during the variable time delay. All these interleaved datasets for each case are combined together and fit to determine T_1 and T_2 . The storage T_1 and T_2 measurements reported in the main text were taken directly after each other. On the other hand, the measurement of the storage-mode phase-flip rate was taken many hours later. Storage T_1 measurements with dissipation at the time of the phase-flip rate measurement yielded $74 \pm 1 \mu\text{s}$.

In Fig. 17(a), we show the results of the storage T_1 measurements intermittently performed over the course of a

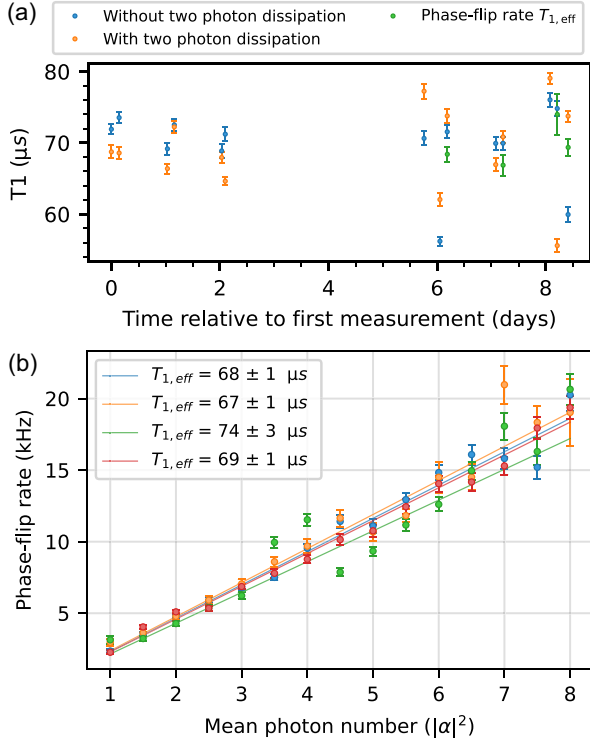


FIG. 17. Temporal variation of storage T_1 . (a) Measurements of the storage T_1 with (orange data points) and without (blue data points) dissipation being applied over time using the same procedure as Fig. 4. We also include points representing the $T_{1,eff}$ (green data points) extracted from fitting the phase-flip rate as a function of $|\alpha|^2$. (b) Data and fits corresponding to the phase-flip rate experiments that $T_{1,eff}$ was determined from in (a). Data are averaged over the even and odd initial cat states.

week. Each data point uses the same interleaving procedure. Finally, we remark that weeks before the data presented here we have observed lower T_1 ($< 50 \mu\text{s}$) on the storage mode. We also include examples of effective lifetimes taken from measurements of phase-flip rates. The phase-flip rate data corresponding to these points are shown in Fig. 17(b).

APPENDIX M: STORAGE-MODE DISPLACEMENT AND AMPLITUDE CALIBRATIONS

Here, we detail the methods we use for calibrating both the storage-mode displacement and buffer drive amplitude. Storage-mode displacement is calibrated by fitting the displaced vacuum state parity which is expected to follow $P(\alpha) \propto e^{-2|\alpha|^2}$, where α is the storage-mode displacement [57]. An example dataset and fit are shown in Fig. 18. In this experiment, we symmetrize the parity measurement readout by using two phases for the final ancilla $\pi/2$ pulse [31].

To calibrate the relation between the buffer drive amplitude and the size (i.e., average photon number $|\alpha|^2$)

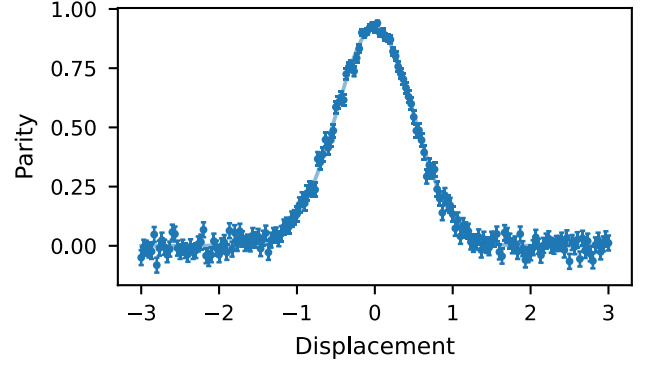


FIG. 18. Storage-mode displacement calibration. Storage-mode displacement is calibrated by fitting the displaced vacuum state parity.

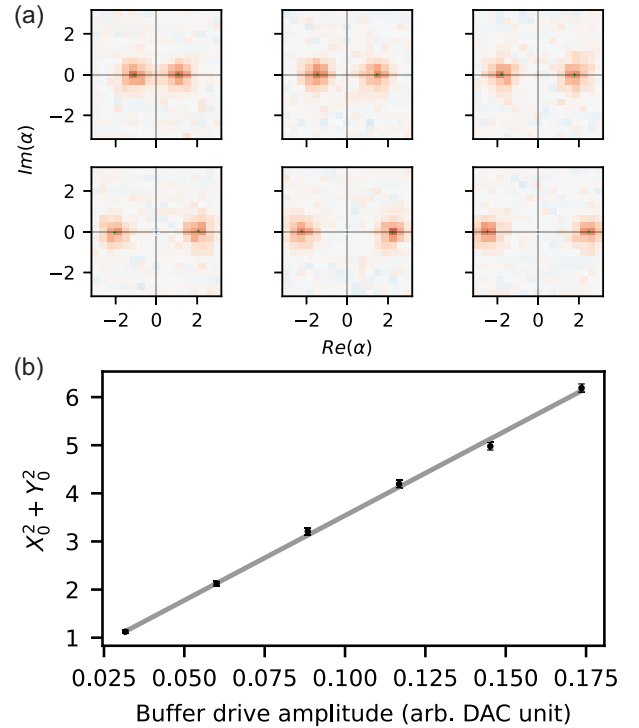


FIG. 19. Buffer drive amplitude calibration. (a) Measured Wigner tomogram of the steady state of the storage mode during two-photon dissipation for a series of buffer drive amplitudes. From left to right and top to bottom, the buffer drive amplitude is increasing from approximately 0.03 to approximately 0.175 (arbitrary units). (b) For each Wigner tomogram in (a), corresponding to a buffer drive amplitude, we plot the squared distance from the origin in phase space ($X_0^2 + Y_0^2$) determined by fitting the Wigner tomogram to two diametrically opposed Gaussians. We use the linear fit to this curve as a calibration of storage mean photon number versus buffer drive amplitude.

of the steady-state cat-qubit states under the two-photon dissipation, we measure the storage-mode Wigner functions for a few different values of the buffer drive amplitude as shown in Fig. 19(a) [16]. We reach the steady-state

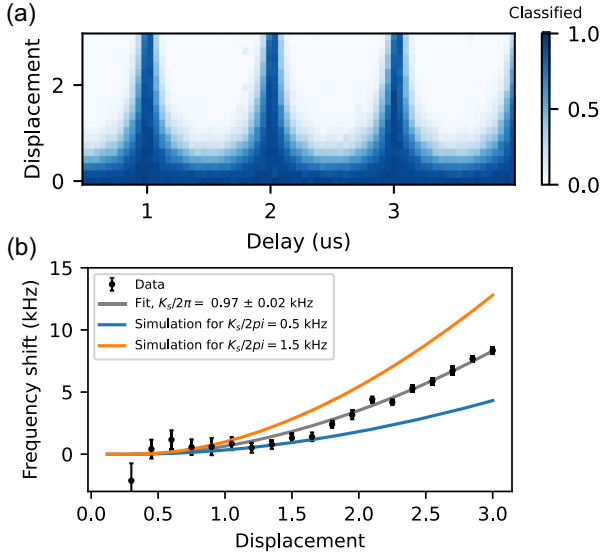


FIG. 20. Storage self-Kerr fitting. (a) Measured occupancy of the storage-mode vacuum (color scale bar) versus displacement amplitude and time delay for a storage Ramsey measurement [58]. (b) Plot of the frequency offset of the storage-mode frequency versus displacement amplitude of the storage mode, from fits to Ramsey measurements as in (a). A fit to a simulation of the experiment yields $K_s/2\pi = 0.97 \pm 0.02$ kHz. Simulations for $K_s/2\pi = 0.5$ and 1.5 kHz are shown as red and blue curves, respectively. Note that two points with very low displacements are excluded, since they have poor contrast for fitting. Also, the frequency offset is defined relative to the average of the first four data points of the curve.

manifold by applying two-photon dissipation for $200 \mu\text{s}$. For each buffer drive amplitude, we fit the storage-mode Wigner function to two diametrically opposed Gaussians. For points near $|\alpha|^2 = 1$, this is not exactly the correct model, because the storage-mode parity is nonzero. In simulations, we find that this effect causes only fits at $|\alpha|^2 = 1$ to disagree with the storage-mode photon number by approximately 3%. We fit the storage-mode steady-state mean photon number as a function of buffer drive amplitude to a linear model to complete the calibration as shown in Fig. 19(b).

APPENDIX N: STORAGE SELF-KERR FITTING

To determine the storage-mode self-Kerr nonlinearity, we perform a storage Ramsey experiment for coherent states of the storage mode with varying coherent-state amplitudes. This involves displacing the storage mode, waiting for a variable time, displacing the storage mode back with a time-dependent phase, and finally measuring the storage-mode vacuum population [58]. Note that, since the storage mode can be highly excited in this experiment, we use a coupler flux position for the number-selective measurements near that of the parity measurement flux position. The vacuum populations from running this

experiment are shown in Fig. 20(a). We determine the average storage frequency for each displacement by fitting to the functional form

$$Ae^{-|\alpha|^2[2-2\cos(2\pi ft)]}. \quad (\text{N1})$$

The free parameters are A , which is a scaling factor, and f , which is the frequency of the storage mode. In Fig. 20(b), we show the fit storage-mode frequency offset as a function of the displacement from the experimental data. Fitting the frequency offset curve to a simulation of the experiment yields a storage self-Kerr nonlinearity of $K_s/2\pi = 0.97 \pm 0.02$ kHz.

APPENDIX O: TWO-PHOTON DISSIPATION WITH ANCILLA AND COUPLER EXCITATIONS

To provide further insights into the saturation behavior of the bit-flip times of our cat qubit, we measure the bit-flip times while we repeatedly excite an aggressor qubit [sequence shown schematically in Fig. 21(a)]. If the frequency shift from the aggressor qubit excitations is sufficiently strong to overcome the confinement rate of the two-photon dissipation, we expect to observe lower saturated bit-flip times than what we observe without the aggressor qubit excitations. In the sequence, we excite the aggressor qubit every $500 \mu\text{s}$, far faster than the observed saturated bit-flip time (approximately 0.5 s) in the absence of the aggressor excitation. Note that, in the experiment here, the value of the buffer detuning (i.e., the difference between the buffer drive frequency and the Stark-shifted buffer frequency) is slightly different from the one used to generate the data in the main text. However, the saturated bit-flip times are not sensitively dependent on the buffer detuning, and we observe similar saturated bit-flip times even with a slightly different buffer detuning in this experiment.

In Fig. 21(b), we show the bit-flip time of our cat qubit when the aggressor qubit is the ancilla. The blue curve is a reference where the ancilla is not excited. We consider the case of repeatedly exciting the ancilla to $|e\rangle$ and $|f\rangle$. In both cases, we observe no significant change in the saturated bit-flip times. This is consistent with the small storage-ancilla dispersive shift (i.e., $|\chi_{sa}| = 2\pi \times 1.4 \text{ kHz} \ll \kappa_2|\alpha|^2$) measured with the coupler in the off position (see Appendix K).

Since the storage mode is directly coupled to the couplers, the dispersive coupling of the storage mode to the coupler resonances is more sizable (e.g., predicted to be approximately -100 kHz in the off position by the model in Appendix H 1) than the dispersive coupling to the ancilla. In Fig. 21(c), we numerically study the impact of such storage-coupler dispersive coupling. In particular, in addition to the storage-buffer Hamiltonian and dissipation terms considered in Appendix H 4, we further add the storage-coupler dispersive coupling $\chi_{sc}\hat{a}^\dagger\hat{a}\hat{c}^\dagger\hat{c}$ to the Hamiltonian, where \hat{c} and \hat{c}^\dagger are the annihilation and

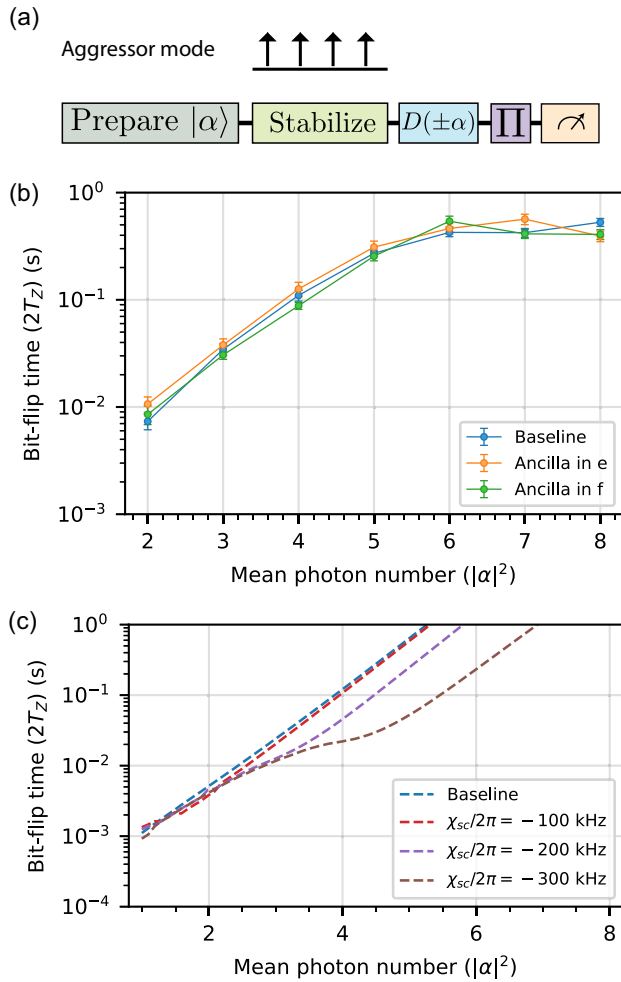


FIG. 21. Bit-flip times with aggressor excitation. (a) Experimental protocol used to perform bit-flip characterization while exciting an aggressor mode. In the data presented here, the aggressor mode is excited every $500 \mu\text{s}$. (b) Measured bit-flip times when the aggressor mode is the transmon ancilla, from which we infer that excitation of the ancilla does not affect the storage-mode bit-flip times. (c) Numerically simulated bit-flip times when a thermally excited coupler mode is included. The different dashed curves (blue, red, purple, and brown) correspond to storage-coupler dispersive couplings of $\chi_{sc}/2\pi = (0, -100, -200, -300)$ kHz, respectively.

creation operators, respectively, of the coupler mode. Additionally, we add additional dissipators associated with the coupler loss $\kappa_c \mathcal{D}[\hat{c}]$ and heating $\kappa_{c,\text{up}} \mathcal{D}[\hat{c}^\dagger]$ to the Lindblad dissipator. Guided by the circuit-quantization predictions, we consider storage-coupler dispersive couplings of $\chi_{sc}/2\pi = 0, -100, -200, -300$ kHz. Also, we assume $\kappa_c = 1/(10 \mu\text{s})$ and $\kappa_{c,\text{up}}/\kappa_c = 10^{-3}$.

The simulation results suggest that the cat-qubit bit-flip times are not degraded when the storage-coupler dispersive coupling strength is $\chi_{sc}/2\pi = -100$ kHz (dashed red curve) compared to the baseline case with $\chi_{sc} = 0$ (dashed blue curve). Notably, these simulations include higher

excited states of the coupler up to the third excited state. Thus, they account for enhanced dispersive shift of the storage-mode frequency under these highly excited coupler states. For larger values of storage-coupler dispersive coupling with $\chi_{sc}/2\pi = -200$ kHz (dashed purple curve) and $\chi_{sc}/2\pi = -300$ kHz (dashed brown curve), a degradation in the bit-flip time starts to become noticeable, especially in the small mean photon number regime, e.g., $|\alpha|^2 \lesssim 4$. However, these coupler-induced limitations are overcome, and the exponentially increasing trend of the bit-flip times is quickly recovered at larger values of $|\alpha|^2 \geq 4$ as the confinement rate of the two-photon dissipation surpasses the storage-coupler dispersive shift. In stark contrast, the cat-qubit bit-flip times in our experiment stay saturated at around 0.3–0.6 s even at large values of $|\alpha|^2$ above 10. These results show that heating of the ancilla and coupler, even to higher excited states beyond the first excited state, are likely not the reason for the saturated bit-flip times of our cat qubit.

APPENDIX P: SIMULATION OF LOOP INDUCTANCE

The inductances $L_{P,1}$ and $L_{P,2}$ in series with the ATS side junctions are estimated during device design from finite-element magnetostatic simulation of the buffer. In this type of simulation, the ATS metal geometry is modeled as a perfectly conducting surface as though the junctions are electrical shorts, with a junction lead on each side interrupted by a current source. For each current source, the magnetic field solution is computed by applying current on that source, leaving the other source electrically open. Each solution, thus, corresponds to current flowing around one side loop.

The self-inductance of each side loop is calculated from the magnetic energy in the corresponding solution, while a mutual inductance between the two loops is calculated based on the excess energy of the superposition of the two solutions. We attribute the loop self-inductances mainly to current in the side and middle junction leads, while the mutual inductance is attributed mainly to the shared current path through the middle. Specifically, we interpret the results as an inductance matrix describing a lumped-element circuit consisting of a loop with inductors $L_{P,1}$ and $L_{P,2}$ on each side, corresponding to the outer SQUID loop geometry of the ATS, which is bisected by an inductance equivalent to the calculated mutual inductance.

To estimate the desired serial inductances, we thus subtract the mutual inductance from each loop's self-inductance. This yields an inductance around 20 pH on each side, compared to 27 pH from the fit to experimental data in Fig. 5(b). We note that other effects (e.g., kinetic inductance in the junction leads) which are not included in the above inductance calculations might account for the discrepancy.

APPENDIX Q: FLUX-LINE-INDUCED LOSS

The flux lines used to control the buffer also present an environment to the buffer without the protection of a bandpass filter. For that reason, photon loss through the flux lines may still affect the storage mode's lifetime, even if the buffer mode's lifetime is dominated by loss through its filtered output line.

To understand loss induced by the flux lines, we consider an isolated buffer as in Appendix H 3, statically parked at a saddle point, but now allowing for time-dependent flux due to fluctuating current in both buffer flux-pump lines. Following the procedure for circuit quantization with time-dependent flux described in Refs. [59,60], we obtain a Hamiltonian in terms of "irrotational" degrees of freedom \tilde{N}_b and $\tilde{\phi}_b$ such that all time-dependent pump fluctuations are assigned to the inductive potential:

$$\begin{aligned}
 H_{\text{irr}}(\tilde{N}_b, \tilde{\phi}_b) &= 4E_C \tilde{N}_b^2 - NE_{J,\text{array}} \cos \frac{1}{N} (\tilde{\phi}_b - \varphi_\Delta(t) + \varphi_\delta(t)) \\
 &\quad - (E_{J,1} + E_{J,2}) \sin \varphi_\Sigma(t) \sin (\tilde{\phi}_b + \varphi_\delta(t)) \\
 &\quad - (E_{J,1} - E_{J,2}) \cos \varphi_\Sigma(t) \cos (\tilde{\phi}_b + \varphi_\delta(t)). \quad (\text{Q1})
 \end{aligned}$$

Above, the sigma-flux fluctuation $\varphi_\Sigma(t)$ is assigned to the side junctions, while the delta-flux fluctuation $\varphi_\Delta(t)$ is assigned to the array. Each branch also sees a common offset $\varphi_\delta(t)$, which can be interpreted as related to magnetic field nonuniformity in the buffer's capacitive gap (not just the ATS loops). We describe these fluxes as arising from inductive couplings to current $I_\Sigma(t)$, $I_\Delta(t)$ in each flux line. To estimate rates of photon emission into the flux lines, we then expand the above Hamiltonian assuming small fluctuations and consider linear couplings $I(t)\tilde{\phi}_b$.

For example, we consider the case where the only time-dependent flux is $\varphi_\Sigma(t) = (2\pi/\Phi_0)M_\Sigma I_\Sigma(t)$, where the inductive coupling M_Σ is required for operating the buffer pump. That is, we ignore both $\varphi_\Delta(t)$ and $\varphi_\delta(t)$, as well as any unintended coupling between the delta-flux line and the sigma flux. This yields the term of interest $-(E_{J,1} + E_{J,2})(2\pi/\Phi_0)M_\Sigma I_\Sigma(t)\tilde{\phi}_b$, corresponding to a loss rate

$$\kappa_\Sigma = \frac{\phi_{\text{zpf}}^2}{\hbar^2} \left(\frac{2\pi}{\Phi_0} (E_{J,1} + E_{J,2}) M_\Sigma \right)^2 S_{I_\Sigma I_\Sigma}(\omega) \quad (\text{Q2})$$

for a mode at frequency ω with zero-point fluctuation ϕ_{zpf} associated with ϕ_b , where $S_{I_\Sigma I_\Sigma}(\omega)$ is the noise power spectrum of the current $I_\Sigma(t)$. For example, if the buffer sees a real load impedance of Z_Σ at zero temperature through the flux line, then $S_{I_\Sigma I_\Sigma}(\omega) = 2\hbar\omega/Z_\Sigma$. Together with finite-element simulation to estimate inductive couplings, further calculation shows that this is the largest expected contribution to flux-line-induced loss. Accordingly, our flux line design targets a minimal M_Σ .

- [1] J. O'Gorman and E. T. Campbell, *Quantum computation with realistic magic-state factories*, *Phys. Rev. A* **95**, 032338 (2017).
- [2] I. D. Kivlichan *et al.*, *Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via Trotterization*, *Quantum* **4**, 296 (2020).
- [3] C. Gidney and M. Ekerå, *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*, *Quantum* **5**, 433 (2021).
- [4] P. T. Cochrane, G. J. Milburn, and W. J. Munro, *Macroscopically distinct quantum-superposition states as a bosonic code for amplitude damping*, *Phys. Rev. A* **59**, 2631 (1999).
- [5] P. Aliferis and J. Preskill, *Fault-tolerant quantum computation against biased noise*, *Phys. Rev. A* **78**, 052331 (2008).
- [6] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, *Ultra-high error threshold for surface codes with biased noise*, *Phys. Rev. Lett.* **120**, 050505 (2018).
- [7] J. Guillaud and M. Mirrahimi, *Repetition cat qubits for fault-tolerant quantum computation*, *Phys. Rev. X* **9**, 041053 (2019).
- [8] J. Guillaud and M. Mirrahimi, *Error rates and resource overheads of repetition cat qubits*, *Phys. Rev. A* **103**, 042413 (2021).
- [9] 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).
- [10] J. P. Bonilla Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, *The XZZX surface code*, *Nat. Commun.* **12**, 2172 (2021).
- [11] C. Chamberland *et al.*, *Building a fault-tolerant quantum computer using concatenated cat codes*, *PRX Quantum* **3**, 010329 (2022).
- [12] F.-M. L. Régent, C. Berdou, Z. Leghtas, J. Guillaud, and M. Mirrahimi, *High-performance repetition cat code using fast noisy operations*, *Quantum* **7**, 1198 (2023).
- [13] E. 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).
- [14] D. Ruiz, J. Guillaud, A. Leverrier, M. Mirrahimi, and C. Vuillot, *LDPC-cat codes for low-overhead quantum computing in 2D*, *Nat. Commun.* **16**, 1040 (2025).
- [15] H. Jeong and M. S. Kim, *Efficient quantum computation using coherent states*, *Phys. Rev. A* **65**, 042305 (2002).
- [16] R. Lescanne *et al.*, *Exponential suppression of bit-flips in a qubit encoded in an oscillator*, *Nat. Phys.* **16**, 509 (2020).
- [17] C. Berdou *et al.*, *One hundred second bit-flip time in a two-photon dissipative oscillator*, *PRX Quantum* **4**, 020350 (2023).
- [18] A. Marquet *et al.*, *Autoparametric resonance extending the bit-flip time of a cat qubit up to 0.3 s*, *Phys. Rev. X* **14**, 021019 (2024).
- [19] U. Réglade *et al.*, *Quantum control of a cat qubit with bit-flip times exceeding ten seconds*, *Nature (London)* **629**, 778 (2024).
- [20] M. Mirrahimi *et al.*, *Dynamically protected cat-qubits: A new paradigm for universal quantum computation*, *New J. Phys.* **16**, 045014 (2014).

- [21] Z. Leghtas *et al.*, *Confining the state of light to a quantum manifold by engineered two-photon loss*, *Science* **347**, 853 (2015).
- [22] S. Touzard *et al.*, *Coherent oscillations inside a quantum manifold stabilized by dissipation*, *Phys. Rev. X* **8**, 021005 (2018).
- [23] V. S. Ferreira *et al.*, *Collapse and revival of an artificial atom coupled to a structured photonic reservoir*, *Phys. Rev. X* **11**, 041043 (2021).
- [24] M. Mirhosseini *et al.*, *Superconducting metamaterials for waveguide quantum electrodynamics*, *Nat. Commun.* **9**, 3706 (2018).
- [25] B. Foxen *et al.*, *Qubit compatible superconducting interconnects*, *Quantum Sci. Technol.* **3**, 014005 (2017).
- [26] R. N. Das *et al.*, *Cryogenic qubit integration for quantum computing*, in *Proceedings of the 2018 IEEE 68th Electronic Components and Technology Conference (ECTC)* (IEEE, New York, 2018), pp. 504–514.
- [27] A. P. M. Place *et al.*, *New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds*, *Nat. Commun.* **12**, 1779 (2021).
- [28] L. G. Lutterbach and L. Davidovich, *Method for direct measurement of the Wigner function in cavity QED and ion traps*, *Phys. Rev. Lett.* **78**, 2547 (1997).
- [29] D. I. Schuster *et al.*, *Resolving photon number states in a superconducting circuit*, *Nature (London)* **445**, 515 (2007).
- [30] 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).
- [31] L. Sun *et al.*, *Tracking photon jumps with repeated quantum non-demolition parity measurements*, *Nature (London)* **511**, 444 (2014).
- [32] N. Ofek *et al.*, *Extending the lifetime of a quantum bit with error correction in superconducting circuits*, *Nature (London)* **536**, 441 (2016).
- [33] K. S. Chou *et al.*, *Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements*, [arXiv:2307.03169](https://arxiv.org/abs/2307.03169).
- [34] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *Tunable coupling scheme for implementing high-fidelity two-qubit gates*, *Phys. Rev. Appl.* **10**, 054062 (2018).
- [35] Y. Sung *et al.*, *Realization of high-fidelity cz and zz-free iswap gates with a tunable coupler*, *Phys. Rev. X* **11**, 021058 (2021).
- [36] H. Putterman, J. Iverson, Q. Xu, L. Jiang, O. Painter, F. G. S. L. Brandao, and K. Noh, *Stabilizing a bosonic qubit using colored dissipation*, *Phys. Rev. Lett.* **128**, 110502 (2022).
- [37] C. Wang *et al.*, *Towards practical quantum computers: Transmon qubit with a lifetime approaching 0.5 milliseconds*, *npj Quantum Inf.* **8**, 3 (2022).
- [38] D. Kafri, C. Quintana, Y. Chen, A. Shabani, J. M. Martinis, and H. Neven, *Tunable inductive coupling of superconducting qubits in the strongly nonlinear regime*, *Phys. Rev. A* **95**, 052333 (2017).
- [39] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret, *Optimizing the nonlinearity and dissipation of a snail parametric amplifier for dynamic range*, *Phys. Rev. Appl.* **10**, 054020 (2018).
- [40] C. Quintana, *Superconducting flux qubits for high-connectivity quantum annealing without lossy dielectrics*, Ph.D. thesis, UC Santa Barbara, 2017, <https://escholarship.org/uc/item/9844c3h3>.
- [41] K. S. Dubovitskii, *Bit-flip errors in dissipative cat qubits: Second-order perturbation theory*, *Phys. Rev. A* **111**, 012617 (2025).
- [42] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological quantum memory*, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [43] S. Puri *et al.*, *Bias-preserving gates with stabilized cat qubits*, *Sci. Adv.* **6**, eaay5901 (2020).
- [44] R. Gautier, A. Sarlette, and M. Mirrahimi, *Combined dissipative and Hamiltonian confinement of cat qubits*, *PRX Quantum* **3**, 020339 (2022).
- [45] H. Putterman *et al.*, *Hardware-efficient quantum error correction via concatenated bosonic qubits*, *Nature (London)* **638**, 927 (2025).
- [46] M. McEwen *et al.*, *Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits*, *Nat. Phys.* **18**, 107 (2022).
- [47] R. Acharya *et al.*, *Quantum error correction below the surface code threshold*, *Nature (London)* **638**, 920 (2025).
- [48] A. Dunswoth *et al.*, *Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits*, *Appl. Phys. Lett.* **111**, 022601 (2017).
- [49] A. J. Keller *et al.*, *Al transmon qubits on silicon-on-insulator for quantum device integration*, *Appl. Phys. Lett.* **111**, 042603 (2017).
- [50] K. D. Crowley *et al.*, *Disentangling losses in tantalum superconducting circuits*, *Phys. Rev. X* **13**, 041005 (2023).
- [51] V. Negîrnea *et al.*, *High-fidelity controlled-z gate with maximal intermediate leakage operating at the speed limit in a superconducting quantum processor*, *Phys. Rev. Lett.* **126**, 220502 (2021).
- [52] J. Koch *et al.*, *Charge-insensitive qubit design derived from the Cooper pair box*, *Phys. Rev. A* **76**, 042319 (2007).
- [53] H. Sambe, *Steady states and quasienergies of a quantum-mechanical system in an oscillating field*, *Phys. Rev. A* **7**, 2203 (1973).
- [54] A. Di Paolo *et al.*, *Extensible circuit-QED architecture via amplitude- and frequency-variable microwaves*, [arXiv:2204.08098](https://arxiv.org/abs/2204.08098).
- [55] F. Reiter and A. S. Sørensen, *Effective operator formalism for open quantum systems*, *Phys. Rev. A* **85**, 032111 (2012).
- [56] S. Puri *et al.*, *Stabilized cat in a driven nonlinear cavity: A fault-tolerant error syndrome detector*, *Phys. Rev. X* **9**, 041009 (2019).
- [57] G. Kirchmair *et al.*, *Observation of quantum state collapse and revival due to the single-photon Kerr effect*, *Nature (London)* **495**, 205 (2013).
- [58] P. Reinhold, *Controlling error-correctable bosonic qubits*, Ph.D. thesis, Yale, 2019, [https://rsl.yale.edu/sites/default/files/files/RSL_Theses/Reinhold-Thesis%20\(1\).pdf](https://rsl.yale.edu/sites/default/files/files/RSL_Theses/Reinhold-Thesis%20(1).pdf).
- [59] X. You, J. A. Sauls, and J. Koch, *Circuit quantization in the presence of time-dependent external flux*, *Phys. Rev. B* **99**, 174512 (2019).
- [60] R.-P. Riwar and D. P. DiVincenzo, *Circuit quantization with time-dependent magnetic fields for realistic geometries*, *npj Quantum Inf.* **8**, 36 (2022).