
Supplementary information

Real-time quantum error correction beyond break-even

In the format provided by the
authors and unedited

Supplementary Information

“Real-time quantum error correction beyond break-even”

V. V. Sivak,^{1,2,*} A. Eickbusch,^{1,2} B. Royer,^{1,2,3} S. Singh,^{1,2} I. Tsioutsios,^{1,2} S. Ganjam,^{1,2} A. Miano,^{1,2}
B. L. Brock,^{1,2} A. Z. Ding,^{1,2} L. Frunzio,^{1,2} S. M. Girvin,^{1,2} R. J. Schoelkopf,^{1,2} and M. H. Devoret^{1,2,†}

¹*Departments of Physics and Applied Physics, Yale University, New Haven, CT 06520, USA*

²*Yale Quantum Institute, Yale University, New Haven, CT 06511, USA*

³*Institut Quantique and Département de Physique,
Université de Sherbrooke, Québec, J1K 2R1, Canada*

CONTENTS

I. Experimental setup and sample parameters	2
II. Calibration and characterization experiments	3
A. Primitive pulses	3
B. Hamiltonian parameters	5
C. Readout and reset	5
D. Conditional displacement	7
E. Oscillator error channels	8
F. Active oscillator cooling	9
III. Quantum control optimization	11
A. Model-based optimization of control circuits	11
B. Model-free reinforcement learning for QEC	13
IV. Quantum error correction of the grid code	14
A. Brief introduction to grid code	14
B. Small-Big-Small (SBS) protocol	15
C. QEC cycle: implementation details	17
D. Learned and scripted parameters	19
E. Syndrome measurement statistics	20
F. Analysis of transmon leakage	22
G. Post-selection of errors	24
H. Wigner tomography of logical states	24
I. Sensitivity to auxiliary qubit errors	27
J. Long-time system stability	29
K. Average channel fidelity	30
References	32

* vladimir.sivak@yale.edu

† michel.devoret@yale.edu

I. EXPERIMENTAL SETUP AND SAMPLE PARAMETERS

Assembly. Our system design follows the hybrid planar-3D circuit QED architecture developed in [1]. The storage oscillator is realized as an electromagnetic mode hosted by a seamless superconducting coaxial stub cavity made of high-purity aluminum and treated with a chemical etch to improve surface quality. This is the same physical cavity as used in [2], although with lower coherence time due to aging during the storage time of ~ 2 years. The cavity is anchored to a copper bracket inside a cryoperm shield. The auxiliary qubit chip is inserted in a tunnel waveguide that connects to the storage cavity, and is secured at one end with a copper clamp. Thermalizing copper braids run from the clamp to the base plate of the dilution refrigerator, see Supp. Fig. 1.

Auxiliary qubit chip fabrication. The auxiliary qubit chip contains a transmon qubit, a stripline readout resonator, and a stripline bandpass Purcell filter. The resonators and transmon are tantalum-based devices, with Josephson junction of the transmon made of a small aluminum section; they are fabricated with a process similar to [3]. Adopting a tantalum-based platform results in improved qubit coherence relative to an all-aluminum platform; however, the reasons for this are still under active investigation. Possible explanations include, but are not limited to: 1) The corrosion resistance allows for the use of rigorous acid-based cleaning techniques to be employed during the fabrication process that improves surface dielectric quality and minimizes the presence of organic residues; 2) The high melting point of tantalum allows for deposition to occur at higher temperatures, where atomic mobility is high enough to enable epitaxial film growth with a high degree of crystalline order; 3) Tantalum has a higher superconducting transition temperature than aluminum, which may lead to increased resistance to quasiparticle loss.


We use a C-plane sapphire wafer produced using the heat exchanger method (HEM), as it was shown to have smaller dielectric loss [4]. The wafer was initially cleaned with a piranha solution ($2 : 1 \text{ H}_2\text{SO}_4 : \text{H}_2\text{O}_2$) for 20 minutes and rinsed with DI water. The wafer was then annealed at 1200°C in an oxygen-rich environment for 1 hour. After cooling down to room temperature, the wafer was immediately transferred to a sputtering system for tantalum deposition. 150 nm of tantalum was deposited by DC magnetron sputtering with the substrate temperature being held at 800°C . The Purcell filter, readout resonator, and transmon pads were subtractively patterned using a positive photoresist mask and reactive ion etching. After tantalum patterning, the Josephson junction was patterned using electron-beam lithography and defined using the Dolan bridge method. The junction was deposited using electron-beam evaporation of aluminum at 2 angles with an interleaved static oxidation step to construct the tunnel barrier. Liftoff was performed in NMP heated to 90°C , followed by sonication in acetone, isopropanol, and DI water. Finally, the wafer was protected with a layer of photoresist before dicing into individual chips, followed by additional cleaning with NMP, acetone, and isopropanol to remove the protective photoresist.

System parameters. The measured parameters of this system are summarized in Supp. Table I.

Cavity mode	Frequency	$\omega_c = 2\pi \times 4.479 \text{ GHz}$
	1st order dispersive shift	$\chi = 2\pi \times 46.5 \text{ kHz}$
	2nd order dispersive shift	$\chi' = 2\pi \times 5.8 \text{ Hz}$
	Kerr nonlinearity	$K = -2\pi \times 4.8 \text{ Hz}$
	Relaxation	$\bar{T}_1^c = 606 \pm 10 \text{ us}$
	Dephasing	$\bar{T}_2^c = 980 \pm 30 \text{ us}$
Auxiliary transmon	Frequency	$\omega_t = 2\pi \times 5.921 \text{ GHz}$
	Anharmonicity	$\alpha = -2\pi \times 222 \text{ MHz}$
	Relaxation	$\bar{T}_1^t = 280 \pm 30 \text{ us}$
	Equilibrium population	$\bar{n}_{\text{th}}^t = 0.043 \pm 0.008$
	Dephasing (Ramsey)	$\bar{T}_{2R}^t = 62 \pm 5 \text{ us}$
	Dephasing (Echo)	$\bar{T}_{2E}^t = 238 \pm 8 \text{ us}$
Readout resonator	Frequency	$\omega_r = 2\pi \times 9.107 \text{ GHz}$
	Dispersive shift	$\chi_{qr} = 2\pi \times 0.60 \text{ MHz}$
	Coupling strength	$\kappa_{r(c)} = 2\pi \times 0.47 \text{ MHz}$
	Internal loss	$\kappa_{r(i)} = 2\pi \times 0.03 \text{ MHz}$

Supplementary Table I. **Measured system parameters.** For transmon parameters ($\bar{T}_1^t, \bar{T}_{2R}^t, \bar{T}_{2E}^t, \bar{n}_{\text{th}}^t$) and cavity parameters ($\bar{T}_1^c, \bar{T}_2^c$) that appreciably fluctuate in time, we provide the mean and standard deviation over a week-long period. The definition of the Hamiltonian parameters can be found in Section II B.

Control wiring. As shown in Supp. Fig. 2, the quantum system is controlled by a classical computer (VPXI-ePC) that hosts two control cards (X6-1000M) from Innovative Integration. Each card integrates digital-to-analog converters (DAC), analog-to-digital converters (ADC), digital inputs and outputs (DIO), and a Xilinx VIRTEX-6 field-programmable gate array (FPGA). This controller was developed in [5] and used in prior bosonic QEC experiments [5, 6]. The baseband control signals are sampled from the DACs at 500 MS/s rate with 16-bit resolution and upconverted



Supplementary Figure 1. **Sample assembly.** (a) Clamped auxiliary qubit chip. (b) Thermalization. (c) Shielding.


to the oscillator, qubit, and readout frequencies through single-sideband modulation of the local oscillators (Agilent N5183A). After amplification, the signals are gated with fast RF switches (9 ns rise time, 5 ns fall time) and filtered before entering the dilution refrigerator. The signals are further attenuated and filtered in the cryogenic environment. A crucial component of the filtering is the eccosorb CR-110 infrared absorber filter [7] located inside the cryoperm shield, and the copper plate, painted with stycast epoxy mixed with black carbon powder, that wraps around the sample. On the output side, the reflected readout signal is amplified at 30 mK stage with a near-quantum-limited Josephson array-mode parametric amplifier (JAMPA) [8], followed at 4 K stage by a low-noise HEMT amplifier. Upon further amplification at 300 K stage and down-conversion to 50 MHz, the readout signal is digitized, demodulated, and integrated with a filter function to obtain I and Q quadratures. Their values are compared to the decision boundaries I_{th} and Q_{th} to obtain two bits of information $s_0 = \Theta(I - I_{th})$ and $s_1 = \Theta(Q - Q_{th})$, where Θ is the Heaviside step function. This allows to classify the measurement outcome as “ g ” if $s_0 = 0$, “ e ” if $(s_0, s_1) = (1, 0)$, and “ f ” if $(s_0, s_1) = (1, 1)$. The bits s_0 and s_1 are redistributed to all control cards which run independent but synchronized control flows that include conditional branching on these bits. Further details of the readout subroutine are described in Section II C.

II. CALIBRATION AND CHARACTERIZATION EXPERIMENTS


A. Primitive pulses

Qubit rotation. The waveform for transmon $g \leftrightarrow e$ and $e \leftrightarrow f$ rotations is a Gaussian with $\sigma = 8$ ns and symmetric chop at 2σ . The pulse amplitude is calibrated with a standard amplitude Rabi experiment, shown in Supp. Fig. 3(a). We find that finite negative detuning of a few MHz is needed to maximize the Rabi contrast in both cases. In a similar manner, we calibrate a selective square pulse of duration $2\pi/\chi \approx 22 \mu\text{s}$ that performs $g \leftrightarrow e$ rotations conditioned on the oscillator in $|0\rangle$.


Oscillator displacement. The waveform for oscillator displacements is a Gaussian with $\sigma = 10$ ns and symmetric chop at 2σ . It is calibrated in several steps, refining the accuracy at each step. First, before the precise value of χ is known, we use a rough calibration by creating a coherent state of unknown amplitude α and measuring the probability of $|0\rangle$ via a selective qubit pulse, with the results shown in Supp. Fig. 3(b). Fitting the data to $P(0) = e^{-|\alpha|^2}$ allows us to calibrate the DAC amplitude for displacement of $|\alpha| = 1$. This first-stage calibration enables us to use active oscillator cooling, see Section II F, which is important for the next calibration step that relies on a vacuum state. Next, after determining the value of χ (using number-resolved qubit spectroscopy, see Section II B), we measure the Wigner function of vacuum $W(\alpha) = (2/\pi)e^{-2|\alpha|^2}$ and adjust the DAC amplitude calibration to obtain the variance of $1/4$, with the results shown in Supp. Fig. 3(c). We find that these two calibrations typically agree within 2%.



Supplementary Figure 2. **Experimental setup.** For simplicity, the diagram omits DC ground connections of all active components, and attenuators placed at different locations at 300 K stage to ensure power levels within specs for amplifiers and switches.



Supplementary Figure 3. **Calibration of primitive pulses.** (a) Amplitude Rabi experiment to calibrate qubit rotations. (b) First-step calibration of displacement: probability of $|0\rangle$ in a coherent state, $P(0) = e^{-|\alpha|^2}$. (c) Second-step calibration of displacement: Wigner function of vacuum, $W(\alpha) = (2/\pi)e^{-2|\alpha|^2}$.



Supplementary Figure 4. **Calibration of Hamiltonian parameters.** (a) Number-resolved qubit spectroscopy with a selective square pulse of duration $\sim 50 \mu\text{s}$ when the oscillator is in the vacuum state (blue) and a coherent state (red). (b) Optimal return phase in the out-and-back experiment (inset) with the qubit in $|g\rangle$ and $|e\rangle$. As seen from the phase dispersion with $\bar{n}$, the effective oscillator nonlinearity is larger when the qubit is in $|e\rangle$. (c) Average oscillator rotation frequency $(\omega_g + \omega_e)/2$, and a linear fit to extract Δ and K . (d) Relative oscillator rotation frequency $\omega_g - \omega_e$, and a linear fit to extract χ and χ' . The star indicates χ obtained in (a).

B. Hamiltonian parameters

Our system is well described with the following Hamiltonian

$$H/\hbar = \Delta (a^\dagger a) + \frac{1}{2}\chi (a^\dagger a) \sigma_z + \frac{1}{2}K (a^\dagger a)^2 + \frac{1}{4}\chi' (a^\dagger a)^2 \sigma_z, \quad (1)$$


where Δ is the oscillator frequency detuning in the chosen rotating frame, χ is the dispersive shift, χ' is the second-order dispersive shift, and K is the Kerr nonlinearity.

We calibrate χ with number-resolved qubit spectroscopy [9] in the presence of a coherent state of small amplitude $\alpha \approx 0.6$ in the oscillator. The spectroscopy data, shown in Supp. Fig. 4(a) in red, is fitted to a 5-component equal-spacing mixture of the spectroscopy lineshapes with the oscillator in vacuum, shown in blue, which results in $\chi = 46.6$ kHz. After additionally performing the cavity mode spectroscopy (data not shown), we set the LO frequency to work in the rotating frame with $\Delta = 0$.

After calibrating the displacement amplitude, as described in Section II A, we perform an out-and-back experiment [10] to determine the higher order nonlinearities K and χ' . In this experiment, shown in the inset of Supp. Fig. 4(b), we create a coherent state $|\alpha\rangle$ with an average number of $\bar{n} = |\alpha|^2$ photons, wait for some time while it rotates in phase space, and attempt to return it back to the origin with a displacement of variable phase. The optimal return phase for qubit in $|g\rangle$ and $|e\rangle$ is shown in Supp. Fig. 4(b). Performing this experiment for different wait times allows to extract the effective oscillator rotation frequencies $\omega_g(\bar{n})$ and $\omega_e(\bar{n})$. The linear fit of the average rotation frequency $(\omega_g + \omega_e)/2 = \Delta + K\bar{n}$ yields the values of the detuning Δ and Kerr nonlinearity K , as shown in Supp. Fig. 4(c). The linear fit of the relative rotation frequency $\omega_g - \omega_e = \chi + \chi'\bar{n}$ yields the values of the dispersive shift χ and the second-order dispersive shift χ' , as shown in Supp. Fig. 4(d). We find that the value of χ predicted with this method typically agrees with the value obtained via number-resolved spectroscopy to within 1%.

C. Readout and reset

The transmon measurement consists of a readout pulse of duration 700 ns with 40 ns ramp-up and ramp-down. The reflected microwave signal is acquired (after 300 ns delay to account for signal travel time) for the duration of 1400 ns. After acquiring the readout signal, FPGA performs digital signal processing, which consists of demodulation, integration of the signal with a filter function, and thresholding, all of which takes 332 ns. Next, the bits s_0 and s_1 that encode the measurement outcome [“g” if $s_0 = 0$, “e” if $(s_0, s_1) = (1, 0)$, and “f” if $(s_0, s_1) = (1, 1)$] are distributed to all control cards, which takes 100 ns. For a schematic of this measurement process, see Supp. Fig. 5(a). When the readout is used to realize the reset of the auxiliary qubit, additional time is spent on branching on the s_0 and s_1 signals



Supplementary Figure 5. **Auxiliary qubit measurement.** (a) Timing of different components of readout and reset. (b) Logarithmic histogram of the integrated readout signal for different transmon initial states. (c) Markov transition matrix derived from the histogram in (b). It shows the probability of transition from any initial state to any final state during the measurement. (d) Time trace of $|e\rangle$ -state lifetime (color coded). The resonator field amplitude in the steady state is proportional to the DAC amplitude (horizontal axis). The amplitude used for the actual readout corresponds to $\sqrt{\bar{n}} = 0.22$. When a spurious resonance around $\sqrt{\bar{n}} = 0.1$ reappears, the readout fidelity significantly reduces.

to apply appropriate feedback pulses. The branching is done as shown in Supp. Fig. 8(a), taking additional 200 ns to complete the reset. Due to the slow ringdown of the readout photons on a time scale of $1/(\kappa_{r(c)} + \kappa_{r(i)}) \approx 320$ ns, the readout resonator is not fully empty when the feedback pulses are applied, partly limiting their fidelity through measurement-induced dephasing mechanism [11]. This limitation could be addressed in the future by using a strongly coupled readout resonator with photon lifetime on the order of ten nanoseconds [12], or, alternatively, by using an active resonator depletion protocol [13] as was done in the grid-code experiment [14].

To characterize the readout, we perform a two-measurement experiment in which the transmon state is prepared with post-selection on the outcome of the first measurement [15]. The second measurement follows with a 500 ns delay after the first one. Its outcome is histogrammed, as shown in Supp. Fig. 5(b) for $|g\rangle$, $|e\rangle$ and $|f\rangle$ initial states. The SNR of our readout is large enough to not be a dominant cause of the readout infidelity. Instead, the fidelity is limited by state transitions during the finite readout time. Some transitions are expected due to the finite lifetime of the transmon excited states, and the excess is induced by the readout pulse itself [16].

The Markov matrix in Supp. Fig. 5(c) describes the transition probabilities in this characterization experiment. It is obtained by integrating the parts of a histogram on various sides of two thresholds. The diagonal elements of this matrix can be interpreted as readout fidelities of different transmon states, with precision of about $\sim 10^{-3}$ for $|e\rangle$ and $|f\rangle$ states due to possible decay during the 500 ns delay between the two measurements. The readout fidelity of the ground state $\mathcal{F}_r^{(g)} = 0.9997$ is significantly better than that of the excited state $\mathcal{F}_r^{(e)} = 0.9914$ – a crucial feature exploited in our QEC protocol, where the dominant “no error” syndrome is mapped to the g outcome

The readout fidelity of $|g\rangle$ is also more stable over time, see Section IV J. The readout fidelity of $|e\rangle$ fluctuates over time due to fluctuations of $T_1^t(\bar{n})$ (qubit lifetime in the presence of $\bar{n}$ readout photons). A sample drift of $T_1^t(\bar{n})$ is shown in Supp. Fig. 5(d). The exact reason for this effect is still not well understood. It is possible that the dependence $T_1^t(\bar{n})$ comes from drive-induced hybridization of the transmon energy levels [17]. Higher levels are sensitive to offset charge, and thus fluctuations of environmental charges can affect the hybridization strength and lead to fluctuations of $T_1^t(\bar{n})$. Another possible explanation is that fluctuating $T_1^t(\bar{n})$ dependence comes from a spectral overlap of the Stark-shifted qubit frequency with a spurious degree of freedom (not necessarily charged), e.g. a two-level defect, which itself fluctuates [18]. The correlation between the logical qubit performance, the readout infidelity of $|e\rangle$ state, and the fluctuating $T_1^t(\bar{n})$ is further discussed in Section IV J.

D. Conditional displacement

We create an echoed conditional displacement gate $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$ using the approach described in Ref. [10]. As illustrated in Supp. Fig. 6(a), this gate consists of the following steps: (i) the oscillator is displaced out in phase space by large amplitude α ; (ii) the conditional rotation is accumulated during the time interval τ along the arc of a large radius $|\alpha|$ – this is equivalent to accumulation of the conditional displacement in the direction orthogonal to α at an enhanced rate $\chi|\alpha|$; (iii) the oscillator is returned back towards the origin of phase space with displacement of amplitude $-\alpha \cos(\chi\tau/2)$; (iv) the qubit state is flipped with an echo π -pulse; (v) an analogous large displacement sequence is repeated in the symmetrically opposite direction in phase space. Under the dispersive coupling model and in the limit of instantaneous rotation and displacement pulses, this protocol results in a net conditional displacement of amplitude $\beta = -2i\alpha \sin(\chi\tau)$. Due to deviations from this idealized scenario, such as finite pulse durations and higher order Hamiltonian terms, we need to experimentally calibrate the amplitude $\alpha(\beta)$ of the large displacement required to achieve a desired conditional displacement.

Calibration of amplitude. Starting with qubit in $|g\rangle$ (or $|e\rangle$, with similar results) and oscillator in $|0\rangle$, we apply the ECD gate with fixed delay τ in the displaced state and varying amplitude α , and then attempt to undo the effect of the gate and return the oscillator to vacuum with a simple displacement $D(-\beta/2)$. This out-and-back sequence is repeated N times to increase the resolution. At the end of the experiment, the qubit is probed with a selective pulse conditioned on oscillator in $|0\rangle$. The complete sequence is illustrated in Supp. Fig. 6(b), and the experimental data for the gate with delay $\tau = 600$ ns is shown in Supp. Fig. 6(c). From this calibration measurement, we fit the dependence $\alpha(\beta|\tau)$, and we perform this calibration for a set of different wait times τ .

During the optimization of the QEC performance, our RL agent is asked to pick the optimal values of the large displacement amplitude α and of the conditional displacement amplitude β . Therefore, we need to have a calibrated inversion function $\tau(\alpha, \beta)$ that predicts the wait time τ to realize a gate with these parameters. We find that the empirical relation

$$\tau_e(\alpha, \beta) = \beta \left(p_0 + \frac{p_1}{2\alpha} \right) - p_2, \quad (2)$$

with fit parameters $\vec{p} = \{p_0, p_1, p_2\}$, is able to simultaneously fit all ECD calibration datasets, such as the one shown in Supp. Fig. 6(c), sufficiently well to be used with the training of the RL agent. Note that in the idealized model, we would have $\vec{p} = \{0, 1/\chi, 0\}$. The empirical fit results are shown in Supp. Fig. 6(d), where the shaded region indicates the prohibited parameter values, including the limited dynamic range of the DAC that allows $\alpha \in [0, 26]$ given our choice of fixed-duration displacement pulses.

Calibration of qubit phase. As explained in Ref. [10], this experimental implementation of the $\text{ECD}(\beta)$ gate results in additional qubit phase accumulation $\Theta[\beta] = \xi|\beta|^2$, i.e. we implement $\bar{\text{ECD}}(\beta) = \exp(-i\sigma_z \Theta[\beta]/2) \text{ECD}(\beta)$. The amplitude calibration experiment described above is not sensitive to this phase, because the qubit always remains in the eigenstate of σ_z . However, this phase is important when conditional displacements are concatenated, e.g. in the ECD control unitaries, as described in Section III A.

To calibrate this phase $\Theta[\beta]$, we perform the following “cat-and-back” experiment

$$\overline{\text{ECD}}(-\beta) R_x(\pi) \overline{\text{ECD}}(\beta), \quad (3)$$


also shown in Supp. Fig. 6(e), which is ideally equivalent to $\sigma_x \exp(i\Theta[\beta]\sigma_z)$. Starting with a qubit in $|+\rangle$, the final state will satisfy $\langle\sigma_y\rangle = \sin(2\xi|\beta|^2)$ and $\langle\sigma_x\rangle = \cos(2\xi|\beta|^2)$ irrespective of the initial oscillator state.

However, due to decoherence and control imperfections in the ECD implementation, we find that the auxiliary qubit also experiences loss of purity. Under the assumption that the losses of purity during the two conditional displacement gates $\overline{\text{ECD}}(\beta)$ and $\overline{\text{ECD}}(-\beta)$ are uncorrelated and independent of the direction in phase space, we model it as a uniform contraction of the Bloch vector by $\sqrt{1 - p[\beta]}$ per ECD gate, where $p[\beta] = \eta_0 + \eta_2 |\beta|^2 + \eta_4 |\beta|^4$. Hence, we fit the cat-and-back experiment to the following model:

$$\begin{pmatrix} \langle\sigma_x\rangle \\ \langle\sigma_y\rangle \end{pmatrix} = (1 - p[\beta]) \begin{pmatrix} \cos(2\Theta[\beta]) \\ \sin(2\Theta[\beta]) \end{pmatrix}, \quad (4)$$

with fit parameters $\{\eta_0, \eta_2, \eta_4, \xi\}$, of which only ξ is used in the ECD control compilation method, see Section III A.

The results of qubit state tomography together with the fit to the model in Eq. (4) are shown in Supp. Fig. 6(f) for the same ECD gate as in Supp. Fig. 6(c). As explained in Ref. [10], the value of ξ depends on the shape of the phase space trajectory during the ECD gate, and thus we calibrate it independently for every choice of delay time τ .




E. Oscillator error channels

Relaxation and excitation. To measure the oscillator relaxation rate $\gamma_1^c = \gamma_\downarrow^c + \gamma_\uparrow^c$, we first prepare Fock state $|1\rangle$ using a unitary control circuit with 5 layers, see Section III A. After a time delay of varying length, we measure the remaining occupation of $|1\rangle$ and fit it to an exponential decay with time constant $T_1^c = 1/\gamma_1^c$. To measure this occupation, we apply a spectrally selective auxiliary qubit pulse which flips the qubit conditioned on one photon in the oscillator. Monitoring the oscillator over a week-long period, we find the mean and standard deviation of $T_1^c = 606 \pm 10 \mu\text{s}$. As seen from the histogram in Supp. Fig. 7(a), the relative fluctuations of T_1^c are small compared to the relative fluctuations of other error channels in the same time frame. We attribute this stability to the fact that most of the electromagnetic field of this mode resides in the vacuum of the cavity.

To bound the rate of thermal excitation $\gamma_\uparrow^c$, we apply the feedback cooling technique described in Section II F, to the oscillator in its steady state. Since we find no detectable difference in the qubit number-resolved spectroscopy contrast of the zeroth peak after feedback cooling, the resolution of this measurement of $\sim 1\%$ provides a bound on the oscillator excitation rate of $\gamma_\uparrow^c < 1/(60 \text{ ms})$. This rate is negligible compared to all other rates in the system and is ignored in the rest of the analysis.

Dephasing. To measure the rate of dephasing γ_2^c within the $\{|0\rangle, |1\rangle\}$ manifold, we prepare a superposition $|0\rangle + |1\rangle$ using the Y90 gate realized with a unitary control circuit with 8 layers, see Supp. Table II. After a time delay of varying length we apply the Y90 gate again and measure the occupation of $|0\rangle$. In the reference frame of the LO, the oscillator state rotates with angular frequency $\chi/2$ during the time delay, which results in Ramsey oscillations modulating the exponential decay with decay time constant $T_2^c = 1/\gamma_2^c$. We adjust the sampling rate to make the oscillations appear slow. We find the one-week mean and standard deviation of $T_2^c = 980 \pm 30 \mu\text{s}$.

One possible source of oscillator dephasing is stochastic rotations acquired due to dispersive coupling with the transmon combined with transmon stochastic excitation and relaxation events [19]. The dephasing rate due to this effect was predicted to be $\gamma_\varphi^{c,t} \approx n_{\text{th}}^t \gamma_\downarrow^t$ in the limit $\chi \gg \gamma_1^t$ and $\gamma_\downarrow^t \gg \gamma_\uparrow^t$, where n_{th}^t is the steady-state population of $|e\rangle$. In our system, the correlation between $\gamma_\varphi^c = \gamma_2^c - \gamma_1^c/2$ and $\gamma_\varphi^{c,t}$ is difficult to measure because these rates are small and their estimators are subject to strong relative fluctuations. By comparing the medians of their marginal distributions, $\bar{\gamma}_\varphi^{c,t} = 1/(6.5 \text{ ms})$ and $\bar{\gamma}_\varphi^c = 1/(5.1 \text{ ms})$, shown in Supp. Fig. 7(b), we find the remaining unexplained



Supplementary Figure 7. **Fluctuating error channels.** (a) Histogram of T_1 and T_2 times of the transmon and the oscillator, and logical lifetimes of error-corrected grid states. The histogram is derived from a week-long scan described in Chapter IV J. (b) Oscillator pure dephasing time extracted from the measured oscillator parameters and predicted from the dispersive coupling model.

contribution to dephasing at a rate $\gamma_\varphi^c = 1/(24\text{ ms})$ whose source is not yet identified. It is plausibly related to second-order excitations from $|e\rangle$ to $|f\rangle$ [2].

F. Active oscillator cooling

Given the long relaxation time $\bar{T}_1^c = 606\text{ }\mu\text{s}$ of our oscillator, passive cooling that relies on the natural interaction with the cold environment is impractically long. For example, starting with a Fock state $|1\rangle$, it would take approximately $4.6\bar{T}_1^c = 2.8\text{ ms}$ to reduce the average population to 0.01 photons. In practice, since we work with the grid states, the required cooling time is even longer. Therefore, the goal of our active cooling routine is to reduce the experimental duty cycle time and also to remove any residual thermal population. We achieve these goals via a two-step procedure which consists of an engineered dissipative pre-cooling and subsequent feedback cooling.

Dissipative pre-cooling. We introduce a novel oscillator cooling method based on the conditional displacements, auxiliary qubit rotations, and auxiliary qubit resets. This protocol can also be realized in trapped ions, as was hinted in Ref. [20].

To derive this protocol, we apply the same dissipation engineering framework [21] as used in Ref. [22] to derive the SBS stabilization of the GKP manifold. The dissipator $\gamma\mathcal{D}[a]$ can be approximated with a sequence of discrete entangling interactions $U(t)$ between the auxiliary qubit and the oscillator, and auxiliary qubit resets. For $\gamma\mathcal{D}[a]$, the interaction should be of the form $U(t) = \exp[-i\sqrt{\gamma t}(a\sigma_+ + a^\dagger\sigma_-)]$, where the constraint $\langle a^\dagger a \rangle \gamma t \ll 1$ controls the validity of this discrete approximation. To further approximate this unitary as a multi-layer circuit with gates from our gate set, we perform the first order Trotter decomposition:

$$U = \exp\left(-i\sqrt{\frac{\gamma t}{2}}(x\sigma_x + p\sigma_y)\right) \quad (5)$$

$$= \exp\left(-i\sqrt{\frac{\gamma t}{2}}x\sigma_x\right) \exp\left(-i\sqrt{\frac{\gamma t}{2}}p\sigma_y\right) + O(\gamma t) \quad (6)$$


$$\approx R_y^\dagger(\pi/2) \text{ECD}(-i\varepsilon) R_x^\dagger(\pi/2) \text{ECD}(\varepsilon) R_y(\pi/2) R_z(\pi/2), \quad (7)$$

where we defined the “trimming amplitude” $\varepsilon = \sqrt{\gamma t}$. Furthermore, since the auxiliary qubit is assumed to always start in $|g\rangle$, we can omit the first gate $R_z(\pi/2)$. The resulting unitary part of the dissipative cooling circuit is:

$$R_y^\dagger(\pi/2) \text{ECD}(-i\varepsilon) R_x^\dagger(\pi/2) \text{ECD}(\varepsilon) R_y(\pi/2), \quad (8)$$

also summarized in Supp. Table II. To achieve uniform cooling in all directions in phase space, the orientation of the ECD gates needs to cycle between position and momentum quadratures. A single cycle, including the pulse sequence in (8), auxiliary qubit reset, and subsequent virtual rotation gate on the FPGA, has a duration of $3.38\text{ }\mu\text{s}$.

To demonstrate the performance of this cooling protocol, we start with a $|+Z\rangle$ grid state with $\Delta = 0.3$ and apply varying number of cooling cycles, monitoring the population of $|0\rangle$ with a selective qubit pulse. As seen in Supp. Fig. 8(c), dissipative cooling allows the state to shrink towards vacuum significantly faster than passive cooling.



Supplementary Figure 8. **System cooling.** (a) Transmon reset subroutine with measurement-based three-state feedback. (b) Oscillator feedback cooling subroutine adapted from Ref. [5]. (c) Demonstration of dissipative cooling of the oscillator starting from GKP $|+Z\rangle$ state with $\Delta = 0.3$. A single cooling cycle consists of a pulse sequence in Eq. (8), auxiliary qubit reset as in (a), and virtual rotation gate to the orthogonal quadrature for the next cycle. The duration of a single such cycle is $3.38 \mu\text{s}$. The case $\epsilon = 0$ is equivalent to passive cooling. Dashed lines represent the contrast of the zeroth photon number peak in qubit spectroscopy after passive cooling of 5 ms and after feedback cooling with $Y = 3$.

With $\epsilon = 0.4$, the cooling rate is 20 times faster than energy relaxation time of the oscillator. For small $\epsilon \leq 0.3$ the steady-state thermal occupation after dissipative cooling is similar to passive cooling of this state of duration 5 ms. Larger ϵ allows for faster cooling, but at the expense of significant residual thermal occupation.

Feedback cooling. To remove the residual thermal photons, we further apply the feedback cooling protocol introduced in Ref. [5] and shown in Supp. Fig. 8(b). With the help of a selective qubit pulse conditioned on $|0\rangle$ and qubit measurement, the protocol repetitively asks the question “Is the oscillator in vacuum?” and terminates only when it receives Y consecutive “yes” answers. It would be inefficient to run this feedback protocol starting with an arbitrary initial oscillator state, since the probability p_y of obtaining “yes” can be very small. The dissipative pre-cooling quickly boosts this probability to a level essentially limited by the fidelity of the selective qubit pulse, and thereby decreases the run time of the subsequent feedback cooling step. The run time of feedback cooling is non-deterministic, but the expected number of rounds in a model with constant p_y is

$$\overline{N}_{fc}(p_y, Y) = \frac{p_y^{-Y} - 1}{1 - p_y} \quad (9)$$

Our final routine, called “active cooling” throughout this work, consists of 25 cycles of dissipative pre-cooling (50 cycles, if counting each quadrature individually) with $\epsilon = 0.4$ followed by the feedback cooling with $Y = 3$. We estimate that with $p_y = 0.87$, achieved after the pre-cooling, the expected run time of the whole routine is approximately $50 \times 3.38 \mu\text{s} + \overline{N}_{fc}(0.87, 3) \times 25 \mu\text{s} = 270 \mu\text{s}$, which in our system corresponds to $0.45 T_1^c$ (and could potentially be reduced further).

From the contrast of the zeroth photon number peak in the qubit spectroscopy [dashed lines in Supp. Fig. 8(c)], we see that passive cooling of duration 5 ms starting from the $|+Z\rangle$ grid state still leaves a residual thermal population larger than what our protocol achieves in a much shorter time. However, when active cooling is applied to an oscillator in its steady state (nominally, vacuum) we find no resolvable improvement of the spectroscopy contrast, which leads us to conclude that the residual thermal population after active cooling is at the sub-percent level where it cannot be resolved with our spectroscopy. This observation is used in Section II E to derive an upper bound on the oscillator thermal excitation rate.

III. QUANTUM CONTROL OPTIMIZATION

A. Model-based optimization of control circuits

Circuit decomposition. Our control gate set consists of two parametrized gates: (i) echoed conditional displacement of the oscillator $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$, where $D(\alpha) = \exp[\alpha a^\dagger - \alpha^* a]$ is the displacement operator, and (ii) rotation of the qubit $R(\varphi, \theta) = \exp[-i(\theta/2)(\sigma_x \cos \varphi + \sigma_y \sin \varphi)]$. Recently, it was shown that this gate set is well suited for the universal control of an oscillator with weak dispersive coupling to a qubit [10]. Most unitary operations in our experiment are decomposed as parametrized multilayer circuits of the form

$$\text{circuit}(\beta, \varphi, \theta) = \underbrace{\text{ECD}(\beta_T)R(\varphi_T, \theta_T)}_{\text{layer } T} \cdots \underbrace{\text{ECD}(\beta_1)R(\varphi_1, \theta_1)}_{\text{layer } 1}, \quad (10)$$

where $\beta \in \mathbb{C}^T$ is a vector of conditional displacement amplitudes, and $\varphi, \theta \in \mathbb{R}^T$ are vectors of qubit rotation phases and angles respectively. For example, we utilize this decomposition as part of the following operations:

- Dissipative cooling of the oscillator, see Section II F.
- Preparation of the GKP states, see Section IV H.
- Small-Big-Small protocol, see Section IV C.
- Preparation of the Fock state $|1\rangle$, see Section II E.
- Y90 gate on Fock $\{|0\rangle, |1\rangle\}$ encoding, see Section II E.

Circuit optimization. A circuit optimization method for this gate set was developed in Ref. [10]. Here, we present a simplified modular framework based on the Keras library [23], which allows to optimize circuit parameters in a manner similar to training of the neural networks. The parametrized control circuits (10) are created as instances of the `tf.keras.Sequential` class which is commonly used to concatenate multiple neural network layers. Here, we instead use custom layers that represent the parametrized gates $\text{ECD}(\beta)$ and $R(\varphi, \theta)$ as subclasses of `tf.keras.layers.Layer`. This allows us to exploit flexible and user-friendly application-programming interface of the Keras library to optimize the circuit parameters and automatically monitor various aspects of the optimization progress. To illustrate the accessibility of such an approach, in Supp. Fig. 9 we provide an example code for optimization of the Y90 gate on the $\{|0\rangle, |1\rangle\}$ qubit. Complete code with dependencies and further examples is available in Ref. [24]. Such optimization, which is performed for a batch of $B = 300$ circuit candidates in parallel on a graphics processing unit (GPU), takes about 10 minutes to finish. In Supp. Table II, we list circuit parameters for some of the control operations in our experiment. Curiously, some of the numerically optimized parameter values are clearly interpretable, e.g. in GKP state preparation circuit the rotations at steps $t = 1, 6, 10, 11$ seem to be by an angle π . Detailed inspection of these circuits can lead to improved analytic constructions, which is left for future research.

	Dissipative cooling			SBS protocol			$ +Z\rangle$ grid state prep.			Y90 gate for $\{ 0\rangle, 1\rangle\}$ qubit		
t	β	φ	θ	β	φ	θ	β	φ	θ	β	φ	θ
1	+0.4	$+\pi/2$	$+\pi/2$	+0.2i	$+\pi/2$	$+\pi/2$	+0.52 + 2.54i	-1.28	+1.57	+0.64 + 0.11i	-1.06	+1.58
2	-0.4i	0	$-\pi/2$	$+\sqrt{2}\pi$	0	$-\pi/2$	-0.83 - 0.36i	+2.85	-2.76	-0.15 - 1.00i	+2.64	-1.44
3	0	$+\pi/2$	$-\pi/2$	+0.2i	0	$+\pi/2$	-0.36 + 0.85i	+0.29	+0.55	+1.02 + 0.05i	+0.58	-1.97
4				0	$+\pi/2$	$-\pi/2$	-0.86 + 1.61i	-0.29	+1.43	+1.55 + 1.02i	-1.84	-1.55
5							-2.16 + 0.12i	+0.29	+0.92	+0.34 + 1.06i	+2.74	+0.26
6							-0.09 + 1.73i	+2.85	-1.56	-0.26 - 0.92i	-0.01	-1.25
7							+2.05 + 0.73i	+0.29	+1.08	-0.61 - 0.05i	-2.91	-1.75
8							+0.22 - 0.66i	-0.29	-2.71	+0.02 + 0.05i	-1.79	-1.67
9							-0.08 - 1.56i	+0.29	+2.06			
10							+0.19 + 0.04i	+2.85	+1.60			
11							0	+1.86	+1.57			

Supplementary Table II. **Circuit parameters.** Parameters for dissipative cooling and SBS protocol are created based on the models described in Section II F and Ref. [22] respectively. Parameters for $|+Z\rangle$ grid state preparation and Fock Y90 gate on $\{|0\rangle, |1\rangle\}$ qubit are numerically optimized with Keras.

```

import tensorflow as tf
from quantum_control.layers import QubitRotation, ConditionalDisplacement
from utils import basis, Kronecker_product, batch_dot, normalize

N = 100 # oscillator truncation in Fock basis
T = 8 # circuit depth
B = 300 # batch size

# build the circuit as Keras model
circuit = tf.keras.Sequential(name='fockY90')
for i in range(T):
    echo_pulse = True if i < T-1 else False
    circuit.add(QubitRotation(N, [1,B]))
    circuit.add(ConditionalDisplacement(N, [1,B], echo_pulse=echo_pulse))

# Create cardinal points on the Bloch sphere
# Hilbert space convention is qubit x oscillator
states = {'+Z': Kronecker_product([basis(0,2), basis(0,N)]),
          '-Z': Kronecker_product([basis(0,2), basis(1,N)]),
          '+X': normalize(states['+Z'] + states['-Z'])[0],
          '-X': normalize(states['+Z'] - states['-Z'])[0]}

# specify the action of Y90 gate
gate_map = {'+X': '-Z', '-X': '+Z'}

# define input and output states
inputs = tf.stack([states[i] for (i,f) in gate_map.items()])
targets = tf.stack([states[f] for (i,f) in gate_map.items()])


# define the loss function and the optimizer
def loss(state, target):
    return -tf.math.reduce_sum(tf.math.real(batch_dot(state, target)))

optimizer = tf.keras.optimizers.Adam(learning_rate=1e-3)

# compile this Keras model and run the optimizer
circuit.compile(optimizer=optimizer, loss=loss)
circuit.fit(x=inputs, y=targets, epochs=10000)

```

Supplementary Figure 9. **Circuit optimization.** Example of a Python script for optimization of the circuit parameters for Y90 gate on the $\{|0\rangle, |1\rangle\}$ qubit. Gates are represented as custom Keras layers, and the circuit is compiled as sequential model. Optimization utilizes TensorFlow backend for automatic differentiation of the model.



Supplementary Figure 10. **Waveform for $|+Z\rangle$ grid state preparation.** The parametrized control circuit is decomposed into primitive gates: qubit rotations, oscillator displacements, and conditional rotations. The waveform is compiled from this sequence of gates using experimental calibrations. Each qubit rotation and oscillator displacement is replaced with a corresponding Gaussian pulse, and the conditional rotation is replaced with a delay of certain length during which the system freely evolves under the dispersive coupling Hamiltonian.

Pulse compilation. Having obtained the circuit parameters, we compile the waveforms to be played on the qubit and oscillator control lines. Such compilation requires prior calibration of the rotation $R(\varphi, \theta)$ gate, described in Section II A, and the ECD(β) gate, described in Section II D.

As explained in Ref. [10] and in Section II D, our experimental implementation of the ECD(β) gate results in additional qubit phase accumulation $\Theta[\beta] \propto |\beta|^2$, i.e. we implement $\text{ECD}(\beta) = \exp(-i\sigma_z \Theta[\beta]/2) \text{ECD}(\beta)$. We use

the experimental calibration of this phase to adjust the numerically optimized vector φ according to the rule

$$\varphi_t \leftarrow \varphi_t - \sum_{\tau=1}^{t-1} (-1)^{t-\tau} \Theta[\beta_\tau], \quad t > 1. \quad (11)$$

In addition, in many cases of interest the auxiliary qubit at the end of the circuit returns to $|g\rangle$ and disentangles from the oscillator. In such cases, the last conditional displacement $\text{ECD}(\beta_T)$ can be realized as a simple displacement $D(\beta_T/2)$. We use this simplification in state preparation circuits and in the SBS protocol.

In Supp. Fig. 10, we show an example waveform for unitary preparation of the $|+Z\rangle$ grid state using a parametrized circuit with $T = 11$ layers. Each ECD gate is decomposed via large displacements and conditional rotations. For clarity, in this example all conditional rotations are implemented with a constant wait time $\tau = 200$ ns; hence, the whole compiled waveform has a duration of $6.4 \mu\text{s}$. Faster implementations are possible if the wait time is adapted to the magnitude of the conditional displacement, as described in Section IID. For example, in our system the conditional displacement of amplitude $|\beta| < 0.5$ could, in principle, be implemented with zero wait time, see Supp. Fig. 6(d).


B. Model-free reinforcement learning for QEC

While most quantum operations in our experiment are optimized with a model-based approach described above, for quantum error correction we deploy a more powerful framework of model-free optimization. We use a reinforcement learning algorithm called proximal policy optimization (PPO) [25, 26]. For a detailed description of this algorithm in the context of quantum control we refer to Ref. [8]; here, we only provide a basic high-level picture. The complete training loop of our experiment is illustrated in Supp. Fig. 11; it is structured as follows:

Step 1. On training epoch t , neural network produces a probability distribution $\mathcal{N}(\vec{\mu}_t, \vec{\sigma}_t)$, where $\vec{\mu}_t = \vec{\mu}(\theta_t)$, $\vec{\sigma}_t = \vec{\sigma}(\theta_t)$, and θ_t summarizes the values of all weights and biases of the neural network in the current epoch.

Step 2. We sample a batch of $B = 10$ parameter vectors from this distribution. They correspond to different QEC circuit candidates that should be evaluated in experiment. The neural network and sampling are implemented on NVIDIA 2080Ti graphics processing unit (GPU) in a separate computer. The sampled vectors are sent to the control computer via a local area network with negligible communication time.

Step 3. Based on these parameter vectors, we compile QEC circuit candidates, translated into FPGA instructions and DAC waveforms. All circuit candidates follow the same program execution flow, but the control waveforms and the content of FPGA registers is different for every candidate. The FPGA is reset and its wave memory is updated. This time-consuming step is the bottleneck of the training loop.



Supplementary Figure 11. **Reinforcement learning.** (a) Experimental training loop. (b) Training time budget per epoch.

Step 4. Each candidate is evaluated in experiment. To this end, we initialize logical Pauli eigenstates $|+Z\rangle$ and $|+X\rangle$, run the QEC for $T = 160$ cycles, and then perform one-bit phase estimation of the corresponding logical Pauli operators. To suppress the sampling noise, we repeat this $N_{\text{avg}} = 150$ times per Pauli and per circuit candidate. In total, one epoch of training consists of $N_{\text{tot}} = 2BN_{\text{avg}} = 3000$ experimental shots.

Step 5. To produce the reward, we treat the measurement of a Pauli operator after T cycles as a proxy for logical lifetime. While averaging the measurement outcomes, we mask the experimental shots that started with incorrect state initialization, as flagged by a verification transmon measurement after the state initialization.

Step 6. Once the rewards are available, PPO algorithm updates the neural network parameters $\theta_t \rightarrow \theta_{t+1}$ for the next epoch. The gradients of these parameters are computed with automatic differentiation via back-propagation. After updating the neural network, the new training epoch begins.


The time budget of this training is shown in Supp. Fig. 11. All steps outlined above amount to 15.6 s per epoch. In the current implementation, the major bottleneck is Python-to-FPGA transition (step 3). Because of this, the implementation is less optimal in terms of sample efficiency than the proposal in Ref. [8]. The optimal approach would be to spend the total sample budget per epoch to evaluate more circuit candidates with minimal accuracy, instead of evaluation only a few candidates with high accuracy (achieved through averaging). In other words, based on the results of Ref. [8], we expect that a training with $(B, N_{\text{avg}}) = (1000, 1)$ would require fewer experimental shots to reach a given performance level than a training with $(B, N_{\text{avg}}) = (10, 100)$. However, considering the total run time of the training, we had to compromise between bare sample efficiency (number of shots) and the overhead in step 3 of the pipeline. The overhead is independent of N_{avg} but increases with B , and due to a limited FPGA instruction sequence length we can only evaluate $B \leq 10$ candidates per compilation. After paying the compilation overhead in step 3, a certain amount of averaging comes essentially for free and does not considerably affect the run time, hence the choices made here.

In Section IV D, we describe the QEC circuit parametrization, show the evolution of parameter values during the course of training, and provide interpretation of the observed trends.

IV. QUANTUM ERROR CORRECTION OF THE GRID CODE

A. Brief introduction to grid code

Qubit-register stabilizer codes are based on the group of Pauli operators; consider instead a stabilizer code based on the group of oscillator displacement operators. By definition, the $+1$ eigenstates of a displacement operator $D(\alpha)$ are displacement-invariant in phase space along the direction of α with a period $|\alpha|$. Having two code stabilizers $S_0^X = D(\alpha_X)$ and $S_0^Z = D(\alpha_Z)$ imposes displacement invariance along two non-equivalent directions, which means that all codewords are grids in phase space with a unit cell defined by $\{\alpha_X, \alpha_Z\}$. The requirement of commutativity



Supplementary Figure 12. **Wigner functions (numerical) of grid states.**

of S_0^x and S_0^z imposes a constraint

$$\alpha_X^* \alpha_Z - \alpha_X \alpha_Z^* = 2l_s^2 n i, \quad n \in \mathbb{Z}, \quad (12)$$

where $l_s = \sqrt{2\pi}$. Here, we consider encoding a single logical qubit into an oscillator, which corresponds to $n = 1$. By parametrizing the complex-valued displacement amplitudes as $\alpha_X = l_s[M_{22} - iM_{12}]$ and $\alpha_Z = l_s[iM_{11} - M_{21}]$, we obtain a grid code with the following stabilizers:

$$S_0^z = D(l_s[iM_{11} - M_{21}]), \quad (13)$$

$$S_0^x = D(l_s[M_{22} - iM_{12}]). \quad (14)$$

From the constraint (12) we derive a single requirement that a real matrix $M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}$ has a determinant $\det M = 1$. This matrix defines the structure of the grid in phase space. Here, we only consider the square grid code, which is obtained with $M = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. The hexagonal code with $M = \sqrt{\frac{2}{3}} \begin{bmatrix} 1 & 1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}$ was realized in Ref. [14].

The Pauli operators of the logical qubit are defined as

$$X_L = \sqrt{S_0^x} = D(l_s[M_{22} - iM_{12}]/2), \quad (15)$$

$$Z_L = \sqrt{S_0^z} = D(l_s[iM_{11} - M_{21}]/2). \quad (16)$$

They satisfy the standard algebraic properties $X_L^2 = I$, $Z_L^2 = I$, and $X_L Z_L = -Z_L X_L$, inside the code space. Using the identity $Y_L = -iZ_L X_L$, we find the third Pauli operator $Y_L = -iD(l_s[iM_{11} - iM_{12} + M_{22} - M_{21}]/2)$.

The eigenstates of Pauli Z_L of the ideal grid code are shown in Supp. Fig. 12(a). Finite-energy code families can be obtained by regularizing the ideal code through application of an envelope operator [22, 27], with a common choice being $N_\Delta = \exp(-\Delta^2 n)$. We show several members of this code family in Supp. Fig. 12(b-d). Note that such a regularization leads to non-orthogonal states, with fidelity loss due to the finite state overlap that scales as $\exp(-\pi/4\Delta^2)$ [see Eq. S28 in Ref. [22]], which is negligible for our choice of $\Delta = 0.34$.

B. Small-Big-Small (SBS) protocol

Here, we describe the SBS protocol, first proposed in Ref. [20, 22] from a new angle. The full QEC circuit in this protocol is shown in Supp. Fig. 13(a) with nominal parameter values listed in Supp. Table II; it implements a channel $\mathcal{R}_\Delta(\rho) = (\mathcal{R}_\Delta^x \circ \mathcal{R}_\Delta^z)(\rho)$. Let $(K_g^{x/z}, K_e^{x/z})$ denote the Kraus operators of the constituent rank-2 channels $\mathcal{R}_\Delta^{x/z}$ (we omit the Δ subscript from the Kraus operators for simplicity). These operators read:


$$K_g^x = \cos(\sqrt{\pi}p) \cos(\sqrt{\pi}\Delta^2 x) + \sin(\pi\Delta^2/2) \cos(\sqrt{\pi}p), \quad (17)$$

$$K_e^x = -\cos(\pi\Delta^2/2) \sin(\sqrt{\pi}p) + i \cos(\sqrt{\pi}p) \sin(\sqrt{\pi}\Delta^2 x), \quad (18)$$

where $x = (a + a^\dagger)/\sqrt{2}$ and $p = i(a^\dagger - a)/\sqrt{2}$, and (K_g^z, K_e^z) are obtained with a substitution $(x, p) \rightarrow (-p, x)$. Then, the Kraus operators of a composite rank-4 channel are:

$$K_{gg} = K_g^z K_g^x, \quad K_{ge} = K_g^z K_e^x, \quad K_{eg} = K_e^z K_g^x, \quad K_{ee} = K_e^z K_e^x. \quad (19)$$

For $\Delta = 0.34$, these Kraus operators are shown as matrices in the truncated eigenbasis of $K_{gg}^\dagger K_{gg}$ in Supp. Fig. 13(b). This eigenbasis splits into pairs of states $\mathcal{C}_i = \{|0_i^L\rangle, |1_i^L\rangle\}$, $i \in \mathbb{N}$, that define orthogonal replicas of the logical subspace $\mathcal{C}_0$ generated by the errors. We show the Wigner functions of the projectors Π_0 , Π_1 , and Π_2 onto the first three subspaces in Supp. Fig. 13(d). Note that $\Pi_1 \approx a \Pi_0 a^\dagger$ and $\Pi_2 \approx a^\dagger \Pi_0 a$, hence the errors in the first level of hierarchy resemble photon loss (a) and gain ($a^\dagger$) errors. While a and $a^\dagger$ only approximately satisfy the Knill-Laflamme conditions [28] for the finite-energy grid code, the actual error operators that define the subspaces $\mathcal{C}_1$ and $\mathcal{C}_2$ satisfy these conditions exactly (since the eigenspaces of a Hermitian operator $K_{gg}^\dagger K_{gg}$ are orthogonal). Similarly, by inspecting the Wigner functions of the projectors onto higher subspaces, we find that the second level of error hierarchy resembles a^2 , $a^\dagger a$ and $a^{\dagger 2}$. The number of error subspaces in each level is given by the number of unique combinations of a and $a^\dagger$: two subspaces (a and $a^\dagger$) in the first level, and three subspaces (a^2 , $a^\dagger a$ and $a^{\dagger 2}$) in the second level, leading to the blocks of size 4×4 and 6×6 in the Kraus matrices in Supp. Fig. 13(b). Further understanding the structure of the error hierarchy is the subject of ongoing research.



Supplementary Figure 13. **SBS protocol.** (a) Circuit structure of one QEC cycle. (b) Kraus operators of the QEC cycle with $\Delta = 0.34$, written in the eigenbasis of $K_{gg}^\dagger K_{gg}$. This eigenbasis splits into pairs of states $C_i = \{|0_i^L\rangle, |1_i^L\rangle\}$, $i \in \mathbb{N}$, that define replicas of the logical subspace C_0 . Color encodes the absolute value of the matrix elements. (c) Flow diagram corresponding to each Kraus operator. Circles represent error spaces, and arrows show the most relevant matrix elements. The dynamics within the subspaces is discarded in this representation. (d) Numerical Wigner functions of the projectors onto the subspaces C_0 , C_1 , and C_2 . Comparison to the subspaces generated from the code space by the errors a and $a^\dagger$ reveals that the errors in the first level of hierarchy approximately correspond to a and $a^\dagger$. (e) Quantum state trajectories with errors and QEC. The state is represented in the same basis as in (b); color encodes the absolute value of the state components in this basis. Red dotted lines are guides to the eye that separate the error subspaces. The occurrence of errors is indicated with red arrows at the top. The time axis is measured in QEC cycles. The Kraus operators are applied between the time steps, and the syndrome string encodes which Kraus operator was applied on every step. The state transfer fidelity, shown at the bottom, measures the squared overlap of the final and initial state vectors.

Unlike in the standard stabilizer formalism of QEC [29], Kraus operators here do not correspond to a projection of a state onto a single error subspace and its subsequent transfer to the code space. Instead, the transfer here is realized gradually, following an error hierarchy imposed by the QEC circuit. To clarify the action of the Kraus operators, their reduced representation using directional flow of a quantum state between error subspaces is shown in Supp. Fig. 13(c) [this representation ignores the dynamics within each subspace]. We now briefly discuss the interpretation of the processes corresponding to each of the gg , ge , eg , ee outcomes of a QEC cycle. Outcome gg heralds a process in which the state has remained in the same subspace. The probability of emitting gg from within the code space is nearly 1. This property is exploited in Section IV E to extract the expectation value of the code projector $\langle \Pi_0 \rangle$ from the statistics of long strings of the $gg/gg/\dots$ type. Both ge and eg outcomes herald the process in which the quantum state was transferred one level down the error hierarchy. Strings like $eg/eg/eg/\dots$ therefore correspond to processes in which the state directionally hops level by level towards the code space. Finally, the ee outcome heralds a transfer two levels down the error hierarchy.

Besides the transfer between the error spaces, the Kraus operators apply a deterministic logical flip: X_L in the $\mathcal{R}_\Delta^X$ cycles, and Z_L in the $\mathcal{R}_\Delta^Z$ cycles. This flip is visible in the off-diagonal structure of the sub-blocks in the Kraus matrices, see Supp. Fig. 13(b). For example, the lower right 2×2 block in K_{gg} represents the code subspace, and the off-diagonal structure represents the combined effect of $X_L Z_L = -iY_L$ on the codewords. Due to this effect, the lifetime of $+1$ and -1 logical Pauli eigenstates in our QEC protocol are exactly equal. We track the Pauli frame in software, and undo its change in the data reported in Fig. 3 of the main text.

To demonstrate how the errors are corrected by this QEC scheme, we show several examples of quantum state trajectories in Supp. Fig. 13(e). In the first trajectory, the state is initialized as one of the logical basis states, and then evolved for several QEC cycles without any errors. The Pauli frame switching is apparent here from the oscillating pattern within the code space (in this picture, the phase information is not shown, but the QEC process also protects the phase of the logical qubit). In the second trajectory, an error $a^\dagger$ was applied to the state prior to QEC, and then it was almost perfectly corrected, accompanied by the emission of $eg/gg/\dots$ syndrome string. In the third trajectory, this error was instead corrected during the third QEC cycle, and the quantum state spent extra time in the error space $\mathcal{C}_2$. This example explicitly demonstrates that the Pauli frame update is applied correctly irrespective of the subspace, hence Pauli gates done in this manner are transversal. The subsequent trajectories demonstrate that even higher-order errors, such as $a^{\dagger 2}$ or $a^{\dagger 4}$, can be corrected with high fidelity. Moreover, as seen in the fifth trajectory, the state can be recovered even if additional errors happen while the previous errors have not yet been fully corrected. The latter example highlights that the “slowness” of the low-rank error-correction dissipation is not a problem, as long as the error rate is sufficiently small compared to the correction rate.

A few remarks with regards to the simplified interpretation of the QEC process in the main text are in order: (i) The correct interpretation of the action of a QEC cycle requires considering pairs of outcomes, like ge , instead of isolated outcomes, like g or e . We adopted the latter approach in the main text for simplicity of exposition. (ii) The gg outcome does not herald the projection onto the code space, as mentioned in the main text, but rather a process in which “no error was corrected”. Conditioned on the state residing in the code subspace, this outcome will be emitted with probability nearly 1. However, if the state is in one of the error spaces this outcome can still occur with smaller probability starting from about 0.47 at the lowest level in the error hierarchy and reducing for higher levels. (iii) When one of the outcome eg , ge or ee is obtained, there is a small chance that the QEC process has added an error, leading to a random walk among the error spaces that is heavily biased towards the code space.

C. QEC cycle: implementation details

In this section, we detail the implementation of a QEC cycle, whose schematic is shown in Supp. Fig. 13. The various datasets in this work were taken with several different versions of the QEC circuit. All these versions have the same overall structure, but different parameter values obtained from re-training after the system drift has appreciably affected the logical performance (this happens on a time scale of 1-2 weeks, see Section IV J). Below, the quoted durations of various components of a QEC cycle refer to the circuit version that we used to collect the system lifetimes dataset and that achieved the highest reported QEC gain.

SBS unitary. We refer to the unitary part of the circuit U_θ prior to auxiliary qubit measurement as the “SBS unitary” since it is based on the ansatz from Ref. [22]. The SBS unitary is compiled as a four-layer parametrized circuit with nominal parameters shown in Supp. Table II, and is further translated into the pulse sequence with the method described in Section III A. The last circuit layer does not contain an ECD gate, and instead only contains a qubit rotation and oscillator displacement. Since the qubit is reset after the SBS, the function of the latter rotation is to choose the “reset axis”, which can be an arbitrary axis on the qubit Bloch sphere.

As shown in Ref. [22], without any special asymmetries between $|g\rangle$ and $|e\rangle$ it would not matter along which axis the auxiliary qubit reset is done – all choices result in the same completely positive trace-preserving map after averaging

over the measurement outcomes. However, in practice the asymmetry comes from the auxiliary qubit relaxation channel that degrades the readout fidelity of the $|e\rangle$ state. Hence, it is advantageous to choose the reset axis that preferentially returns the $|g\rangle$ outcome. The parameter sequence for the SBS unitary in Supp. Table II takes this choice into account. The choice of reset axis also results in a different unraveling of state trajectories and different Kraus operators. The choice made here enabled the interpretation of e outcomes as syndromes that signal the occurrence and correction of errors, which is utilized in the post-selection experiments described in Section IV G. This is in contrast with Ref. [14], where g and e outcomes are interpreted as left or right displacements of the grid.

The duration of the SBS unitary is not fixed, because its constituent ECD gates can be implemented with different choices of the speed enhancement factor α (amplitude of the intermediate displacement). Since α is included in the action space of the RL agent, all circuit candidates during the training have different durations of the SBS unitary (we will soon comment on how this affects the reward comparison among them). In the final circuit that achieved the highest reported QEC gain, the duration of SBS unitary is $t_{\text{SBS}} = 1546 \text{ ns}$.

Auxiliary qubit reset. In principle, error correction with the SBS protocol could be fully autonomous (without a classical feedback loop) as was envisioned in the proposal [22] and realized in a trapped ion system [20]. The autonomous scheme has an advantage of significantly simplifying the demands on the classical co-processor (in our case, the FPGA). Moreover, there exist various dissipative reset protocols for the transmon [30–32]. However, the disadvantage of a fully autonomous implementation in our system is that it is not able to compensate for a spurious rotation of the oscillator due to the always-on dispersive coupling with the transmon. The back-action of discarding the transmon state during the reset is the dephasing of the oscillator – a particularly harmful error channel for the GKP code [22]. Partly because of this reason, we chose to implement transmon reset through measurement and classical feedback, as described in Section II C, with the total duration of transmon reset subroutine of $t_{\text{reset}} = 2332 \text{ ns}$.

Virtual rotation. Due to the always-on dispersive coupling, the oscillator acquires a spurious rotation during the auxiliary qubit readout time. In experiment [14], a simple echo sequence was used to cancel this rotation. With such an approach, transmon spends half of the time in $|g\rangle$ and half in $|e\rangle$ regardless of the actual syndrome measurement outcome, which is detrimental to the code due to additional error sources associated with the $|e\rangle$ state. Here, we instead chose the reset axis which results in 0.9 probability of detecting $|g\rangle$. Therefore, the ability to compensate for the spurious oscillator rotation without echoing the state back to $|e\rangle$ is crucial to maintain this advantage.

We achieve this by dynamically tracking the oscillator phase that stochastically changes due to random transmon measurement outcomes, and compensating for it with a virtual counter-rotation. The spurious oscillator rotation angle accumulates during the reset time t_{reset} , during the time t_{VR} that it takes to execute the virtual rotation on the FPGA, and during the idle time t_{idle} when the transmon is nominally in $|g\rangle$ (the latter will be explained shortly). Therefore, in the idealistic dispersive coupling model, the oscillator would rotate by $\vartheta_g = \chi(t_{\text{VR}} + t_{\text{idle}} + t_{\text{reset}})/2$ if the transmon is found in $|g\rangle$, and $\vartheta_e = \chi(t_{\text{VR}} + t_{\text{idle}} - t_{\text{reset}})/2$ if it is found in $|e\rangle$. Although the $|f\rangle$ state is not computational, our controller is able to reset it with an accompanying virtual rotation by angle ϑ_f . Instead of relying on the simple dispersive coupling model, in experiment we independently calibrate the angles $\vartheta_{g/e/f}$ with a variation of the out-and-back experiment [10] to account for additional minor timing contributions related to FPGA program entering or exiting a subroutine. These calibrated angles are used to initialize the QEC circuit for training.

Another important aspect of the virtual rotation is the switching between momentum and position quadratures of the oscillator to realize $\mathcal{R}_\Delta^x$ and $\mathcal{R}_\Delta^z$ dissipators. Such switching can be achieved with a rotation of the SBS unitary by $\pi/2$ in phase space. This results in an additional deterministic contribution of ϑ_{SBS} to every virtual rotation gate. The value of ϑ_{SBS} is $\pi/2$ for the square grid code, and $\pi/3$ for hexagonal grid code. Note that this is the only change in the protocol that would be required to stabilize the hexagonal code.

The virtual rotation gate utilizes a floating point register ϑ on the FPGA. During this gate, the FPGA performs the calculation $\vartheta \leftarrow \vartheta + \vartheta_{g/e/f} + \vartheta_{\text{SBS}}$ with subsequent reconfiguration of the dynamic mixer matrix which applies a rotation transformation to the oscillator pulses before they are being streamed at the DAC. The total duration of the virtual rotation gate that includes all these steps is independent of ϑ and is equal to $t_{\text{VR}} = 448 \text{ ns}$.

Idle section. During the agent training, the reward is measured after a fixed number of $T = 160$ cycles. Therefore, if our goal is to optimize for the logical lifetime, it is necessary to keep the duration of a QEC cycle constant across different protocol candidates to ensure a fair reward comparison (i.e. for all candidates the reward should be measured after the same physical duration of time). However, the agent is able to affect the physical duration of the SBS unitary by changing the speed enhancement factor α in the ECD gates. To reconcile these two requirements, after the virtual rotation gate we add a section of idle time that is calculated based on the duration of the SBS unitary in each circuit candidate. We constrain the combined duration of SBS unitary and the idle section to be $2 \mu\text{s}$. In the circuit that achieved the highest reported QEC gain, the duration of the idle section was $t_{\text{idle}} = 452 \text{ ns}$. Note that in the previous QEC experiments that tried to create a long-lived quantum memory [5, 6], the idle section was inserted intentionally to avoid frequently entangling the high-quality oscillator with low-quality auxiliary qubit. Here, we find that inserting any additional idle time degrades the performance, hence we kept it nearly to a minimum while still leaving some room for change of the SBS duration by the RL agent.

Component	Subcomponent	Duration (ns)
	Enter cycle	24
SBS	Enter SBS	24
	Circuit layer 1	502
	Circuit layer 2	708
	Circuit layer 3	262
	Circuit layer 4	76
	Exit SBS	24
Reset	Enter reset	24
	Roundtrip delay	300
	Acquisition	1400
	Signal processing	332
	Distribution of s_0 and s_1	100
	Branching and feedback	200
	Exit reset	24
Virtual rotation	Mixer matrix calculation	400
	Mixer update	48
Idle	Delay	452
	Exit cycle	24

Supplementary Table III. **Timing of the cycle components.**

In Supp. Table III, we provide a detailed timing breakdown of all components of the cycle.

D. Learned and scripted parameters


Our QEC protocol has multiple parameters which could be optimized to improve its performance. Some of these parameters are difficult to incorporate into our optimization framework in its current form, and therefore their values are chosen as an approximate compromise between various tradeoffs and then held constant. The remaining $P = 45$ of them are optimized with reinforcement learning, given a reasonable starting point obtained from independent calibrations. Here, we briefly explain the meaning of these parameters.

Scripted parameters:

- Duration, shape, and amplitude of the readout pulse.
- State classification thresholds.
- Timing of all components of the reset subroutine.
- Durations of primitive pulses (qubit rotations and oscillator displacements).
- Combined duration of the SBS unitary and the idle section.

Learned parameters:

- Virtual rotation angle ($\vartheta_g, \vartheta_e, \vartheta_f$) for each measurement outcome. It is initialized with a result of independent calibration using a variation of the out-and-back experiment.
- Detuning of the transmon $|g\rangle \leftrightarrow |e\rangle$ and $|e\rangle \leftrightarrow |f\rangle$ pulses (same parameter for all pulses). It is initialized with a result of independent calibration, when the oscillator is in the vacuum state (i.e. when there is no Stark shift).
- Spectral corrections to the $|g\rangle \leftrightarrow |e\rangle$ and $|e\rangle \leftrightarrow |f\rangle$ pulses based on derivative reduction by adiabatic gate (DRAG) scheme [33] (same parameter for all pulses). It is initialized with 0.
- Complex-valued amplitudes of the conditional displacement gates in the first three layers of the SBS unitary, and a complex-valued amplitude of the unconditional displacement in the fourth layer. Two small amplitudes are initialized with $\beta_{S1} = \beta_{S2} = 0.2i$, and the big amplitude is initialized with $\beta_B = \sqrt{2\pi}$. The unconditional displacement is initialized with 0.
- Magnitudes of the intermediate large displacements used to execute the ECD gates in the first three layers of the SBS unitary. They are initialized with $\alpha_{S1} = \alpha_{S2} = 6$ for small conditional displacements and $\alpha_B = 16$ for the big conditional displacement. Note that changing these parameters also influences the duration of the ECD gates.
- Angular corrections to intermediate large displacements in the first three layers of the SBS unitary. These heuristic parameters compensate for the effect of the second-order dispersive shift and for the fact that the conditional displacement accumulates along an arc of small curvature instead of a straight line. These corrections are initialized with 0.
- Phases and angles of all auxiliary qubit rotations in the SBS circuit layers (including the echo pulses inside the ECD gates), and in the reset subroutine. These parameters are initialized with nominal values from Supp. Table II.



Supplementary Figure 14. **Evolution of QEC circuit parameters during the training.** Top row: example from one particular training run. The shaded region bounds the minimal and maximal sampled parameter values. Solid lines indicate the mean. Bottom row: evolution of the mean parameter values in several independent training runs performed during a two-day period, showcasing the reproducibility of the training results.


◦ Detuning of the local oscillator (LO) frequency for the cavity mode. This LO is calibrated with spectroscopy and set to be half-way between the number-split oscillator frequencies when the qubit is in the states $|g\rangle$ and $|e\rangle$, corresponding to $\Delta = 0$ in Eq. (1).

Evolution of parameters during training. In Supp. Fig. 14, we show the evolution of several QEC circuit parameters during the training. Most parameters, when initialized well, merely exhibit small fluctuations around the mean. However, some parameters undergo systematic and reproducible changes, as observed in the provided examples. For instance, in Supp. Fig. 14(a), the big conditional displacement amplitude $\text{Re}[\beta_B]$, which we expect to be equal to the size of the grid unit cell, changes from a calibrated value of $\sqrt{2\pi}$ by about 8%, likely indicating the presence of a miscalibration error (the last calibration was done several weeks prior to this training). Similarly, the trend in Supp. Fig. 14(e) towards the negative detuning of the $|e\rangle \leftrightarrow |f\rangle$ pulses could be compensating for an additional Stark shift that was not present at the initial calibration stage (calibration was performed with the oscillator in its vacuum state). In Supp. Fig. 14(b), the trend in amplitudes β_{S1} and β_{S2} of the two small conditional displacements in the SBS unitary is particularly insightful, as it helped us identify a limitation of the proposal in Ref. [22], according to which the amplitudes β_{S1} and β_{S2} should be identical and equal to $i\Delta^2/2$, while the RL agent systematically converges to $|\beta_{S2}| > |\beta_{S1}|$. Using simulations, we verified that in the presence of error channels that act during the execution of the SBS unitary, this is indeed a correct inequality. The optimal ratio of these two amplitudes is found in simulations to be strongly dependent on the error channel. The agent adapts this ratio to the real error channel of our system.


In Ref. [8], it was shown that a similar RL agent is able to converge to correct solutions even starting from completely random parameter initializations. Here, we initialize the parameters close to their expected optimal values through various calibrations, but the compounded effect of small errors (at the level of a few percent) in multiple parameters results in a QEC protocol which, although fully functional, is far from optimal. In particular, we were not able to reach break-even with only the independent calibrations and educated guesses, hence model-free RL can be acknowledged as one of the most crucial factors in the success of this project.

E. Syndrome measurement statistics

A sample of 600 experimental QEC shots is shown in Supp. Fig. 15, where QEC is run for $T = 1000$ cycles in each shot. Consider a string of measurement outcomes going from any chosen time step t_i to a time step t_{i+2n} (it contains n QEC cycles). The probability $P([gg]^n)$ that this string contains only gg outcomes is shown in Supp. Fig. 16(a), where it is averaged over the experimental shots and over initial times t_i (the averaging over t_i is done with a sliding window method which is applicable due to process stationarity). While in the most general case the functional form of $P([gg]^n)$ is a sum of multiple decaying exponentials, we clearly observe only a single dominant exponential contribution. Hence, we fit this probability to $P([gg]^n) = a\lambda^n$, obtaining $a = 0.936 \pm 0.003$ and $\lambda = 0.86517 \pm 0.00013$. By adopting a model for the error process and for the QEC process, we can link the fit parameters $\{a, \lambda\}$ to model parameters. In general, such a model would be quite complex. However, here we are interested in only two characteristic parameters




Supplementary Figure 15. **Syndrome measurement outcomes.** A sample of 600 experimental QEC shots of duration $T = 1000$ cycles each. The “ g ” outcome (green) is prevalent, heralding the no-error process, while occasional “ e ” outcomes (yellow) indicate correction of errors, and “ $\geq f$ ” outcomes (red) indicate leakage. When transmon escapes to a state higher than $|f\rangle$, which is not addressed by our reset scheme, the leakage outcome persists for multiple cycles (streaks of red). In the readout IQ plane such states occur above the Q threshold, and therefore they are conveniently classified as leakage, but without further identification of the exact leakage state. After transmon stochastically drops back to $|f\rangle$, the controller is able to reset it and return the transmon to the computational manifold. However, during the cycles when transmon is effectively inactive, the code is not stabilized. Hence, leakage streaks are often followed by streaks of e outcomes where QEC re-stabilizes the code manifold.



Supplementary Figure 16. **Analysis of syndrome measurement outcomes.** (a) Probability of a string $gg/gg/\dots$ as a function of its length, together with the fit to a single exponential decay. (b) Histogram of durations of leakage events. Events of duration 1 or 2 cycles are predominantly $|f\rangle$ state; longer duration events are likely $|h\rangle$ or higher excited states. The exponential fit gives the effective lifetime of these leakage states. (c) Fraction of shots that experienced a leakage event of duration d up to a given time. Dotted lines are fits to a constant-rate model, see Section IV F.

of the process: the probability $\langle \Pi_0 \rangle$ of occupying the code space in the dynamical equilibrium of the QEC process, and the probability p_{err} of having an error that transfers the state out of the code space. These parameters can be extracted with minimal model assumptions.

Using the transfer matrix approach, it can be shown that $p_{\text{err}} \approx 1 - \lambda$ and $\langle \Pi_0 \rangle \approx a\lambda$ under the following assumptions: 1) The error probability is small $p_{\text{err}} \ll 1$, which is justified since the cycle duration is small compared to all relevant error rates in the system, and is confirmed by the fit results; 2) The conditional probability $P_{gg}^{C_0}$ of emitting gg when the quantum state is in the code space is nearly 1. This is justified, since in the error-free model of the SBS



Supplementary Figure 17. **Correlation of syndrome measurement outcomes.** (a) Correlation matrix r_{ij} for the first 30 cycles, computed from the dataset in Supp. Fig. 16(a). (b) r_{ij} for the full QEC duration of 1200 cycles with zoomed-in color scale to resolve small numbers. (c) r_{ij} after removing leakage events of duration ≥ 2 cycles. (d, e) Cuts of the correlation matrices from (b) and (c) at different locations in the QEC trajectory. In a stationary process, r_{ij} would only depend on $|i-j|$, which is clearly not satisfied in (d). After removing length ≥ 2 leakage events, the deviation from stationarity is not resolvable.

protocol described in Section IV B this probability is 0.999; 3) The conditional probability P_{gg}^{err} of emitting gg when the quantum state is in any of the error spaces is small $P_{gg}^{\text{err}} \ll 1$. This assumption is partially justified, since in the error-free model of the SBS protocol this probability is smaller than 0.5 for the first error space, and then monotonously reduces for the higher levels of the error hierarchy; 4) The probability P_{gg}^{corr} of correcting an error and emitting gg is small $P_{gg}^{\text{corr}} \ll 1$. In the error-free model of the SBS protocol this probability is exactly zero, while in practice it is limited to $\sim 10^{-2}$ due to the readout infidelity of the $|e\rangle$ state.

This result can be intuitively understood as follows: the two most probable system trajectories that generate the string of all gg 's correspond to (i) starting in the code space and remaining there for n steps, which happens with probability $\langle \Pi_0 \rangle (1 - p_{\text{err}})^n$, and (ii) starting in the code space, remaining there for $n - 1$ steps, and transitioning out on the very last step, which happens with probability $\langle \Pi_0 \rangle (1 - p_{\text{err}})^{n-1} p_{\text{err}}$. The sum of these two contributions equals $\langle \Pi_0 \rangle (1 - p_{\text{err}})^{n-1} \equiv a \lambda^n$, leading to $p_{\text{err}} \approx 1 - \lambda$ and $\langle \Pi_0 \rangle \approx a \lambda$. The corrections to these formulas are of the second order in parameters $\{p_{\text{err}}, 1 - P_{gg}^{\text{C}_0}, P_{gg}^{\text{err}}, P_{gg}^{\text{corr}}\}$. From the fit in Supp. Fig. 16(b), we extract $p_{\text{err}} = 0.13 \pm 0.02$ and $\langle \Pi_0 \rangle = 0.81 \pm 0.02$. Note that this result for $\langle \Pi_0 \rangle$ agrees within the error margin with the result obtained by an independent method based on the reconstruction of the density matrix from the measured Wigner functions in Section IV H. Here, the error of the fit is negligible compared to the model approximations, hence the quoted error bars are obtained from an estimate of the second-order corrections $\sim p_{\text{err}}^2 \approx 0.02$.

Here, we only considered the string of a special type $gg/gg/\dots$; an important avenue of future research would include learning the error channel from the full statistics of syndrome outcomes, using the dataset in Supp. Fig. 15.

F. Analysis of transmon leakage

To quantify transmon leakage, we histogram its duration in Supp. Fig. 16(b). The most likely leakage duration is 1 cycle, since the controller resets $|f\rangle$ to $|g\rangle$ with high probability. However, because of the finite readout fidelity of $|f\rangle$,

shown in Supp. Fig. 5(b), the controller sometimes fails to reset this state, resulting in the next most probable leakage duration of 2 cycles. After that, the histogram follows an exponential distribution with decay constant of 17.2 cycles, corresponding to $85 \mu\text{s}$, which we attribute to the effective lifetime of higher leakage states that are not addressed by our reset scheme. This time scale is consistent with our estimate $\sim 280 \mu\text{s}/3 = 93 \mu\text{s}$ of the $|h\rangle$ state lifetime, derived from the bosonic statistics and the average measured lifetime of the $|e\rangle$ state.

To extract the leakage rate, in Supp. Fig. 16(c) we plot the fraction of shots that experienced a leakage event of a certain duration up to a given cycle. We fit the data to a constant-rate model $L(t) = 1 - \exp(-t/\tau_l)$ where t is the cycle index and $p_l = 1/\tau_l$ is the leakage rate (i.e. leakage probability per cycle). We find $\tau_l = 1480 \pm 10$ cycles, corresponding to a leakage rate of $p_l = (6.76 \pm 0.04) \times 10^{-4}$. Similar analysis can be done for leakage events of length ≥ 2 ; with the fit to the same model, we find the time scale of $\tau_{l,\geq 2} = 7820 \pm 10$ cycles, and the corresponding rate of $p_{l,\geq 2} = (1.280 \pm 0.002) \times 10^{-4}$. Note that this measurement was performed at the time of slightly sub-optimal performance, and therefore the leakage rate at the maximal achieved QEC gain might have been smaller.

Next, we study the correlation of syndrome measurement outcomes across time. The correlation matrix is given by

$$r_{ij} = \frac{\mathbb{E}[m_i m_j] - \mathbb{E}[m_i]\mathbb{E}[m_j]}{\sqrt{(\mathbb{E}[m_i^2] - \mathbb{E}[m_i]^2)(\mathbb{E}[m_j^2] - \mathbb{E}[m_j]^2)}}, \quad (20)$$

where m_k is the measurement outcome obtained at cycle k , and empirical expectation values are obtained by averaging across experimental shots. For the dataset of Supp. Fig. 15 this correlation matrix is shown in Supp. Fig. 17(a), where we consider only the first 30 cycles. Overall, the correlation is weak, and the correlation between $\mathcal{R}_\Delta^x$ and $\mathcal{R}_\Delta^z$ channels (separated by odd numbers of cycles) is weaker than the correlation between the same-quadrature channels (separated by even numbers of cycles).

By zooming in the color scale to visually resolve small numbers and considering the full duration of the trajectory of 1200 cycles, as shown in Supp. Fig. 17(b), it becomes evident that the process is not perfectly stationary. To emphasize this, we show in Supp. Fig. 17(d) the correlation coefficient r_{ij} as a function of $|i - j|$ for several choices of j . Further along the QEC trajectory the process acquires a correlation tail. Although quite weak, this correlation stretches over hundreds of cycles.

Previously, it was demonstrated that leakage removal helps to reduce correlated errors in the arrays of transmons [34]. Our QEC protocol already contains a mechanism for leakage removal from the $|f\rangle$ state through measurement-based feedback in every cycle. However, leakage states higher than $|f\rangle$ are not cleared by our reset. The signature of such leakage events to higher states is two or more consecutive leakage syndrome outcomes. To check the hypothesis that this residual leakage to states higher than $|f\rangle$ is responsible for increase of correlation, we post-select trajectories that do not have any length-two or longer leakage events. In the post-selected dataset, the correlation matrix does not display any detectable non-stationarity, as seen in Supp. Fig. 17(c,e), confirming the hypothesis. By fitting the remaining short-time correlations in Supp. Fig. 17(e) to an exponential decay, we conclude that it takes 3.9 ± 0.1 cycles (approximately 2 QEC cycles) to lose the memory of a typical large error. However, the most probable small errors are corrected in a single QEC cycle.

	NP	L	$d \geq 5$	$d \geq 4$	$d \geq 3$	$d \geq 2$	$d \geq 1$
Survival prob. per cycle	1.0000	0.9985	0.9993	0.9986	0.9972	0.9907	0.9396
Improvement of Γ_{GKP}	1.00	1.10	1.18	1.36	1.68	2.44	6.31
Lifetime of $ +Z\rangle$ (ms)	1.874 ± 0.004	2.11 ± 0.01	2.23 ± 0.01	2.55 ± 0.01	3.13 ± 0.02	4.60 ± 0.02	10.0 ± 0.6
Lifetime of $ +Y\rangle$ (ms)	1.147 ± 0.004	1.23 ± 0.01	1.36 ± 0.01	1.56 ± 0.01	1.93 ± 0.01	2.80 ± 0.02	9.4 ± 1.0
Lifetime of $ +Z\rangle$ (cycles)	381 ± 1	427 ± 1	452 ± 2	518 ± 2	636 ± 4	934 ± 4	2000 ± 100
Lifetime of $ +Z\rangle$ (cycles)	233 ± 1	250 ± 1	275 ± 2	317 ± 2	393 ± 3	567 ± 5	1900 ± 200

Supplementary Table IV. **Post-selection results.** Top row labels the post-selection schemes. NP stands for “no post-selection”; L stands for “leakage”; $d \geq N$ means post-selection that discards trajectories containing strings of N or more consecutive e outcomes in the same-quadrature cycles.

G. Post-selection of errors


Here, we provide additional details about the the post-selection experiment that verifies the ability of our QEC scheme to faithfully identify the errors. The post-selection results are summarized in Supp. Table IV. Note that this experiment was performed at the time of slightly sub-optimal system performance, hence the baseline results with no post-selection are lower than in some other experiments reported here, e.g. in Section IV J. The main conclusion of this post-selection experiment is that it enables significant improvement of the error probability at a cost of only a modest rejection probability.

The saturation of lifetimes in the most stringent post-selection scheme (which preserves only the all- g trajectories) can be related to the following mechanisms: (i) Direct logical errors, which are undetectable in any QEC scheme. (ii) Misclassification of e as g (due to transmon decay during the measurement), which means that some of the all- g trajectories that survived the post-selection actually contained errors, and some of those errors might have been close to a logical operation instead of the identity operation. (iii) The non-orthogonality of logical states, which in our case is not a limiting factor.


H. Wigner tomography of logical states

Tomography and its calibration. Wigner tomography is derived from the expression for the Wigner function $W(\alpha) = (2/\pi)\langle\Pi_\alpha\rangle$, where $\Pi_\alpha = D(\alpha)\Pi D^\dagger(\alpha)$ is the displaced parity operator, and $\Pi = \exp(i\pi a^\dagger a)$ is the photon number parity operator. The displaced parity operator is unitary and can be measured with phase estimation. It is also Hermitian, and hence its eigenvalues are constrained to be ± 1 . Therefore, it is particularly convenient to measure displaced parity by mapping it onto the qubit observable [35], which is achieved in our system with a circuit shown in Supp. Fig. 18(a). The conditional rotation gate $\text{CR}(\pi)$ is realized with a delay of duration π/χ under the dispersive coupling Hamiltonian, which amounts to approximately $10\ \mu\text{s}$. Because of such long duration, previous GKP experiments with similarly small χ chose to perform state tomography using the characteristic function instead [10, 14]. However, the long coherence of our system allows us to measure the Wigner function with reasonably high fidelity.

We use several calibration techniques to improve the quality of the subsequent state reconstruction from the tomographic data. First, to symmetrize the effect of transmon relaxation during the readout, we map the $+1$ eigenvalue of Π_α to the g outcome in half of the phase estimation runs and to the e outcome in the other half. This technique eliminates any finite offset in $W(\alpha)$, but maintains the reduced contrast due to transmon relaxation and decoherence. Next, to calibrate the contrast reduction, we perform an experiment with a similar circuit in




Supplementary Figure 18. **State reconstruction.** (a) Wigner tomography experiment. (b) Calibration experiment to extract α -dependent contrast of the Wigner tomography. This calibration relies on the assumption that the loss of contrast in tomography is primarily due to incoherent errors during the parity mapping gate $\text{CR}(\pi)$. (c) Results of the qubit state tomography in the calibration experiment shown in (b). The measurement of $\langle\sigma_x\rangle$ is fit to a quadratic function of $|\alpha|$ (black dashed line), and its square root (black dotted line) is used as a Wigner function measurement contrast in the state reconstruction. (d) Experimental Wigner function of the $|+Z\rangle$ state immediately after initialization. (e) Wigner function of the reconstructed state. (f) Real part of the density matrix of the reconstructed state in the photon number basis.



Supplementary Figure 19. **Wigner tomography after QEC.** Evolution of $|+Z\rangle$ state (1st row) and $|-Z\rangle$ state (2nd row) is followed for 800 cycles. Color scheme is the same as in Supp. Fig. 18, with the range scaled to $[-0.63, 0.63]$. Marginal of the Wigner function along momentum (position) quadrature, which gives probability density of the oscillator position (momentum), is shown in the 3rd (4th) row in blue for $|+Z\rangle$ state, and orange for $|-Z\rangle$ state. The probability density is not normalized.

which $\text{CR}(\pi)$ is replaced with $[\text{CR}(\pi)]^2 = I$, see Supp. Fig. 18(b). We fit the result of this experiment, shown in Supp. Fig. 18(c), to $\langle \sigma_x \rangle = 1 - p[\alpha]$, where $p[\alpha] = \eta_0 + \eta_2 |\alpha|^2$ is the purity loss per $\text{CR}(\pi)$ gate. Under the assumption that reduction in tomography contrast is primarily due to incoherent processes (transmon relaxation and dephasing, and oscillator photon loss), the inferred tomography contrast is $P(\alpha) = \sqrt{1 - p[\alpha]}$. At $\alpha = 0$, this inferred contrast is equal to 0.8, which matches the measured contrast of the Wigner function of vacuum in Supp. Fig. 3(c), justifying the assumptions of this calibration method.

The phase space points α_i for Wigner tomography are chosen on a square 81×81 grid in a complex plane restricted to $|\text{Re}[\alpha_i]|, |\text{Im}[\alpha_i]| \leq 3.2$. We acquire 2400 shots per point in 6 separate acquisition time frames. Between the time frames we perform system performance checks; data acquisition is put on hold if the spurious resonance in $T_1^t(\bar{n})$ reappears [see Supp. Fig. 5(d)]. A single state tomography dataset consists of 15.7 million shots, and takes a long time to acquire – from 6 hours in the case of $T = 0$ cycles, to 26 hours in the case of $T = 800$ cycles. Therefore, conclusions derived from the analysis of tomography data apply to long-time average system performance.



Supplementary Figure 20. **QEC subspaces.** (a) Spectrum of the reconstructed density matrices of the $|+Z\rangle$ state evolving under the QEC process. The spectrum separates into two pairs of eigenvalues: one pair corresponds to the code space $\mathcal{C}_0$, and another pair corresponds to an error space obtained from the code space by application of an error operator E . Dashed lines show the sum of the eigenvalues within each pair, which gives the probability of occupying the code space and the error space. (b) Wigner function of the projector Π_0 onto the first pair of eigenvectors (taken at $T = 800$), defining the code space. (c) Wigner function of the projector $E\Pi_0 E^\dagger$ onto the second pair of eigenvectors, defining the error space. This error space corresponds to $\mathcal{C}_2$ subspace in Section IV B.

State reconstruction. Tomographic data is used to produce a best guess for the density matrix of the state. We parametrize the density matrix as $\rho = C^\dagger C / \text{Tr}[C^\dagger C]$, where $C = A + iB$, and A and B are real-valued matrices. Such parametrization ensures that ρ is positive semi-definite with trace 1. We truncate the density matrix to dimension $N = 32$ in the photon number basis. Coefficients of the matrices A and B are optimized using the least squares fit of the Wigner tomography and contrast data, with the cost function given by

$$\text{cost} = \sum_{i=1}^{N_\alpha} \left(\frac{2}{\pi} \text{Tr}[\rho D(\alpha_i) \Pi D^\dagger(\alpha_i)] \cdot P(\alpha_i) - W(\alpha_i) \right)^2. \quad (21)$$

An example Wigner tomography of the $|+Z\rangle$ grid state together with its reconstruction is shown in Supp. Fig. 18(d-f). As will be described shortly, using the reconstructed density matrix we extract various parameters of the state: its purity, mean photon number, and envelope size.

Evolution of logical states. We visualize the evolution of logical $|+Z\rangle$ and $|-Z\rangle$ grid states during the QEC by taking Wigner tomography snapshots after 0, 100, 200, 400, and 800 cycles, with results shown in Supp. Fig. 19.


The marginal of the Wigner function along the momentum quadrature gives the probability density of the oscillator position, shown in the third row of Supp. Fig. 19. The $|+Z\rangle$ and $|-Z\rangle$ states have non-overlapping support in the position representation, clearly observed in the data at $T = 0$ cycles. During the QEC process these basis states mix under the logical Pauli channel, which manifests in the appearance of position peaks of the opposite state, until finally they become almost (but not completely) indistinguishable after $T = 800$ cycles.

On the other hand, the marginal of the Wigner function along position quadrature gives the probability density of the oscillator momentum, shown in the last row of Supp. Fig. 19. In the momentum representation, $|+Z\rangle$ and $|-Z\rangle$ states share the same support, but have a different pattern of phases associated with the peaks of the wavefunction. The phase information is discarded in the probability density function, which looks identical for both states.

Spectral analysis of reconstructed states. Focusing on the time evolution of the $|+Z\rangle$ state, we perform spectral decomposition of its reconstructed density matrices at $T = 100, 200, 400, 800$. We find that the eigenvalues of the density matrix are arranged in pairs corresponding to the images of this state and of its complement $|-Z\rangle$ in different subspaces of the QEC, see Supp. Fig. 20(a). In particular, we identify only two subspaces with a substantial presence of the state during the QEC process: the code space $\mathcal{C}_0$, shown in Supp. Fig. 20(b), and the error space $\mathcal{C}_2$ corresponding to an error E that most closely resembles $a^\dagger$, see Supp. Fig. 20(c) and Section IV B.

While the QEC circuit imposes the structure of the error subspaces, as described in Section IV B, the properties of the “thermal” distribution across these subspaces in the dynamical equilibrium is defined by the strength of the various error mechanisms in our system as well as the rate at which these errors are corrected. The probability of occupying the code space, given by the sum of the first two eigenvalues, remains constant over time and equal to $\langle \Pi_0 \rangle = 0.825 \pm 0.003$, where error bar represents the standard deviation with respect to different durations of the QEC process. This value agrees well with an independent analysis in Section IV E. Having only one relevant error subspace in the steady-state distribution also qualitatively agrees with an observation in Section IV E that errors are rare. We believe that other error subspaces are populated with probability $< 1\%$, which is beyond the resolution power of this method. Developing a more accurate and sample-efficient reconstruction technique for characterizing the distribution across the error spaces is an important direction left for the future.

Extracting the code envelope size. For each QEC duration T , the reconstructed density matrix is used to find the fidelity of the experimental states to the family of finite-energy codewords $\{|\pm Z_\Delta\rangle\}$ parametrized by the envelope



Supplementary Figure 21. **Analysis of $|+Z\rangle$ evolution.** (a) Expectation value $\mathcal{F}_\Delta^\pm$ of the state projectors $|\pm Z_\Delta\rangle\langle\pm Z_\Delta|$ for a range of values of Δ and for different durations of the QEC process. The fidelity $\mathcal{F}_\Delta^+$ decreases, while $\mathcal{F}_\Delta^-$ increases as a function of time, as expected for a logical Pauli channel. The expectation value of the code projector $\langle \Pi_\Delta \rangle = \mathcal{F}_\Delta^+ + \mathcal{F}_\Delta^-$ (black dotted lines) remains nearly time-independent for $T > 0$. (b) Purity of the reconstructed state ρ and of its projection onto the code space ρ_Δ as a function of time. (c) Average photon number as a function of time.

size Δ . For $T > 0$, we additionally displace the target codewords by $(0.08 - 0.12i)\sqrt{\pi/2}$ to account for a small shift visible in the tomography. Since these target states are pure, the fidelity is given by $\mathcal{F}_\Delta^+ = \text{Tr}[|+Z_\Delta\rangle\langle+Z_\Delta|\rho]$ and $\mathcal{F}_\Delta^- = \text{Tr}[|-Z_\Delta\rangle\langle-Z_\Delta|\rho]$. For experiments that start with a preparation of $|+Z\rangle$, these fidelities are shown in Supp. Fig. 21(a) for each QEC duration from the dataset in Supp. Fig. 19. Immediately after the initialization, the fidelity $\mathcal{F}_\Delta^+$ is maximized for $\Delta = 0.36$, where it reaches 0.85. During the QEC process, $\mathcal{F}_\Delta^+$ gradually reduces while $\mathcal{F}_\Delta^-$ increases, consistent with the logical Pauli channel. The sum $\langle \Pi_\Delta \rangle = \mathcal{F}_\Delta^- + \mathcal{F}_\Delta^+$, which is equal to the expectation value of the code projector, remains nearly constant for $T > 0$. It is maximized at $\Delta = 0.34$ (9.4 dB), where it is equal to $\langle \Pi_\Delta \rangle = 0.817 \pm 0.003$. This value is close to $\langle \Pi_0 \rangle = 0.825 \pm 0.003$ extracted from the density matrix spectrum, which indicates that the code $\mathcal{C}_0$ stabilized in the experiment is indeed the one defined by the envelope operator $\exp(-\Delta^2 n)$ with $\Delta = 0.34$.


The purity of the $|+Z\rangle$ grid state is shown in Supp. Fig. 21(b). During the QEC process, it reduces below 0.5, since the steady state contains a mixture of the codewords with their images in the error spaces. However, the part of this mixed state that resides within the code space should approach a purity of 0.5. To confirm that this is the case, for every state ρ we define its projection onto the code space as $\rho_\Delta = \Pi_\Delta \rho \Pi_\Delta / \text{Tr}[\Pi_\Delta \rho \Pi_\Delta]$, where we only consider the code with the optimal envelope $\Delta = 0.34$. As seen in Supp. Fig. 21(b), the purity of ρ_Δ after initialization is close to 1, and after hundreds of cycles it approaches 0.5, as expected for the logical Pauli channel.

Lastly, in Supp. Fig. 21(c) we plot the evolution of the average photon number $\langle n \rangle = \text{Tr}[a^\dagger a \rho]$ during the QEC process. The extracted steady-state photon number is $\langle n \rangle = 4.67 \pm 0.02$.


I. Sensitivity to auxiliary qubit errors

Auxiliary qubit phase flips. Our QEC circuit is fault-tolerant with respect to auxiliary qubit phase flips by design [22]. Consider the effect of a σ_z error. If it happens during the auxiliary qubit readout, it would have no effect because the readout projects the auxiliary qubit onto an eigenstate of σ_z . Likewise, during the virtual rotation gate or the idle time the auxiliary qubit is nominally in the $|g\rangle$ state, which is an eigenstate of σ_z . Finally, the effect of a σ_z error on the SBS unitary can be understood by propagating it through the circuit layers. For example, if such an error happens during the big conditional displacement, it is equivalent to changing the circuit parameters from $\vec{\beta} = l_S \times (i\Delta^2/2, 1, i\Delta^2/2)$ to $\vec{\beta} = l_S \times (i\Delta^2/2, 1, -i\Delta^2/2)$. Since $\Delta^2 \ll 1$, this change is equivalent to a small error that will be corrected in the following QEC cycles.

Auxiliary qubit bit flips. In contrast, auxiliary qubit bit flips can detrimentally affect the logical qubit in several ways. If such an error happens during the middle half of the big conditional displacement of amplitude l_S , it will with high probability generate a logical error. This mechanism accounts for a significant fraction of logical errors in the experiment. For example, its contribution to the Pauli error probability $p_X = 1.8 \times 10^{-3}$ (per QEC cycle), is estimated to be $\sim 0.5 \times 0.5 \times 0.5 \times (700 \text{ ns}/280 \mu\text{s}) \approx 0.3 \times 10^{-3}$, where the factors of 0.5 account for (i) half of the superposition state being sensitive to auxiliary qubit decay, (ii) half of the QEC cycle is devoted to position quadrature, (iii) half of the big conditional displacement gate. In practice, the relaxation time of the auxiliary qubit is likely degraded during the execution of the conditional displacement due to the large number of intermediate photons in the oscillator, see



Supplementary Figure 22. **Effect of auxiliary qubit errors.** Logical error rates γ_Z and γ_Y as a function of physical error rates γ_ϕ^t and γ_1^t of the transmon (γ_X is expected to behave identically to γ_Z). Physical error rates are varied with noise injection. The error sensitivities $d\gamma_Z/d\gamma_1^t = 0.17$, $d\gamma_Y/d\gamma_1^t = 0.25$, $d\gamma_Z/d\gamma_\phi^t = 0.0027$, and $d\gamma_Y/d\gamma_\phi^t = 0.0050$ are extracted by linear fits in the low-error region. The derived sensitivity of $\Gamma_{\text{GKP}} = (\gamma_X + \gamma_Y + \gamma_Z)/3$ to phase flips is 65 times smaller than to bit flips.



Supplementary Figure 23. **Noise verification experiments.** (a) Component error rates as a function of the root mean square (RMS) voltage of the injected noise (at the generator plane). Here, noise is up-converted to the qubit frequency to increase γ_1^t . (b) Same as in (a), but with the baseband noise that increases γ_ϕ^t . (c) Readout infidelity of the transmon $|g\rangle$ and $|e\rangle$ states in these two noise injection settings.

[10] and evidence in Section IV J. Hence, this estimate provides an optimistic lower bound. Auxiliary qubit bit flips can also create detrimental back-action on the oscillator if they happen during the readout time. Since the readout outcome is used in a feedback loop to implement a virtual rotation gate, misclassification of the auxiliary qubit state generates rotational errors that the GKP code is not well suited to correct. These erroneous rotation angles are distributed in the range $0.0 - 0.6$ radians. To estimate the contribution of such errors to the logical error rate, consider that a rotation by $\sim \Delta/(l_S/2) \approx 0.3$ radians would diminish the overlap of the blobs in the Wigner function; therefore, a significant fraction of misclassification-induced rotation errors cause a large disturbance of the stabilized code space. The transmon $|e\rangle$ state is the most prone to misclassification. Since readout fidelity of the $|e\rangle$ state is close to 99%, and this outcome is generated 10% of the time, we estimate an additional $\sim 0.5 \times 10^{-3}$ contribution to logical error probability per QEC cycle from this mechanism. The two contributions described here account for half of the logical error probability p_X ; the remaining half is not yet well understood.

Transmon noise injection. To investigate the effect of auxiliary qubit errors on the logical performance in a controllable way, we perform noise injection experiments that selectively increase the transmon phase-flip rate γ_ϕ^t or bit-flip rate γ_1^t , the results of which are shown in Supp. Fig. 22. With noise injection, we are able to increase γ_1^t by a factor of 14 (spoiling T_1^t from $290 \mu\text{s}$ to $20 \mu\text{s}$), and γ_ϕ^t by a factor of 140 (spoiling T_ϕ^t from $430 \mu\text{s}$ to $3 \mu\text{s}$). Using linear fits in the low-error region, we extract the error sensitivities $d\gamma_Z/d\gamma_1^t = 0.17$, $d\gamma_Y/d\gamma_1^t = 0.25$, $d\gamma_Z/d\gamma_\phi^t = 0.0027$, and $d\gamma_Y/d\gamma_\phi^t = 0.0050$. The derived sensitivity of $\Gamma_{\text{GKP}} = (\gamma_X + \gamma_Y + \gamma_Z)/3$ to auxiliary qubit phase flips is 65 times smaller than the sensitivity to auxiliary qubit bit flips, confirming the qualitative arguments provided above. Mitigating the effect of auxiliary qubit bit flips on the logical performance is one of the most important future directions in grid-code QEC.

Verifying noise injection. In the following, we explain how the noise injection experiments were conducted and how we verified that the noise affects the system as intended, i.e. selectively tunes γ_ϕ^t or γ_1^t . We are able to achieve a high degree of selectivity, with negligibly small spurious effects. To spoil γ_1^t , we inject noise at the transmon frequency, and to spoil γ_ϕ^t , we inject noise at low frequency [36]. The baseband white noise with flat spectral density

up to 80 MHz is sourced from an Agilent 33250A arbitrary waveform generator. In the γ_1^t tuning experiment, it is upconverted to the qubit frequency using a double-balanced mixer with an LO blue-detuned by 30 MHz from the qubit frequency. After this pre-processing, the noise is filtered and combined with the qubit control line after the switch (in contrast to all control pulses, the noise is not gated).

In Supp. Fig. 23(a), we inject resonant noise to tune γ_1^t . This noise couples to the σ_x operator and therefore changes $\gamma_\downarrow^t$ and $\gamma_\uparrow^t$ symmetrically, which results in increased steady-state population of the qubit, approaching 0.5 at the largest applied noise power. Note that this noise also affects the dephasing rate γ_{2E}^t , but the changes in γ_{2E}^t are explained by changes in γ_1^t : the extracted pure dephasing rate $\gamma_\phi^t = \gamma_{2E}^t - \gamma_1^t/2$ remains independent of the noise power, as intended in this experiment. The error bars on γ_ϕ^t increase at large noise power, because this small rate is extracted as a difference of two large rates. The oscillator dephasing rate γ_2^c is also affected by the noise, which is explained by the increased rate of qubit up- and down-transitions that dephase the oscillator through the dispersive coupling [19]. The pure dephasing rate $\gamma_\phi^c = \gamma_2^c - \gamma_1^c/2$ of the oscillator agrees reasonably well with the prediction $\gamma_\phi^{c,t} = n_{\text{th}}^t \gamma_\downarrow^t$ derived from this mechanism (black dotted line). The disagreement at high noise power is under investigation; it likely comes from the breakdown of the simple formula for $\gamma_\phi^{c,t}$ in the limit where n_{th}^t is not small.

In Supp. Fig. 23(b), we inject baseband noise to tune γ_ϕ^t . In addition to this desired effect, within the same dynamic range of the noise we observe an undesired increase of the qubit excited state population by a factor of 2 (data not shown), likely due to the heating of the attenuators by the dissipated noise power. Since $\gamma_\uparrow^t/\gamma_1^t \ll 1$, the qubit lifetime is not significantly affected by this heating. The lifetime and coherence of the oscillator also remain independent of the noise power. The increase of the error bars on γ_1^c and γ_2^c with the noise power is related to strong degradation of the fidelity of the transmon selective pulse used to read out the population of the oscillator $|0\rangle$ and $|1\rangle$ states as described in Section II E. This pulse has a duration of $\sim 20 \mu\text{s}$ and it is directly sensitive to the transmon coherence; at the highest injected noise power, where coherence time is spoiled down to $3 \mu\text{s}$, the fidelity of this selective pulse is only a few percent.

In Supp. Fig. 23(c), we show the effect of the noise on the readout fidelity of the transmon $|g\rangle$ and $|e\rangle$ states. In principle, the noise that induces phase flips (σ_z errors) should not affect the readout of σ_z . However, due to the aforementioned heating of the qubit, we observe a weak degradation of $\mathcal{F}_r^{(g)}$. On the other hand, noise at the qubit frequency couples to σ_x and results in significant degradation of both $\mathcal{F}_r^{(g)}$ and $\mathcal{F}_r^{(e)}$.


J. Long-time system stability

We next investigate the stability of our quantum system over time with repetitive measurements of the lifetimes of the $\{|0\rangle, |1\rangle\}$ qubit, $\{|g\rangle, |e\rangle\}$ qubit, and error-corrected GKP qubit. The results of a week-long scan are shown in Supp. Fig. 24. We find that the $\{|0\rangle, |1\rangle\}$ qubit is the most stable. In contrast, the $\{|g\rangle, |e\rangle\}$ qubit exhibits notable fluctuations of the $|e\rangle$ state lifetime. Such fluctuations are often observed in transmons [18, 37], and although they are typically attributed to two-level defects in the amorphous dielectric, their source in our system is not yet understood.

We also find significant fluctuations of the lifetime of the error-corrected GKP qubit. Periods of relative stability are regularly interrupted with sudden drops and resurgences of performance, correlated with the appearance and disappearance of a resonant feature in the $T_1^t(\bar{n})$ dependence, see Supp. Fig. 24(c). The behavior of the readout infidelity of the $|e\rangle$ state is also correlated with this feature, see Supp. Fig. 24(b). We find that the correlation coefficient between the logical error rate and the readout infidelity is $r = 0.81$. However, preliminary simulations indicate that the degradation of the readout fidelity alone is not sufficient to explain the collapses of the logical performance. We therefore believe that the presence of the spurious resonance affects not only the readout fidelity, but also the fidelity of the SBS unitary.

A plausible causal chain is the following: 1) for unknown reasons, the spurious degrees of freedom (defects) appear and disappear; 2) when the transmon $|g\rangle \leftrightarrow |e\rangle$ transition frequency is resonant with the defect, their interaction strength is enhanced, which reduces the lifetime of the $|e\rangle$ state; 3) during readout, the transmon is Stark-shifted into resonance with the defects by the readout photons, which increases the probability of readout errors; 4) when the transmon state is misclassified, the virtual rotation gate is executed with an incorrect angle, inducing a phase-space rotation error on the oscillator; 5) during the conditional displacement gates, the transmon is also Stark-shifted into resonance with the defects by the intermediate photons of the cavity mode; 6) when transmon decay happens during the big conditional displacement gate, it has a significant chance of inducing a logical error. This proposed connection between spurious defects and fluctuations of the logical performance could be verified with detailed system-level simulations that take into account time-dependent Stark shift of the transmon and Stark-shift-dependent degradation of T_1^t , which is left for future analysis.

Apart from the stochastic fluctuations, we observe a systematic drift that warrants periodic retraining of the QEC



Supplementary Figure 24. **Quantum system stability.** (a) Lifetimes of Pauli Y and Z eigenstates of $\{|0\rangle, |1\rangle\}$ qubit, $\{|g\rangle, |e\rangle\}$ qubit, and an error-corrected GKP qubit. (b) Inverse readout infidelity of the transmon $|g\rangle$ and $|e\rangle$ states. Logical lifetime is strongly correlated with $(1 - \mathcal{F}_r^{(e)})^{-1}$. (c) Transmon lifetime T_1^t as a function of the number $\bar{n}$ of the steady-state photons in the readout resonator. The dashed line denotes the DAC amplitude used for the actual readout. The correlated degradation of the system performance and appearance of a spurious resonance that degrades $T_1^t(\bar{n})$ around $\sqrt{\bar{n}} = 0.1$ is indicated with purple arrows.

circuit. This drift can be seen by comparing the initial and final data points of the scan in Supp. Fig. 24, where all the monitored physical error sources are similar in magnitude, but the logical lifetime is reduced in the final point as compared to the initial point. Due to this effect, the various datasets reported in our work were acquired with different versions of the QEC circuit that were retrained every 1-2 weeks.

K. Average channel fidelity

The *average channel fidelity* of a quantum channel $\mathcal{E} : \rho \rightarrow \mathcal{E}(\rho)$ to a target unitary channel $\mathcal{U} : \rho \rightarrow U\rho U^\dagger$ is

$$\overline{\mathcal{F}}(\mathcal{E}, \mathcal{U}) = \int d\psi \langle \psi | U^\dagger \mathcal{E}(|\psi\rangle\langle\psi|) U | \psi \rangle, \quad (22)$$

where the integral is over the uniform measure on the state space, normalized so that $\int d\psi = 1$. Henceforth, we refer to this metric simply as the *fidelity*.

To derive an equivalent but experimentally-compatible expression, we make use of the Pauli transfer matrix (PTM) representation of a channel $R_{ij}[\mathcal{E}] = \frac{1}{2} \text{Tr}(\sigma_i \mathcal{E}[\sigma_j])$, where $\{\sigma_k, k = I, X, Y, Z\}$ are Pauli matrices. This representation has several useful properties, e.g. that the composition of channels corresponds to a product of their PTMs [38]. In terms of the PTM, we have the following expression for fidelity:

$$\overline{\mathcal{F}} = \frac{2\mathcal{F}_e + 1}{3}, \quad \mathcal{F}_e = \frac{1}{4} \text{Tr}(R^T[\mathcal{U}] R[\mathcal{E}]), \quad (23)$$

where $\mathcal{F}_e$ is often called the *entanglement fidelity*.

To benchmark a quantum error correction channel, we compare it to an identity channel $\mathcal{I} : \rho \rightarrow \rho$ with $R_{ij}[\mathcal{I}] = \delta_{ij}$.

In this case, Eq. (23) can be further simplified to

$$\overline{\mathcal{F}} = \frac{1}{12} \sum_{P=X,Y,Z} (\text{Tr}[P\mathcal{E}(|+P\rangle\langle+P|)] - \text{Tr}[P\mathcal{E}(|-P\rangle\langle-P|)]) + \frac{1}{2}, \quad (24)$$

where we made use of the identities $\text{Tr}[\sigma_P\mathcal{E}(\sigma_P)] = \text{Tr}[P\mathcal{E}(|+P\rangle\langle+P|)] - \text{Tr}[P\mathcal{E}(|-P\rangle\langle-P|)]$ for $P \in \{X, Y, Z\}$, and $\text{Tr}[\mathcal{E}[I]] = 2$. The complete derivation of this formula starting from Eq. (22) can be found in Ref. [39]. In experiment, the right-hand side of Eq. (24) is measured by preparing ± 1 Pauli eigenstates, passing them through the channel $\mathcal{E}$, and then measuring the corresponding Pauli operator, as cartooned in Supp. Fig. 25(a). Such a procedure is applicable to an arbitrary duration t of the channel $\mathcal{E}(t)$.

We focus on the comparison of three different qubits in our system: $\{|0\rangle, |1\rangle\}$, $\{|g\rangle, |e\rangle\}$, and the error-corrected GKP qubit. The free evolution of the two passive qubits is modeled using a composite amplitude damping and white-noise dephasing channel, while the evolution of an error-corrected GKP qubit is modeled using a logical Pauli channel. Given these well-justified assumptions on the error channels, from Eq. (24) we find:

$$\overline{\mathcal{F}}_{\{01\}}(t) = \frac{1}{6}e^{-\gamma_1^c t} + \frac{1}{3}e^{-\gamma_2^c t} + \frac{1}{2}, \quad (25)$$

$$\overline{\mathcal{F}}_{\{ge\}}(t) = \frac{1}{6}e^{-\gamma_1^t t} + \frac{1}{3}e^{-\gamma_{2E}^t t} + \frac{1}{2}, \quad (26)$$

$$\overline{\mathcal{F}}_{\text{GKP}}(t) = \frac{1}{6}e^{-\gamma_X t} + \frac{1}{6}e^{-\gamma_Y t} + \frac{1}{6}e^{-\gamma_Z t} + \frac{1}{2}. \quad (27)$$

We show the time evolution of the fidelity given by Eqs. (25–27) in Supp. Fig. 25(b), using experimentally extracted decay rates at the highest QEC gain measured in our experiment.

As seen above, in general the fidelity decays to its steady-state value in a way that cannot be characterized by a single time constant even in the simplest error models such as Pauli noise or amplitude damping. Therefore, fitting the fidelity decay to a single exponential is not strictly valid, although this heuristic approach was adopted in previous works on bosonic QEC [5, 6, 40]. To avoid such an inconsistency, we consider the channel $\mathcal{E}$ acting for only a short time δt . Any time dependence of the fidelity, even if it contains multiple exponentially decaying contributions, at short times is equivalent to a linear decay:


$$\overline{\mathcal{F}}(\delta t) = 1 - \frac{1}{2}\Gamma \delta t, \quad (28)$$

where Γ is an effective depolarization rate, and $1/\Gamma$ is the fidelity lifetime. For a depolarizing channel $\mathcal{E}_{\text{dep}}(\rho) = (1-p)\rho + p\frac{I}{2}$ with a depolarization probability $p = 1 - e^{-\gamma t}$, we have $\Gamma = \gamma$, motivating the name and the coefficient $1/2$ in Eq. (28). For qubits considered here, the effective depolarization rates are:

$$\Gamma_{\{01\}} = \frac{2\gamma_2^c + \gamma_1^c}{3}, \quad \Gamma_{\{ge\}} = \frac{2\gamma_{2E}^t + \gamma_1^t}{3}, \quad \Gamma_{\text{GKP}} = \frac{\gamma_X + \gamma_Y + \gamma_Z}{3}. \quad (29)$$

We define the coherence gain G as the improvement of the effective depolarization rate of an error-corrected logical qubit over the best physical qubit in the same system (with the break-even point corresponding to $G = 1$). In a bosonic circuit QED system, the latter is typically the $\{|0\rangle, |1\rangle\}$ qubit, hence $G = \Gamma_{\{01\}}/\Gamma_{\text{GKP}}$. The highest gain achieved in our experiment is $G_{\text{max}} = 2.27 \pm 0.07$. During a scan discussed in Section IV J, gain remained above break-even 100% of this week-long time window, with a median of $\overline{G} = 2.0$.

Lastly, we acknowledge that the decay constant of the average channel fidelity is not the only relevant metric that we expect to be correlated with the future ability of such systems to participate in quantum computations. Other metrics, such as the SPAM fidelity and the fidelity of gates, are important as well, and we leave their optimization and detailed characterization for future work.



Supplementary Figure 25. **Average channel fidelity.** (a) Illustration of channel action on the Pauli eigenstates. By linearity, the evolution of these six cardinal points is sufficient to predict the average effect across the whole Bloch sphere. (b) Expected time evolution of average channel fidelity for three different qubits, calculated using experimentally extracted lifetimes of Pauli eigenstates.

- [1] C. Axline, M. Reagor, R. Heeres, P. Reinhold, C. Wang, K. Shain, W. Pfaff, Y. Chu, L. Frunzio, and R. J. Schoelkopf, *An architecture for integrating planar and 3D cQED devices*, Applied Physics Letters **109**, 042601 (2016).
- [2] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, *Fault-tolerant detection of a quantum error*, Science **361**, 266 (2018).
- [3] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, et al., *New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds*, Nature communications **12**, 1 (2021).
- [4] A. P. Read, B. J. Chapman, C. U. Lei, J. C. Curtis, S. Ganjam, L. Krayzman, L. Frunzio, and R. J. Schoelkopf, *Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity*, arXiv:2206.14334 (2022).
- [5] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, et al., *Extending the lifetime of a quantum bit with error correction in superconducting circuits*, Nature **536**, 441 (2016).
- [6] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, C.-L. Zou, et al., *Quantum error correction and universal gate set operation on a binomial bosonic logical qubit*, Nature Physics **15**, 503 (2019).
- [7] K. Serniak, S. Diamond, M. Hays, V. Fatemi, S. Shankar, L. Frunzio, R. Schoelkopf, and M. Devoret, *Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons*, Physical Review Applied **12**, 014052 (2019).
- [8] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, *Model-free quantum control with reinforcement learning*, Physical Review X **12**, 011059 (2022).
- [9] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, et al., *Resolving photon number states in a superconducting circuit*, Nature **445**, 515 (2007).
- [10] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. Girvin, R. J. Schoelkopf, and M. H. Devoret, *Fast universal control of an oscillator with weak dispersive coupling to a qubit*, Nature Physics **18**, 1464 (2022).
- [11] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting*, Physical Review A **74**, 042318 (2006).
- [12] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, et al., *Rapid high-fidelity single-shot dispersive readout of superconducting qubits*, Physical Review Applied **7**, 054020 (2017).
- [13] D. T. McClure, H. Paik, L. S. Bishop, M. Steffen, J. M. Chow, and J. M. Gambetta, *Rapid driven reset of a qubit readout resonator*, Physical Review Applied **5**, 011001 (2016).
- [14] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zayats-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, et al., *Quantum error correction of a qubit encoded in grid states of an oscillator*, Nature **584**, 368 (2020).
- [15] S. Touzard, A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret, *Gated conditional displacement readout of superconducting qubits*, Physical Review Letters **122**, 080502 (2019).
- [16] D. Sank, Z. Chen, M. Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, et al., *Measurement-induced state transitions in a superconducting qubit: beyond the rotating wave approximation*, Physical Review Letters **117**, 190503 (2016).
- [17] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. Lewis, G. Vidal, and A. Blais, *Dynamics of transmon ionization*, Physical Review Applied **18**, 034031 (2022).
- [18] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, et al., *Fluctuations of energy-relaxation times in superconducting qubits*, Physical Review Letters **121**, 90502 (2018).
- [19] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, et al., *Quantum memory with millisecond coherence in circuit QED*, Physical Review B **94**, 014506 (2016).
- [20] B. de Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, *Error correction of a logical grid state qubit by dissipative pumping*, Nature Physics **18**, 296 (2022).

- [21] J. A. Gross, C. M. Caves, G. J. Milburn, and J. Combes, *Qubit models of weak continuous measurements: Markovian conditional and open-system dynamics*, Quantum Science and Technology **3**, 024005 (2018).
- [22] B. Royer, S. Singh, and S. M. Girvin, *Stabilization of finite-energy Gottesman-Kitaev-Preskill states*, Physical Review Letters **125**, 260509 (2020).
- [23] <https://keras.io>.
- [24] https://github.com/v-sivak/tf_quantum_simulator.
- [25] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, *Proximal policy optimization algorithms*, arXiv:1707.06347 (2017).
- [26] S. Guadarrama, A. Korattikara, O. Ramirez, P. Castro, E. Holly, S. Fishman, K. Wang, E. Gonina, N. Wu, E. Kokiopoulou, et al., *TF-Agents: A library for reinforcement learning in tensorflow*, <https://github.com/tensorflow/agents> (2018).
- [27] I. Tzitrin, J. E. Bourassa, N. C. Menicucci, and K. K. Sabapathy, *Progress towards practical qubit computation using approximate Gottesman-Kitaev-Preskill codes*, Physical Review A **101**, 032315 (2020).
- [28] E. Knill and R. Laflamme, *Theory of quantum error-correcting codes*, Phys. Rev. A **55**, 900 (1997).
- [29] D. Gottesman, *Class of quantum error-correcting codes saturating the quantum Hamming bound*, Phys. Rev. A **54**, 1862 (1996).
- [30] K. W. Murch, U. Vool, D. Zhou, S. J. Weber, S. M. Girvin, and I. Siddiqi, *Cavity-assisted quantum bath engineering*, Physical Review Letters **109**, 1 (2012).
- [31] K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, *Demonstrating a driven reset protocol for a superconducting qubit*, Physical Review Letters **110**, 1 (2013).
- [32] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J. C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, et al., *Fast and unconditional all-microwave reset of a superconducting qubit*, Physical Review Letters **121**, 60502 (2018), 1801.07689.
- [33] Z. Chen, J. Kelly, C. Quintana, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Lucero, et al., *Measuring and suppressing quantum state leakage in a superconducting qubit*, Physical Review Letters **116**, 020501 (2016).
- [34] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. Satzinger, C. Quintana, P. V. Klimov, D. Sank, C. Gidney, A. Fowler, et al., *Removing leakage-induced correlated errors in superconducting quantum error correction*, Nature communications **12**, 1 (2021).
- [35] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Deterministically encoding quantum information using 100-photon Schrodinger cat states*, Science **342**, 607 (2013).
- [36] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, *Introduction to quantum noise, measurement, and amplification*, Reviews of Modern Physics **82**, 1155 (2010).
- [37] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, *Dynamics of superconducting qubit relaxation times*, npj Quantum Information **8**, 1 (2022).
- [38] D. Greenbaum, *Introduction to quantum gate set tomography*, arXiv:1509.02921 (2015).
- [39] M. A. Nielsen, *A simple formula for the average gate fidelity of a quantum dynamical operation*, Physics Letters A **303**, 249 (2002).
- [40] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, *Protecting a bosonic qubit with autonomous quantum error correction*, Nature **590**, 243 (2021).