

Real-time quantum error correction beyond break-even

<https://doi.org/10.1038/s41586-023-05782-6>

Received: 16 November 2022

Accepted: 2 February 2023

Published online: 22 March 2023

 Check for updates

V. V. Sivak^{1,2,3,6}✉, A. Eickbusch^{1,2,3}, B. Royer^{1,2,3,4,5}, S. Singh^{1,2,3}, I. Tsioutsios^{1,2,3}, S. Ganjam^{1,2,3}, A. Miano^{1,2,3}, B. L. Brock^{1,2,3}, A. Z. Ding^{1,2,3}, L. Frunzio^{1,2,3}, S. M. Girvin^{1,2,3}, R. J. Schoelkopf^{1,2,3} & M. H. Devoret^{1,2,3}✉

The ambition of harnessing the quantum for computation is at odds with the fundamental phenomenon of decoherence. The purpose of quantum error correction (QEC) is to counteract the natural tendency of a complex system to decohere. This cooperative process, which requires participation of multiple quantum and classical components, creates a special type of dissipation that removes the entropy caused by the errors faster than the rate at which these errors corrupt the stored quantum information. Previous experimental attempts to engineer such a process^{1–7} faced the generation of an excessive number of errors that overwhelmed the error-correcting capability of the process itself. Whether it is practically possible to utilize QEC for extending quantum coherence thus remains an open question. Here we answer it by demonstrating a fully stabilized and error-corrected logical qubit whose quantum coherence is substantially longer than that of all the imperfect quantum components involved in the QEC process, beating the best of them with a coherence gain of $G = 2.27 \pm 0.07$. We achieve this performance by combining innovations in several domains including the fabrication of superconducting quantum circuits and model-free reinforcement learning.

Implementing a single correctable logical qubit requires a physical system with a large state space. It should accommodate the code subspace and its redundant replicas where the logical information will be transferred without distortion when physical errors occur⁸. This redundancy is inextricably associated with an additional operational cost of QEC, known as the control overhead. In the search for an efficient way to alleviate the detrimental effects of the overhead, bosonic codes^{9–11} based on the state space of a harmonic oscillator have been proposed as a promising alternative to the standard approach based on registers of physical qubits^{12–14}. In hybrid architectures, these two approaches are complementary, with qubit-register codes built on logical qubits dynamically protected with efficient base-layer bosonic QEC^{15,16}.

Although some aspects of QEC have been demonstrated with superconducting circuits^{1–7,17}, trapped ions^{18–20} and spins in solid-state systems^{21–23}, the control overhead has prevented current-day experiments from getting to the heart of what QEC promises to achieve—extending the lifetime of quantum information stored in the system. This extension is quantified by the gain G , defined as the ratio between the coherence time of an actively error-corrected logical qubit and the best passive qubit encoding in the same system. The break-even point is reached at $G = 1$. A bosonic cat-code experiment¹⁷ managed to achieve $G = 1.1$, but with a code that continuously shrinks to the vacuum state. Other experiments with various bosonic codes^{1–3} and qubit-register codes^{4–7} have achieved $G = 0.1–0.9$.

We demonstrate full code stabilization and error correction with gain $G = 2.27 \pm 0.07$ using the Gottesman–Kitaev–Preskill (GKP) encoding⁹ of a logical qubit into grid states of an oscillator. The QEC of this code was previously realized in superconducting circuits³ and trapped ions¹⁸. In our work, similarly to ref. ³, the oscillator is an electromagnetic mode of a superconducting cavity whose quantum state is manipulated using a transmon auxiliary qubit (Fig. 1a). Our system has an average relaxation and dephasing time of $\bar{T}_1^t = 280 \mu\text{s}$ and (echo) $\bar{T}_{2E}^t = 240 \mu\text{s}$ for the tantalum-based transmon²⁴, and $\bar{T}_1^c = 610 \mu\text{s}$ and $\bar{T}_2^c = 980 \mu\text{s}$ for the high-purity aluminium cavity²⁵. We implement in this system a ‘trickle-down’ QEC scheme based on the proposals in refs. ^{18,26}, which includes real-time classical processing and measurement-based feedback. We train the QEC circuit parameters in situ with reinforcement learning (RL)^{27–29}, ensuring their adaptation to the real error channels and control imperfections of our system. At peak performance, the achieved lifetimes of logical Pauli eigenstates are $T_X = T_Z = 2.20 \pm 0.03 \text{ ms}$ and $T_Y = 1.36 \pm 0.03 \text{ ms}$, and the logical Pauli error probabilities per QEC cycle are $p_Y = (4.3 \pm 0.4) \times 10^{-4}$ and $p_X = p_Z = (1.81 \pm 0.04) \times 10^{-3}$. With such low logical error probabilities, we explore the QEC process on a previously inaccessible timescale of thousands of cycles, subjecting to scrutiny the standard assumptions of the theory of QEC, such as the stationarity of error rates and absence of leakage-induced correlations. Finally, we carry out error injection experiments to identify the main factors limiting logical performance and chart the path towards the next-generation logical qubit.

¹Department of Physics, Yale University, New Haven, CT, USA. ²Department of Applied Physics, Yale University, New Haven, CT, USA. ³Yale Quantum Institute, Yale University, New Haven, CT, USA. ⁴Institut Quantique, Université de Sherbrooke, Sherbrooke, Quebec, Canada. ⁵Département de Physique, Université de Sherbrooke, Sherbrooke, Quebec, Canada. ⁶Present address: Google AI Quantum, Santa Barbara, CA, USA. ✉e-mail: vladisivak@google.com; michel.devoret@yale.edu

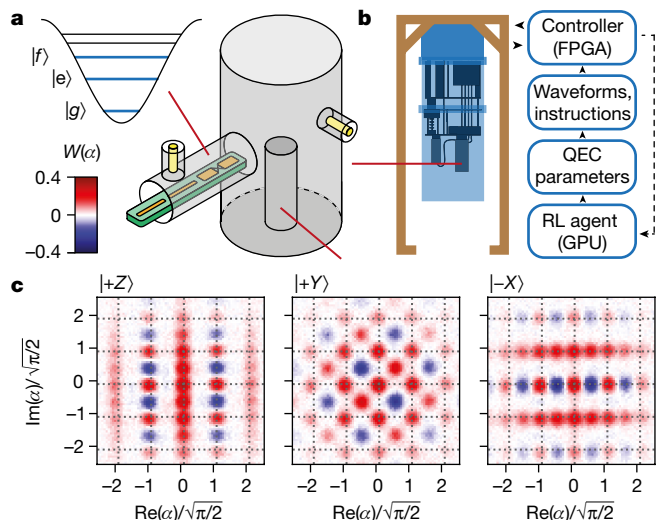


Fig. 1 | Experimental system. **a**, The sample consists of a superconducting aluminium cavity and a sapphire chip with a transmon circuit, readout resonator and Purcell filter. The electromagnetic mode of the cavity implements a harmonic oscillator, and $\{|g\rangle, |e\rangle\}$ levels of the transmon are used as an auxiliary qubit to assist in oscillator QEC. **b**, The sample is cooled in a dilution refrigerator and controlled with microwave and digital electronics. The QEC process is orchestrated by a field-programmable gate array (FPGA), and its parameters are optimized in situ by an RL agent implemented on a graphics processing unit (GPU). **c**, Experimental Wigner functions of the Pauli eigenstates of a grid code with $\Delta = 0.34$ measured after six QEC cycles. Image of the dilution refrigerator was adapted from ref. ³⁷.

Engineering error correction

We now explain the principles of our experiment. Its core idea is to realize an artificial error-correcting dissipation that removes the entropy from the system in an efficient manner by prioritizing the correction of frequent small errors, while not neglecting rare large errors. This idea is illustrated in Fig. 2a for a cartoon system in which redundancy is achieved with only four orthogonal subspaces in total, in which C_0 is the code subspace and C_1 – C_3 are the error subspaces. The error subspaces of our actual infinite-dimensional system are described in Supplementary Section IV.B. In this cartoon example, the standard dissipation scheme, no. 1, is maximally efficient from the perspective of entropy removal, because it corrects any error in a single step. Such an approach is taken by all qubit-register stabilizer codes, for which measurement of the stabilizers, syndrome decoding and recovery, when composed, realize a dissipation channel of high Kraus rank. Although this approach can also be applied to the oscillator grid code (Methods), its implementation entails large control overhead, which in practice might bring more errors than it is designed to correct. By contrast, the trickle-down dissipation scheme, no. 2, has the capacity to correct all of the same errors, but it is not able to do so in a single step. Importantly, the most probable small errors, corresponding to the error space C_1 , are still corrected in a single step. Owing to this simplification, such an approach reduces control overhead in the grid code, and therefore it was adopted in our work. The continuous-time version of approach no. 2 was also demonstrated for other bosonic codes in refs. ^{2,30}.

The stabilizer generators of an ideal square grid code are $S_0^X = D(l_S)$ and $S_0^Z = D(il_S)$, in which $l_S = \sqrt{2\pi}$ is the length of a grid unit cell, and $D(\alpha) = \exp(\alpha a^\dagger - \alpha^* a)$ is the displacement operator for an oscillator with creation and annihilation operators a^\dagger and a . Logical Pauli operators of the ideal code are defined as $X_L = \sqrt{S_0^X}$ and $Z_L = \sqrt{S_0^Z}$. The ideal codewords obey perfect translation symmetry in phase space and thus contain an infinite amount of energy. The finite-energy code is obtained by applying a normalizing envelope operator $N_\Delta = \exp(-\Delta^2 a^\dagger a)$ to the

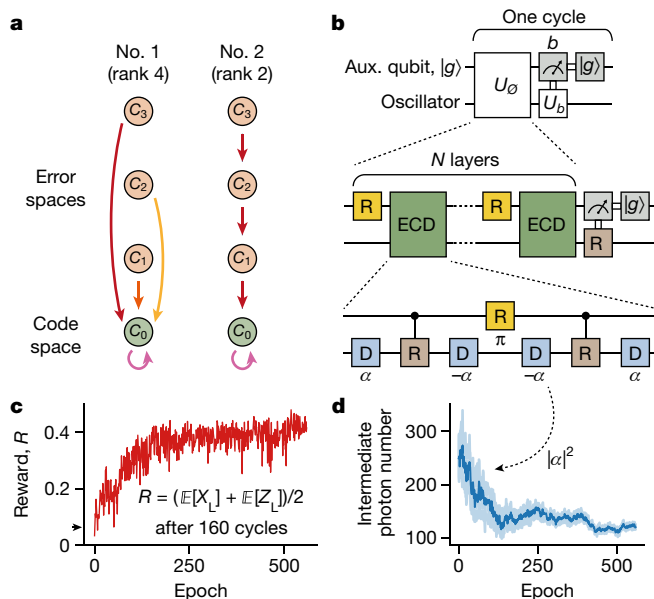


Fig. 2 | QEC implementation and optimization. **a**, Cartoon comparison of error-correcting dissipation channels. The standard dissipation scheme, no. 1, corrects any error in a single step, whereas the ‘trickle-down’ dissipation approach, no. 2, can be viewed as directional hopping between error spaces that eventually brings the quantum state to the code space C_0 . The colours of the arrows correspond to unique Kraus operators, whose number is equal to the channel rank. Higher-rank dissipation removes entropy more efficiently, but incurs larger control overhead. **b**, Implementation of a general rank-2 channel on the oscillator using a single auxiliary (aux.) qubit. The unitary U_θ is approximated as a parametrized circuit consisting of N layers of qubit rotations and oscillator conditional displacements. Each conditional displacement gate utilizes a large intermediate displacement of magnitude $|\alpha|$ to enhance the gate speed. **c**, Evolution of reward of the RL agent during the training. The black arrow indicates the start performance based on independent calibrations. Expectations (denoted with \mathbb{E}) of Pauli operators are taken in their respective eigenstates and include state preparation and measurement (SPAM) errors. **d**, One realization of the learning trajectory of the intermediate photon number used to execute the big conditional displacement gate (‘B’ in the SBS circuit). Light blue shade shows the variance of the sampled parameter values during the training, and dark blue line shows the mean.

ideal codewords, in which Δ parametrizes the code family that approaches the ideal code in the $\Delta \rightarrow 0$ limit. In phase space, this parameter controls the extent of the codewords and the squeezing of their probability peaks. Our experimental Wigner functions of the codewords with $\Delta = 0.34$ are shown in Fig. 1c. The operators of the finite-energy code are obtained through the similarity transformation induced by the envelope operator²⁶ (for example, $S_\Delta^{X/Z} = N_\Delta S_0^{X/Z} N_\Delta^{-1}$).

To realize an error-correcting dissipation channel \mathcal{R}_Δ for the finite-energy code, there is at our disposal a single auxiliary qubit and a classical controller. In principle, with such resources, it is possible to implement arbitrary quantum channels of Kraus rank 2^M by recycling the auxiliary qubit M times and using feedback operations conditioned on the state of the classical M -bit memory of the controller^{31,32}. Here we construct a rank-4 error correction channel as a composition of two rank-2 dissipators $\mathcal{R}_\Delta = \mathcal{R}_\Delta^X \circ \mathcal{R}_\Delta^Z$ that drive the system towards the +1 eigenspace of the finite-energy code stabilizers $S_\Delta^{X/Z}$. A general rank-2 dissipation can be implemented as a unitary U_θ that entangles the system with the auxiliary qubit, followed by a projective measurement of the auxiliary qubit with outcome b and a classically conditioned unitary U_b (Fig. 2b).

In our experiment, any unitary is compiled down to a set of primitive operations: qubit rotations around any equatorial axis $R_\varphi(\theta) = \exp[-i(\theta/2)(\cos\varphi \sigma_x + \sin\varphi \sigma_y)]$ implemented as 32-ns Gaussian pulses

with spectral corrections³³; oscillator displacements $D(\alpha)$ implemented as 40-ns Gaussian pulses; relatively slow conditional rotations $\text{CR}(\theta) = \exp(i\theta\sigma_z a^\dagger a)$ implemented by waiting a certain amount of time under the dispersive coupling Hamiltonian $H_d/\hbar = \chi\sigma_z a^\dagger a/2$, with $\chi = 2\pi \times 46.5$ kHz; and virtual oscillator rotations $R_V(\vartheta) = \exp(i\vartheta a^\dagger a)$ implemented dynamically on the field-programmable gate array in 448 ns. These primitives are used to construct a fast echoed conditional displacement gate $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$ as shown in Fig. 2b, whose speed $\partial_t |\beta\rangle = |\alpha|\chi$ is enhanced compared to the native interaction strength χ by a large factor $|\alpha|$ —the magnitude of the intermediate displacement in phase space^{3,34}.

Both rank-2 dissipators are then implemented as follows: the unitary U_ϕ is decomposed as a parametrized circuit consisting of layers of qubit rotations $R_\phi(\theta)$ and entangling $\text{ECD}(\beta)$ gates, whereas the unitary U_b is realized as only a virtual rotation (Fig. 2b). The role of U_b is twofold: to implement switching between \mathcal{R}_Δ^X and \mathcal{R}_Δ^Z by changing the quadrature of the oscillator by $\pi/2$, and to compensate for a spurious rotation due to the always-on dispersive coupling H_d . The role of U_ϕ is to approximate the mapping of the finite-energy stabilizer onto the state of the auxiliary qubit together with autonomous back-action that pushes the state from the error spaces towards the code space. Several ansätze for the decomposition of U_ϕ were proposed in ref. ²⁶. We adopt a modified version of the so-called small–big–small (SBS) protocol, named to reflect the relative amplitudes of the three conditional displacement gates that it contains: $\beta = l_s \times (i\Delta^2/2, 1, i\Delta^2/2)$ (see Supplementary Section IV.C for further details).

A single application of the resulting composite dissipator \mathcal{R}_Δ realizes a QEC cycle; we refer to applications of constituent dissipators $\mathcal{R}_\Delta^{X/Z}$ as even/odd cycles. In our implementation, the duration of a QEC cycle is $t_c = 2 \times 4.924 \mu\text{s}$, which includes execution of unitary gates, measurements of the auxiliary qubit, and real-time processing and decision-making by the controller.

Learning QEC circuit parameters

Although the SBS ansatz and gate calibrations lead to a functioning QEC process, the highest level of performance cannot be achieved with a crude model of the system based on a few independently calibrated parameters—any such model will inevitably contain unrealistic assumptions. Some model inaccuracies and unknown control imperfections can be compensated by closed-loop optimization with direct feedback from the experimental setup. Previously, pulse-level optimization was successfully utilized to improve gate fidelities^{35,36}, but it was never applied to enhance the performance of QEC. Here we apply a real-time RL agent to this task, as illustrated in Fig. 1b. We use the proximal policy optimization algorithm^{28,29}, which was shown in simulations to outperform other approaches when applied to high-dimensional problems with stochastic objectives that arise in quantum control³⁷. We parametrize the QEC circuit with $P = 45$ parameters that include the amplitudes of various primitive pulses in the circuit decomposition, parameters of the auxiliary qubit reset, and so on.

The training episodes begin with dissipative pre-cooling of the oscillator followed by feedback cooling to prepare the system ground state $|g\rangle|0\rangle$ (Methods). Then, a logical Pauli eigenstate $|+X\rangle$ or $|+Z\rangle$ is initialized with a method from ref. ³⁴, and a candidate QEC protocol is run for $T = 160$ cycles. We chose this duration to enhance the signal-to-noise ratio of the reward, similar to the technique used to sample randomized benchmarking cost functions^{35,36}. At the end of the episode, the reward for the RL agent is obtained by measuring the logical Pauli operator X_L or Z_L (depending on the initial state), which provides a proxy for the logical lifetime. This logical measurement is carried out with one-bit phase estimation of the ideal-code Pauli operators^{33,38}, and its fidelity is intrinsically limited to $(1 + e^{-\pi\Delta^2/4})/2$ (ref. ¹⁶). Although there exist methods of logical readout adapted to the finite code envelope^{18,26,39}, we use the phase estimation method to avoid biasing the RL agent

towards a particular finite envelope size and to let it pick the optimal size given the error channels of our system.

By construction, the reward incentivizes the RL agent to find a QEC protocol that leads to the longest logical qubit lifetime. The typical evolution of the average reward during the training is shown in Fig. 2c. The performance level indicated with a black arrow is achieved with independent calibrations of the system and control parameters (Supplementary Section II). The RL agent substantially improves on this baseline performance in two stages: typically, in the first hundred training epochs, the agent corrects large errors in the initial parameter values, and in the subsequent few hundreds of epochs, it fine-tunes the circuit parameters to achieve the highest performance.

Several trends in the learning trajectories showcase the benefits of the model-free RL approach (see Supplementary Section IV.D for further details). Here we highlight only a single illustrative example. In our implementation of the ECD gate, there exists a nontrivial tradeoff between coherent and incoherent errors: the gate can be implemented faster by displacing the oscillator further in phase space (that is, populating it with more intermediate photons), but this makes the gate more susceptible to high-order nonlinear effects³⁴. Moreover, some choices of this intermediate photon number can result in a Stark shift of the auxiliary qubit into resonance with a spurious degree of freedom (for example, a two-level defect⁴⁰). How these tradeoffs translate into logical qubit performance is difficult to model, but the RL agent can learn the optimal value of the large intermediate displacement without a model. As shown in Fig. 2d, it chose to reduce the intermediate photon number, improving the performance of QEC at the cost of a much slower gate.

Observing QEC beyond break-even

After the training is finished, we pick the best-performing QEC circuit for further characterization. Here we focus on the ability of QEC to create a good quantum memory (that is, to convert the effect of passage of time into an identity channel $\mathcal{I}: \rho \rightarrow \rho$ that preserves all qubit states).

A metric quantifying the deviation of any quantum channel \mathcal{E} from the identity is the average channel fidelity, $\overline{\mathcal{F}}[\mathcal{E}] = \int d\psi \langle \psi | \mathcal{E}(|\psi\rangle\langle\psi|) | \psi \rangle$, in which the integral is over the uniform measure on the qubit state space, normalized so that $\int d\psi = 1$. In general, this fidelity decays over time in a nontrivial way, but to leading order it evolves as $\overline{\mathcal{F}}(t) \approx 1 - \frac{1}{2}\Gamma t$, in which the decay rate Γ is equivalent to an average decoherence rate of all pure states on the qubit Bloch sphere. Conveniently, it suffices to average across the six Pauli eigenstates alone⁴¹, leading to an experimental procedure for extracting Γ that can be applied to any kind of qubit irrespective of its error channel. In Fig. 3, we show the results of such an experiment, conducted for three different qubit encodings in our system: the $\{|g\rangle, |e\rangle\}$ subspace of the transmon, the $\{|0\rangle, |1\rangle\}$ subspace of the oscillator, and the grid code of the oscillator (with and without QEC).

Both the $\{|0\rangle, |1\rangle\}$ and $\{|g\rangle, |e\rangle\}$ qubits are subject to amplitude damping and white-noise dephasing channels, captured by their respective T_1 and T_2 times, with a fidelity decay constant given by $\Gamma = (1/T_1 + 2/T_2)/3$. From the perspective of a quantum memory, the best uncorrectable physical qubit in our system is $\{|0\rangle, |1\rangle\}$, shown in Fig. 3b, which achieves $\Gamma_{\{|0\rangle, |1\rangle\}} = (800 \mu\text{s})^{-1}$. The $\{|g\rangle, |e\rangle\}$ qubit, shown for completeness in Fig. 3a, achieves only $\Gamma_{\{|g\rangle, |e\rangle\}} = (250 \mu\text{s})^{-1}$.

Higher excited states of the oscillator have a shorter lifetime due to bosonic enhancement of spontaneous emission. Therefore, as with any QEC code, encoding a qubit using grid states incurs an immediate penalty in the fidelity decay rate. Moreover, this natural decay, shown in Fig. 3c with open circles, takes the grid states outside the logical manifold and eventually towards the vacuum state $|0\rangle$.

Our error-correcting dissipation stabilizes the grid-code manifold and, together with naturally occurring dissipation, leads to a logical Pauli channel within this manifold, with the lifetimes of logical Pauli eigenstates of $T_X = T_Z = 2.20 \pm 0.03$ ms and $T_Y = 1.36 \pm 0.03$ ms. Under the Pauli

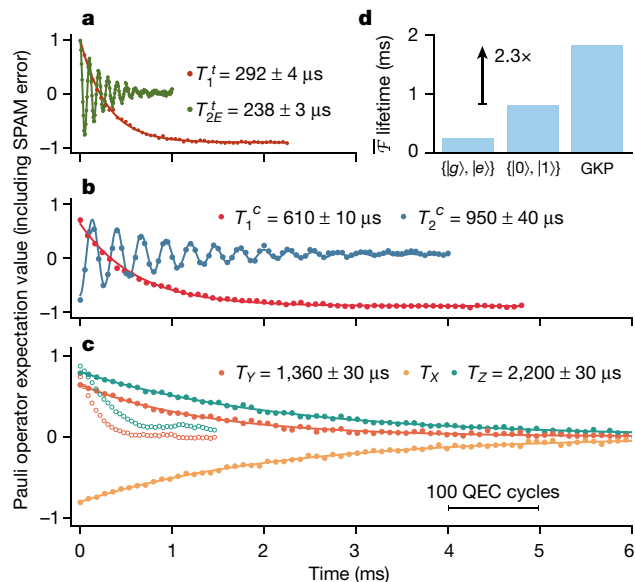


Fig. 3 | System coherence. **a–c**, For the $\{|g\rangle, |e\rangle\}$ qubit (**a**), the $\{|0\rangle, |1\rangle\}$ qubit (**b**) and the GKP qubit (**c**), we initialize Pauli eigenstates, let them evolve freely or under QEC for a variable amount of time, and measure the respective Pauli operators. The data for $\{|g\rangle, |e\rangle\}$ and $\{|0\rangle, |1\rangle\}$ qubits are fitted to amplitude damping and white-noise dephasing channels, and the data for the error-corrected GKP qubit are fitted to a Pauli channel. In **c**, the $|+\rangle$ data are symmetrically reflected with respect to 0 for better visibility. Open circles represent evolution in the absence of QEC, when grid states decay towards vacuum. **d**, Lifetime of average channel fidelity for these three qubits.

channel, the fidelity decay constant is given by $\Gamma = (1/T_x + 1/T_y + 1/T_z)/3$, which in our experiment amounts to $\Gamma_{\text{GKP}} = (1.82 \text{ ms})^{-1}$.

The principal metric characterizing the quality of QEC from the perspective of quantum memory is the coherence gain of an actively error-corrected logical qubit over the best passive qubit encoding. In our experiment, the highest achieved gain is $G = \Gamma_{\{|0\rangle, |1\rangle\}}/\Gamma_{\text{GKP}} = 2.27 \pm 0.07$, confidently beyond break-even.

QEC process characterization

Having characterized the logical qubit as a quantum memory, we next examine the properties of the QEC process. Auxiliary qubit measurement outcomes, referred to as syndromes, inform us which stochastic path the QEC process has taken in each cycle. In Fig. 4a we show a (statistically unrepresentative) sample of these outcomes that comprise trajectories of different experimental shots. Such a dataset contains an immense amount of information about the QEC process, not available in previous experiments with the grid-code QEC^{3,18}.

To interpret this dataset, we adopt here a simplified model of trickle-down dissipation such as depicted in Fig. 2a, which captures the essence of our QEC process. The caveats of this model and the exact Kraus decomposition of our QEC circuit are provided in Supplementary Section IV.B. In this simplified model, the g outcome indicates that the state was projected onto the code space, whereas an e outcome indicates that the state was transferred one level down the error hierarchy, partially or completely correcting an error.

From the dataset in Fig. 4a, we observe that most outcomes are g (green), which means that errors are rare. The stochastic pattern of e outcomes (yellow) reflects randomly occurring errors. Most errors are small and, when corrected, leave single isolated e outcomes. An example syndrome string probably generated by a large error in one quadrature is indicated with an arrow: it has a characteristic $eg/eg/\dots$ pattern. We also observe isolated auxiliary qubit leakage events (red). Leakage to $|f\rangle$ is reset in the same cycle with high probability. Sometimes,

leakage persists for multiple cycles (streak of red), owing to the transmon escaping to a state higher than $|f\rangle$, which is not addressed in our reset scheme.

The average probability of each outcome as a function of time is shown Fig. 4b, in which the process starts from a $|+\rangle$ state. After about 10 cycles of initial state correction, the process settles into a dynamical equilibrium that persists for at least a hundred thousand cycles (the longest measured here) without any notable increase of the error rates over time. Detailed analysis reveals that the QEC process is nearly stationary, with residual deviations from stationarity caused by the transmon leakage to states higher than $|f\rangle$ at a rate 1.3×10^{-4} per cycle (Supplementary Section IV.F).

In this dynamical equilibrium, physical errors excite the quantum state out of the code space with probability $p_{\text{err}} = 0.13 \pm 0.02$ per QEC cycle, as deduced from the statistics of syndrome outcomes. The competition between physical errors and error-correcting dissipation results in a ‘thermal’ distribution across the subspaces with probability $\langle \Pi_0 \rangle = 0.82 \pm 0.02$ of occupying the code space (Methods). Having $p_{\text{err}} \ll 1$ justifies the use of low-rank error-correcting dissipation in our system, which is sufficient to prevent physical errors from accumulating and causing logical errors. At the highest achieved QEC gain, the logical Pauli error probabilities per QEC cycle are $p_y = (4.3 \pm 0.4) \times 10^{-4}$ and $p_x = p_z = (1.81 \pm 0.04) \times 10^{-3}$. By comparing the total logical error probability, $p_x + p_y + p_z$, to the physical error probability, p_{err} , we conclude that 97% of the errors are successfully corrected by our process.

As rare large errors require several cycles to be corrected, the QEC process is weakly time-correlated with a correlation length of 3.9 ± 0.1 cycles (Supplementary Section IV.F). To understand these correlations, in Fig. 4c we inject displacement errors along the position quadrature and monitor the syndromes that they produce as a function of time. Such errors leave traces of e outcomes in proportion to their distance to the closest logical operation. For example, a displacement of length 0, equivalent to a logical identity, leaves no syndrome trace; a displacement of length $l/2$ is close to a logical bit flip of the finite-energy code, and hence it leaves only a small syndrome trace; on the other hand, a midway displacement of length $l/4$ makes a large-distance error that takes the longest time to correct with a low-rank dissipator, generating a lasting trace of e outcomes.

This displacement error injection experiment confirms that errors indeed generate the e syndromes. To verify whether these syndromes herald the occurrence of errors, we carry out post-selection of trajectories with different syndrome patterns. In particular, we discard trajectories that have $\geq d$ consecutive e outcomes in the same-quadrature cycles, with resulting post-selected decay of Pauli eigenstates shown in Fig. 4d. In the case $d = 5$, post-selection eliminates rare large-distance errors and improves the fidelity lifetime only by a factor 1.2, but at the cost of rejection probability of 7×10^{-4} per cycle. On the other hand, in the case $d = 1$, post-selection eliminates relatively frequent small errors that are close to identity, as well as rare large uncorrectable errors that are close to a logical operation. It is because of the latter that the fidelity lifetime in this setting improves by a factor 6.3, but with a more severe rejection probability of 6×10^{-2} per cycle. These favourable post-selection results indicate that such a method can be used for probabilistic preparation of high-fidelity logical states, including the magic states required for universal quantum computing⁴², which is left for future investigation.

Conclusion and outlook

In this work, we used real-time error correction to realize a fully stabilized logical qubit whose lifetime is more than doubled compared to the best passive qubit encoding in the system, marking the transition of QEC from proof-of-principle studies to a practical tool for enhancing quantum memories. Our work improves on previous QEC experiments, which do not protect the logical identity operator I_L (ref. 17), protect only

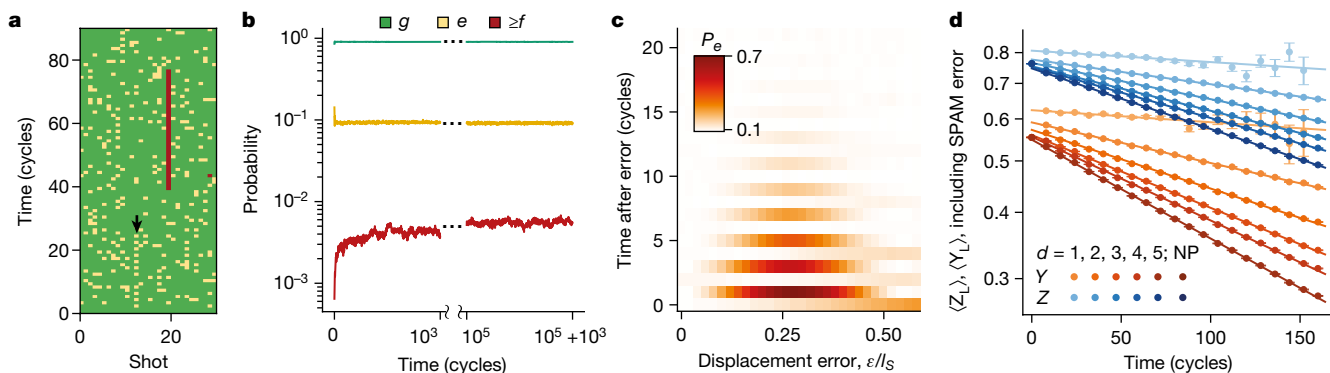


Fig. 4 | Analysis of error syndromes. **a**, A sample of the auxiliary qubit measurement outcomes during QEC. The e outcome (yellow) indicates correction of physical errors. The black arrow points to a syndrome string of the type $eg/eg/\dots$ probably left by a large error in one oscillator quadrature. Red indicates transmon leakage out of the $\{|g\rangle, |e\rangle\}$ qubit subspace. **b**, Average probability of each measurement outcome as a function of time. After correcting state initialization errors, QEC settles into a steady state that persists for at least a hundred thousand cycles. **c**, Probability of e outcome as a function of time after injecting position displacement errors of varying

amplitude. As a logical gate in the ideal code comprises a displacement of amplitude $l_s/2$, a displacement of amplitude $l_s/4$ makes a large-distance error, which takes several cycles to correct with our low-rank QEC channel. **d**, Decay of Pauli eigenstates after eliminating the experimental shots with strings of $\geq d$ consecutive e outcomes in the same-quadrature cycles. Error bars represent statistical uncertainty from finite sampling. Data for X_L eigenstates are not shown; they are expected to behave similarly to those for Z_L . The improvement in lifetime indicates that e outcomes are indeed correlated with occurrence of errors. NP, no post-selection.

one of the logical Pauli operators X_L or Z_L (refs.^{30,43,44}), implement correction in post-processing^{4,5,7}, require post-selection⁴⁵ and do not reach break-even¹⁻⁷. Instrumental for this achievement, among other factors, was the adoption of a model-free learning framework, improved fabrication techniques for the transmon auxiliary qubit and a new grid-code QEC protocol.

Carrying out additional experiments, we identified the core challenges that need to be addressed to ensure future progress of grid-code QEC. In particular, by studying long-time system stability, we found that occasional collapses of the logical performance are strongly correlated with appearance of spurious degrees of freedom in the system. Their resonant interaction with the Stark-shifted transmon qubit degrades the fidelity of our operations (Supplementary Section IV.). In the short term, this effect could be mitigated by adopting a tunable auxiliary qubit and periodically re-training the QEC circuit to find better spectral locations. In the long term, the behaviour of these defects needs to be understood, as they pose even greater danger for scaled-up quantum devices^{4,5,7}.

In addition, we expect that considerable enhancement can be gained by tailoring the QEC process not only to error channels of the oscillator, but also to those of the auxiliary qubit. Our QEC circuit is fault-tolerant with respect to auxiliary qubit phase-flip errors by design²⁶. With the transmon qubit used here, the sensitivity of the logical lifetime to auxiliary qubit phase flips is 65 times smaller than to auxiliary qubit bit flips, as found with noise injection experiments (Supplementary Section IV.I). Future development should incorporate robustness against auxiliary qubit bit flips, either through path-independent control^{46,47} or by adopting an auxiliary qubit with biased noise⁴⁸.

Note added in proof: In parallel with our work, a gain $G = 1.16$ was demonstrated with a binomial code realized in a microwave cavity⁴⁹.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-023-05782-6>.

- Hu, L. et al. Quantum error correction and universal gate set operation on a binomial bosonic logical qubit. *Nat. Phys.* **15**, 503–508 (2019).
- Gertler, J. M. et al. Protecting a bosonic qubit with autonomous quantum error correction. *Nature* **590**, 243–248 (2021).

- Campagne-Ibarcq, P. et al. Quantum error correction of a qubit encoded in grid states of an oscillator. *Nature* **584**, 368–372 (2020).
- Krinner, S. et al. Realizing repeated quantum error correction in a distance-three surface code. *Nature* **605**, 669–674 (2022).
- Zhao, Y. et al. Realization of an error-correcting surface code with superconducting qubits. *Phys. Rev. Lett.* **129**, 030501 (2022).
- Sundaresan, N. et al. Matching and maximum likelihood decoding of a multi-round subsystem quantum error correction experiment. Preprint at <https://arxiv.org/abs/2203.07205> (2022).
- Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* **614**, 676–681 (2023).
- Knill, E. & Laflamme, R. Theory of quantum error-correcting codes. *Phys. Rev. A* **55**, 900–911 (1997).
- Gottesman, D., Kitaev, A. & Preskill, J. Encoding a qubit in an oscillator. *Phys. Rev. A* **64**, 012310 (2001).
- Mirrahimi, M. et al. Dynamically protected cat-qubits: a new paradigm for universal quantum computation. *N. J. Phys.* **16**, 045014 (2014).
- Michael, M. H. et al. New class of quantum error-correcting codes for a bosonic mode. *Phys. Rev. X* **6**, 031006 (2016).
- Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* **52**, R2493–R2496 (1995).
- Steane, A. M. Error correcting codes in quantum theory. *Phys. Rev. Lett.* **77**, 793–797 (1996).
- Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).
- Noh, K. & Chamberland, C. Fault-tolerant bosonic quantum error correction with the surface–Gottesman–Kitaev–Preskill code. *Phys. Rev. A* **101**, 012316 (2020).
- Terhal, B. M., Conrad, J. & Vuillot, C. Towards scalable bosonic quantum error correction. *Quantum Sci. Technol.* **5**, 043001 (2020).
- Ofek, N. et al. Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature* **536**, 441–445 (2016).
- de Neeve, B., Nguyen, T.-L., Behrle, T. & Home, J. P. Error correction of a logical grid state qubit by dissipative pumping. *Nat. Phys.* **18**, 296–300 (2022).
- Ryan-Anderson, C. et al. Realization of real-time fault-tolerant quantum error correction. *Phys. Rev. X* **11**, 041058 (2021).
- Egan, L. et al. Fault-tolerant control of an error-corrected qubit. *Nature* **598**, 281–286 (2021).
- Waldherr, G. et al. Quantum error correction in a solid-state hybrid spin register. *Nature* **506**, 204–207 (2014).
- Abobeih, M. et al. Fault-tolerant operation of a logical qubit in a diamond quantum processor. *Nature* **606**, 884–889 (2022).
- Xue, X. et al. Quantum logic with spin qubits crossing the surface code threshold. *Nature* **601**, 343–347 (2022).
- Place, A. P. et al. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nat. Commun.* **12**, 1779 (2021).
- Reagor, M. et al. Quantum memory with millisecond coherence in circuit QED. *Phys. Rev. B* **94**, 014506 (2016).
- Royer, B., Singh, S. & Girvin, S. M. Stabilization of finite-energy Gottesman–Kitaev–Preskill states. *Phys. Rev. Lett.* **125**, 260509 (2020).
- Sutton, R. S. & Barto, A. G. *Reinforcement Learning: an Introduction* (MIT Press, 2018).
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A. & Klimov, O. Proximal policy optimization algorithms. Preprint at <https://arxiv.org/abs/1707.06347> (2017).
- Guadarrama, S. et al. TF-Agents: a library for reinforcement learning in TensorFlow. GitHub <https://github.com/tensorflow/agents> (2018).
- Lescanne, R. et al. Exponential suppression of bit-flips in a qubit encoded in an oscillator. *Nat. Phys.* **16**, 509–513 (2020).
- Lloyd, S. & Viola, L. Engineering quantum dynamics. *Phys. Rev. A* **65**, 010101 (2001).

32. Shen, C. et al. Quantum channel construction with circuit quantum electrodynamics. *Phys. Rev. B* **95**, 134501 (2017).
33. Chen, Z. et al. Measuring and suppressing quantum state leakage in a superconducting qubit. *Phys. Rev. Lett.* **116**, 020501 (2016).
34. Eickbusch, A. et al. Fast universal control of an oscillator with weak dispersive coupling to a qubit. *Nat. Phys.* **18**, 1464–1469 (2022).
35. Kelly, J. et al. Optimal quantum control using randomized benchmarking. *Phys. Rev. Lett.* **112**, 240504 (2014).
36. Werninghaus, M. et al. Leakage reduction in fast superconducting qubit gates via optimal control. *npj Quantum Inf.* **7**, 14 (2021).
37. Sivak, V. V. et al. Model-free quantum control with reinforcement learning. *Phys. Rev. X* **12**, 011059 (2022).
38. Terhal, B. M. & Weigand, D. Encoding a qubit into a cavity mode in circuit QED using phase estimation. *Phys. Rev. A* **93**, 012315 (2016).
39. Hastrup, J. & Andersen, U. L. Improved readout of qubit-coupled Gottesman-Kitaev-Preskill states. *Quantum Sci. Technol.* **6**, 035016 (2021).
40. Klimov, P. V. et al. Fluctuations of energy-relaxation times in superconducting qubits. *Phys. Rev. Lett.* **121**, 90502 (2018).
41. Nielsen, M. A. A simple formula for the average gate fidelity of a quantum dynamical operation. *Phys. Lett. A* **303**, 249–252 (2002).
42. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* **71**, 022316 (2005).
43. Grimm, A. et al. Stabilization and operation of a Kerr-cat qubit. *Nature* **584**, 205–209 (2020).
44. Chen, Z. et al. Exponential suppression of bit or phase errors with cyclic error correction. *Nature* **595**, 383–387 (2021).
45. Andersen, C. K. et al. Repeated quantum error detection in a surface code. *Nat. Phys.* **16**, 875–880 (2020).
46. Ma, W.-L. et al. Path-independent quantum gates with noisy ancilla. *Phys. Rev. Lett.* **125**, 110503 (2020).
47. Rosenblum, S. et al. Fault-tolerant detection of a quantum error. *Science* **361**, 266–270 (2018).
48. Puri, S. et al. Stabilized cat in a driven nonlinear cavity: a fault-tolerant error syndrome detector. *Phys. Rev. X* **9**, 041009 (2019).
49. Ni, Z. et al. Beating the break-even point with a discrete variable-encoded logical qubit. *Nature* <https://doi.org/10.1038/s41586-023-05784-4> (2023).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2023

Methods

QEC of the ideal grid code

To understand the error-correcting properties of the ideal code, consider an error channel \mathcal{E} decomposed in the displacement basis. An ideal grid code with code projector Π_0 satisfies the Knill–Laflamme conditions⁸ $\Pi_0 D^\dagger(\varepsilon_\alpha) D(\varepsilon_\beta) \Pi_0 \propto \delta(\varepsilon_\alpha - \varepsilon_\beta) \Pi_0$ for all errors in a correctable set $E_+ = \{D(\varepsilon) : |\operatorname{Re}(\varepsilon)|, |\operatorname{Im}(\varepsilon)| < l_S/4\}$. A displacement error of amplitude ε creates an error state $|\psi_\varepsilon\rangle = D(\varepsilon)|\psi\rangle$, in which $|\psi\rangle$ is any state from the code space. As a displaced grid state is still translationally invariant, it remains an eigenstate of the ideal-code stabilizers, and the phase of its eigenvalue encodes a continuous error syndrome: $S_0^Z |\psi_\varepsilon\rangle = \exp(2il_S \operatorname{Re}[\varepsilon]) |\psi_\varepsilon\rangle$ and $S_0^X |\psi_\varepsilon\rangle = \exp(-2il_S \operatorname{Im}[\varepsilon]) |\psi_\varepsilon\rangle$. Error correction of an ideal grid code can be carried out in a similar manner to that for any stabilizer code: first, measure the stabilizers to obtain the error syndrome, which here corresponds to phase estimation of $S_0^{X/Z}$ that yields the error amplitude ε . This step projects the state onto one of the orthogonal error spaces. Then, apply the recovery operation, here a simple displacement $D(-\varepsilon)$, to correct the error. This procedure realizes an artificial dissipation \mathcal{R} of an infinite rank that corrects any error from E_+ in a single cycle, $(\mathcal{R} \circ \mathcal{E})(\rho) \propto \rho$, analogously to the cartoon high-rank dissipation in Fig. 2a. In contrast to this approach, our experiment realizes low-rank dissipation that asymptotically satisfies $([\mathcal{R}]^{n \rightarrow \infty} \circ \mathcal{E})(\rho) \propto \rho$.

Dissipative cooling to vacuum

We utilize the dissipation engineering framework⁵⁰ to design fast cooling of the oscillator to the vacuum state in the weak-coupling regime for which previous known cooling methods⁵¹ fail. We also expect this new method to be applicable to cooling of trapped ions, for which conditional displacement can be realized through sideband driving, and auxiliary qubit reset through internal state repumping³⁸. As in error-correcting dissipation, we realize this cooling as a composition of two rank-2 channels that shrink the oscillator state in the orthogonal quadratures. The unitary $U_\mathcal{O}$ in this case is realized as a three-layer circuit obtained from the first-order Trotter decomposition of $U = \exp[-i\varepsilon(a\sigma_+ + a^\dagger\sigma_-)]$, in which $\varepsilon \ll 1$ controls the cooling rate. This unitary swaps the excitations of the oscillator into the auxiliary qubit, which is reset in every cycle. The duration of one full cooling cycle (including both quadratures) is $t_c = 2 \times 3.38 \mu\text{s}$. With $\varepsilon = 0.4$, we achieve cooling at a rate 20 times faster than the natural energy damping rate of the oscillator. In our experiment, 25 full cycles of such a dissipative cooling are then followed with a feedback cooling protocol adapted from ref.¹⁷ to remove any residual thermal population. See Supplementary Section II.F for more details.

RL implementation

The QEC circuit is parametrized with a vector \mathbf{p} . Instead of optimizing \mathbf{p} directly, the RL agent learns parameters of the probability distribution from which \mathbf{p} is stochastically sampled during the training to ensure adequate exploration of parameter space. To this end, we use a factorized multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ with mean $\boldsymbol{\mu}$ and covariance matrix $\operatorname{diag}[\boldsymbol{\sigma}]^2$. To capture the pattern of relations between different components of \mathbf{p} , the mean and covariance are represented as parametrized functions $\boldsymbol{\mu}(\theta)$ and $\boldsymbol{\sigma}(\theta)$ of common hidden variables θ . In this work, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are produced at the output of a neural network with two fully connected layers of 50 and 20 rectifier linear unit neurons. Starting with the initial vector of parameters $\boldsymbol{\mu}_i$ found with independent calibrations, during the course of learning the agent gradually deforms the distribution and localizes it on the new vector $\boldsymbol{\mu}_f$, the final result of the optimization. Typically, as it proceeds, the agent also reduces the entropy of the distribution to have a finer control over the mean. These features of learning are observed in the example evolution of one component of \mathbf{p} in Fig. 2d. During one training epoch, we evaluate 10 QEC circuit candidates with 300 episodes (that

is, experimental shots) per candidate. The collected information is used to update the neural network parameters θ according to the proximal policy optimization algorithm, which completes the epoch. One epoch takes approximately 16 s, with most time spent on recompilation of instruction sequences for the field-programmable gate array, and its reinitialization. See Supplementary Section III.B for more details.

Steady state of the QEC process

We carry out Wigner tomography of the logical states after a varying duration of the QEC process, reconstruct the density matrix, and from its spectral decomposition extract the expectation value of the code projector $\langle \Pi_0 \rangle = 0.825 \pm 0.003$, in which the uncertainty represents the standard deviation with respect to different process durations of 100, 200, 400 and 800 cycles. In addition to the code space, only one error space is occupied in the steady state with an appreciable probability of 0.170 ± 0.005 . The logical decoherence within this error space happens at the same rate as within the code space. For more details, see Supplementary Section IV.H.

The expectation value of the code projector in the steady state can be estimated independently, using the statistics of syndrome outcomes. Under the approximations discussed in Supplementary Section IV.E, the probability that a syndrome string of length $2n$ consists only of g outcomes asymptotically approaches $(\langle \Pi_0 \rangle (1 - p_{\text{err}}))^{n-1}$ for large n . Using this method, we extract $\langle \Pi_0 \rangle = 0.81 \pm 0.02$ and $p_{\text{err}} = 0.13 \pm 0.02$. The uncertainty in this case represents the inaccuracy of the model for the string probability, which is valid to first order in p_{err} . The value of $\langle \Pi_0 \rangle$ quoted in the main text is the average of the two methods. Constructing a detailed error budget of the aggregate error probability p_{err} based on the system-level simulation of the known error processes is an avenue left for future work.

Data availability

The data that support the findings of this study are available from the corresponding authors upon a request.

Code availability

The open-source implementation of the proximal policy optimization algorithm is available in ref.²⁹. The custom code used for quantum control optimization, data acquisition, analysis and visualization is available from the corresponding authors upon a request.

50. Gross, J. A., Caves, C. M., Milburn, G. J. & Combes, J. Qubit models of weak continuous measurements: Markovian conditional and open-system dynamics. *Quantum Sci. Technol.* **3**, 024005 (2018).
51. Pfaff, W. et al. Controlled release of multiphoton quantum states from a microwave cavity memory. *Nat. Phys.* **13**, 882–887 (2017).

Acknowledgements We acknowledge discussions with R. Cortiñas, J. Claes and A. Mi. We thank B. Huard for feedback on the manuscript. This research was supported by the US Army Research Office under grants W911NF-18-1-0212 and W911NF-16-1-0349, and by the US Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under contract number DE-SC0012704. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The use of fabrication facilities was supported by the Yale Institute for Nanoscience and Quantum Engineering and the Yale SEAS Cleanroom.

Author contributions V.V.S., A.M. and A.E. built the experimental setup. R.J.S. contributed to experimental apparatus. I.T., S.G. and L.F. fabricated the transmon chip. B.R., S.S. and S.M.G. developed the theory. B.R., V.V.S., A.E. and B.L.B. developed dissipative oscillator cooling. A.E., V.V.S. and A.Z.D. developed the state initialization technique. V.V.S. implemented RL, carried out the experiments and analysed data. V.V.S., A.E., B.R. and M.H.D. regularly discussed the project and provided insight. M.H.D. supervised the project. V.V.S. and M.H.D. wrote the manuscript with feedback from all authors.

Article

Competing interests R.J.S., L.F. and M.H.D. are founders, and R.J.S. and L.F. are shareholders of Quantum Circuits, Inc.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-023-05782-6>.

Correspondence and requests for materials should be addressed to V. V. Sivak or M. H. Devoret.

Peer review information *Nature* thanks Atsushi Noguchi and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.