Demonstrating Efficient and Robust Bosonic State Reconstruction via Optimized Excitation Counting

Tanjung Krisnanda,[1,†] Clara Yun Fontaine,[1,†] Adrian Copetudo,[1,†] Pengtao Song,[1,†] Kai Xiang Lee,[2] Ni-Ni Huang,[1] Fernando Valadares,[1] Timothy C.H. Liew,[2] and Yvonne Y. Gao[1,3,*]

[1] *Centre for Quantum Technologies, National University of Singapore, Singapore 117543, Singapore*
[2] *School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371, Singapore*
[3] *Department of Physics, National University of Singapore, Singapore 117542, Singapore*

(Received 22 March 2024; revised 29 November 2024; accepted 3 December 2024; published 6 January 2025)

Quantum state reconstruction is an essential element in quantum information processing. However, efficient and reliable reconstruction of nontrivial quantum states in the presence of hardware imperfections can be challenging. This task is particularly demanding for high-dimensional states encoded in continuous-variable (CV) systems, where a large number of grid-based measurements are often used to adequately sample relevant regions of phase space. In this work, we introduce an efficient and robust technique of optimized reconstruction with excitation-number sampling (ORENS) based on the idea of the generalized Q function. We use a standard bosonic circuit quantum electrodynamics (cQED) setup to experimentally demonstrate effective state reconstruction using the theoretically minimum number of measurements. Our investigation highlights that ORENS is naturally free of parasitic system dynamics and resilient to decoherence effects in the hardware, enabling it to outperform the conventional reconstruction techniques in cQED such as Wigner tomography. Finally, ORENS relies only on the ability to accurately measure the excitation number of a given CV state, making it a versatile and accessible tool for a wide range of CV platforms and readily scalable to multimode systems. Thus, our work provides a crucial and valuable primitive for practical quantum information processing using bosonic modes.

DOI: 10.1103/PRXQuantum.6.010303

I. INTRODUCTION

Continuous-variable (CV) quantum systems offer the rich and versatile dynamics of a large Hilbert space [1–4], with applications ranging across quantum computation [5,6], metrology [7], and simulation [8]. To take full advantage of these systems, it is essential to develop techniques to accurately characterize the properties, interactions, and evolutions of their quantum states. However, reconstructing the density matrix of an arbitrary CV state in a large Hilbert space is a challenging task. Not only are many measurement observables needed to capture features spread across the large phase space but, experimentally, the observables must often be mapped to an auxiliary element (e.g., a qubit) via nonideal and error-prone operations to extract the relevant measurement outcomes.

Current techniques used in CV quantum information platforms include the Wigner [9,10], Husimi Q [11], and generalized Q [12,13] function measurements. They rely on over-sampling the phase-space features of the CV state to achieve high reconstruction accuracy. However, as these quantum systems advance in both scale and complexity, it is necessary to design more optimal reconstruction techniques that use the fewest measurements of observables that are resilient against experimental errors. While many different approaches to tackling this scalability challenge have been theoretically proposed or experimentally demonstrated [14–18], these strategies often come at the cost of versatility, measurement quality, engineering convenience, and optimization complexity.

In this work, we present a method to robustly and efficiently reconstruct arbitrary bosonic states with the optimized fewest measurements of the excitation number. Our technique, named optimized reconstruction with excitation number sampling (ORENS), augments the classic idea of the generalized Q function with a carefully designed optimization procedure to allow effective reconstruction of arbitrary CV states with minimal measurements. For instance, compared to an earlier work [13] that reconstructs states within the truncation dimension $D = 8$, ORENS

*Contact author: yvonne.gao@nus.edu.sg
†These authors contributed equally to this work.

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

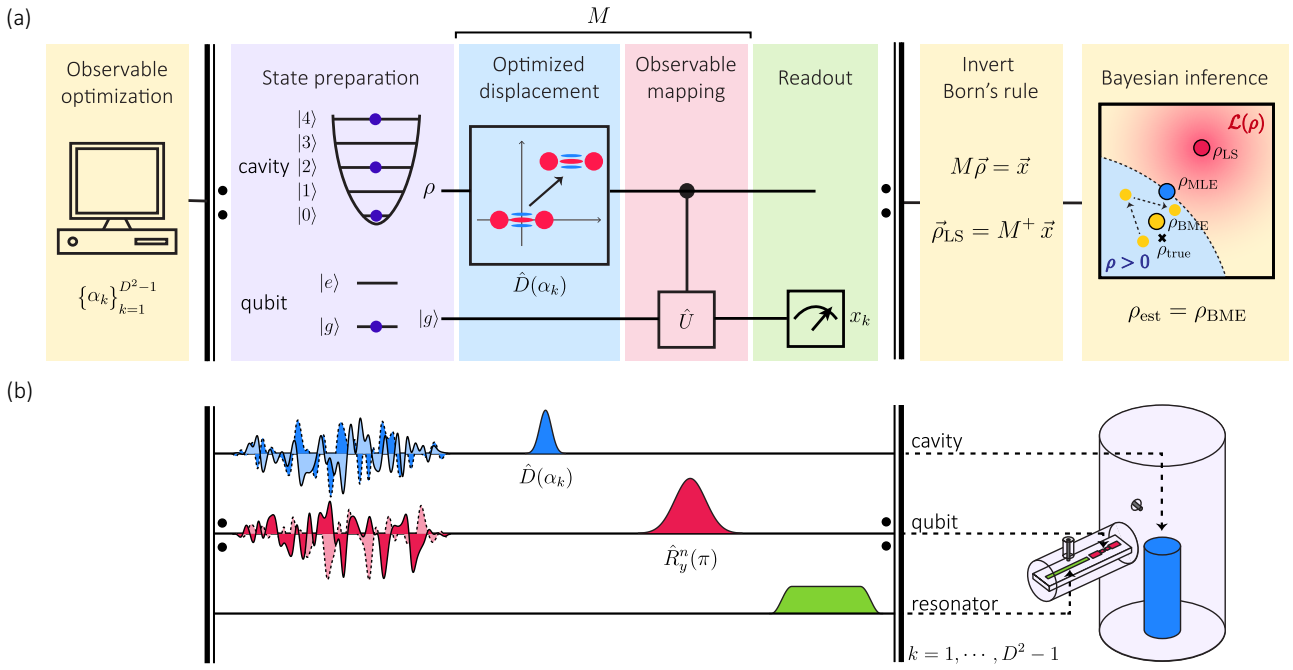


FIG. 1. An overview of bosonic state reconstruction and its translation to ORENS in cQED. (a) To reconstruct an arbitrary CV state ρ of dimension D, measurements of at least $D^2 - 1 := N_{\rm obs}^*$ independent real parameters must be made to be informationally complete. In bosonic systems, this is typically achieved through displacements, $\hat{D}(\alpha_k)$, to $N_{\rm obs}^*$ distinct points in phase space, which effectively changes the measurement observable. These displacement points, $\{\alpha_k\}$ are classically optimized before the experiment for a specific truncation dimension, without any knowledge of the state itself. A Bayesian-inference algorithm is applied to compute the optimal estimator $\rho_{\rm BME}$ based on the resulting measurement outcomes [27]. (b) A prototypical bosonic cQED device consists of a storage cavity (blue), an ancillary qubit (red), and a readout resonator (green). State preparation is implemented with numerically optimized pulses played to the cavity and qubit. Displacements on the cavity state are implemented by driving the cavity at its resonance frequency using a microwave pulse with a Gaussian envelope. The excitation-number sampling for ORENS is enacted by a π pulse on the qubit conditioned on a specific photon number, n, in the bosonic mode. The optimal choice of n is made before the experiment, based only on the truncation dimension. The measurement outcome is extracted with the standard dispersive readout of the resonator.

offers a 50-fold reduction in the number of measurement observables without compromising reconstruction accuracy. The technique can be readily implemented across CV platforms, where these excitations take the form of optical photons [19], microwave photons [20–23], and motional phonons of trapped ions [24–26]. We experimentally demonstrate the efficacy of ORENS in a cQED platform to showcase its performance for high-dimensional CV states, even under severe decoherence. We show that ORENS outperforms the state-of-the-art Wigner reconstruction technique [9,10], owing to its inherent robustness against both coherent and incoherent errors. Our contribution to bosonic state reconstruction, a key research pillar in CV applications, will reinforce the development and analysis of more complex bosonic states and dynamics across different physical devices.

II. ORENS PROTOCOL

Conceptually, reconstructing an arbitrary quantum state ρ consists of measuring identically prepared copies along many different bases. To accurately obtain information about the state, these bases must be informationally complete and the measurements have to be resilient against errors. For CV systems, when the state does not extend beyond a certain dimension D, its Hilbert space can be truncated [28,29]. As such, only $D^2 - 1 := N_{\rm obs}^*$ independent real parameters must be obtained from measurements for informational completeness, setting the optimal number of measurements for efficient reconstruction [27]. To achieve independence between measurements for CV systems, distinct displacement transformations $\hat{D}(\alpha) = \exp\left(\alpha\hat{c}^\dagger + \alpha^*\hat{c}\right)$ can be conveniently applied on ρ to sample different regions of phase space [28]. Upon choosing a base observable and a set of displacements, the set of measurements is written as a matrix M and the measurement process is described with Born's rule as $\vec{x} = M\vec{\rho}$, where $\vec{x}$ is the measurement outcomes and $\vec{\rho}$ is the vectorized density matrix. Born's rule can be inverted to find the least-squares estimator $\vec{\rho}_{\rm LS} = M^+\vec{x}$, where M^+ is the Moore-Penrose pseudoinverse of M (Appendix D 1), which is constrained to be physical to realize the final estimator $\rho_{\rm est}$ (Appendix D 3). An overview of the key bosonic-state-reconstruction stages is illustrated in Fig. 1(a).

Ideally, measurements of the minimal N_{obs}^{*} independent real parameters enable perfect reconstruction of ρ. However, with experimental imperfections, the accuracy of the estimated state critically depends on the choice of M, which amplifies measurement errors to varying degrees upon estimation. The robustness to error is characterized by the condition number (CN) of the measurement matrix M, where a CN of 1 corresponds to the absence of error amplification and grants the optimal reconstruction [30].

Our proposed method, ORENS, leverages the sampling of the excitation number across phase space, which is not only a readily accessible measurement observable in many CV platforms but also has built-in robustness against both dephasing and nonideal coherent dynamics. Prior to the experiment, a minimal set of N_{obs}^{*} measurement observables is classically optimized using a standard gradient-descent method for a given truncation dimension D. This provides a set of optimal displacement points and excitation numbers that effectively minimizes the CN (Appendix D 2) and ensures the accuracy of the reconstruction. This classical optimization process only needs to be done once and the resulting set of displacements can be used to reconstruct any arbitrary CV states within the same truncation dimension.

ORENS is conceptually based on the technique of the generalized Q function [12,13,31,32], where $Q_n(\alpha) = \text{Tr}\left(|n\rangle\langle n|\hat{D}(\alpha)^{\dagger}\rho\hat{D}(\alpha)\right)$. This is the generalization of the Husimi-Q function, $Q_0(\alpha) = \langle\alpha|\rho|\alpha\rangle$, to an arbitrary number of excitations n. Sampling higher n overcomes the limitations of Husimi Q by boosting sensitivity to phase-space oscillations of ρ. This can be understood graphically by considering that the Q_n function of a given state ρ is the convolution of the Wigner function of ρ with that of $|n\rangle$ [33]. For the specific case of vacuum, $n = 0$, the Wigner of vacuum is a Gaussian distribution centered in the origin of the phase space and thus acts as a Gaussian filter on ρ, erasing fast phase-space oscillations and resulting in a strictly non-negative version of the Wigner function. However, for larger $n > 0$, the features of ρ are better preserved by Q_n, enabling robust reconstruction.

We validate the resilience and efficiency of ORENS by reconstructing arbitrary CV states in cQED, where the excitation number is synonymous with the photon number. We demonstrate that the excitation-number measurement is inherently free of parasitic system dynamics and is robust against decoherence. In our hardware, illustrated in Fig. 1(b), the CV states are stored in the electromagnetic field of superconducting LC resonators, realized as a three-dimensional (3D) high-Q coaxial cavity machined out of high-purity (4N) aluminum. The states are prepared, transformed, and measured via the engineered dispersive interaction with an auxiliary qubit. The qubit is a standard transmon dispersively coupled to an on-chip readout resonator and both elements are fabricated out of aluminum on a sapphire substrate (Appendix A 2). The full Hamiltonian of this qubit-cavity system can be found in Appendix A 4.

III. EXCITATION COUNTING IN CQED

Extracting the excitation number of the cavity consists of conditionally exciting the qubit depending on the number of excitations in the cavity. This conditional excitation leverages the dispersive interaction between the cavity and qubit, which can be understood by looking at the system Hamiltonian in the rotating frame of the qubit drive,

$$\frac{\hat{H}}{\hbar} = \frac{\Omega_d}{2}\hat{\sigma}_y + \Delta|e\rangle\langle e| - \chi\hat{n}|e\rangle\langle e|, \tag{1}$$

where $\hat{n}$ is the excitation-number operator of the cavity mode, χ is the dispersive coupling between the cavity and the qubit, $\hat{\sigma}_y$ is the $\hat{y}$ Pauli operator, Δ is the drive detuning from the qubit frequency, and Ω_d is the drive amplitude. The first term in Eq. (1) yields a rotation of the qubit state around the y axis of the qubit Bloch sphere, where the rotation angle depends on the drive amplitude Ω_d and the pulse length. A π rotation corresponds to a complete flip of the qubit state. By using a long spectrally selective drive with duration denoted as t_π, such that $t_\pi > 1/\chi$, the individual shifts of the qubit frequency $(\omega_q - n\chi)$ corresponding to n excitations in the cavity, are resolved and can be individually addressed [22]. Choosing a drive detuning $\Delta = \chi n$ enables robust mapping of the excitation number of the cavity to the qubit:

$$p_n = \text{Tr}\,(\rho|n\rangle\langle n|) \approx p_e(\Delta = \chi n), \tag{2}$$

where p_n is the probability of the cavity having n excitations and p_e is the probability of the qubit being in the excited state. More details can be found in Appendix B 1.

By design, this mapping of the excitation number does not experience significant parasitic Hamiltonian dynamics, making it an excellent choice of measurement observable for state reconstruction. To verify this, we prepare a given Fock state $|n\rangle$ in the cavity using numerically optimized gradient-ascent pulse engineering (GRAPE) pulses [34]. We flip the state of the qubit with a $t_\pi = 1\ \mu$s Gaussian pulse selective at a frequency $\omega_q - \chi n$ to map the excitation number corresponding to $|n\rangle$ to the state of the qubit, which is read out with a single-shot measurement. We repeat the experiment 1000 times and use the average outcomes to estimate $Q_n(0)$. The results, shown in Fig. 2(a), demonstrate effective mapping of the excitation number from the cavity to the qubit state with $p_n > 0.93$. This is in excellent agreement with the simulated outcomes, where the effects of qubit decoherence, qubit thermal population, and finite readout discrimination are taken into account.

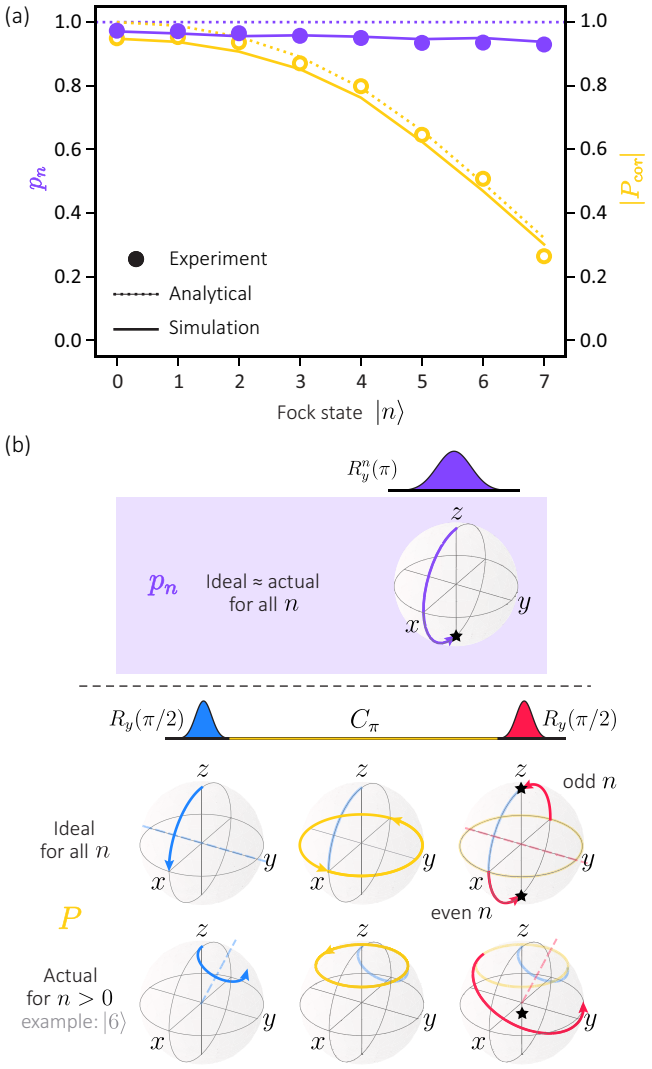


FIG. 2. The excitation-number and parity mapping of Fock states in cQED. (a) The measurement outcomes of excitation number p_n (purple) and the absolute value of the corrected parity $|P_{\text{cor}}|$ (yellow) as a function of the Fock state $|n\rangle$ prepared in the cavity. The data are obtained with qubit T_ϕ of 15.3 μs. They show good agreement with analytical [Eqs. (2) and (3)] and simulated trends based on real device parameters. (b) Bloch spheres of the qubit state at each step of the excitation number (top) and parity (bottom) mapping. The measurement outcomes are marked with black stars.

In contrast to excitation-number measurement, the cavity parity measurement commonly used in the ubiquitous Wigner reconstruction is highly prone to coherent errors in the system. For the parity measurement, an $\pi/2$-rotation pulse brings the state of the qubit, originally in $|g\rangle$, to the equator of the Bloch sphere, $|+\rangle$, by evolving the qubit state under the Hamiltonian $\hat{H} = (\Omega/2)\hat{\sigma}_y$ for a time $\pi/(2\Omega)$. Then, a conditional phase gate, $\hat{C}_\pi = |g\rangle\langle g| \otimes \mathbb{1} + |e\rangle\langle e| \otimes e^{i\pi\hat{n}}$, is implemented by waiting for a time $t_w = \pi/\chi$. As such, the state of the qubit acquires a relative phase depending on the photon-number parity, with its resulting state being in $|+\rangle(|-\rangle)$ when the cavity contains an even (odd) number of excitations. A final $\pi/2$-rotation pulse maps $|+\rangle(|-\rangle)$ to $|e\rangle(|g\rangle)$. Despite this protocol being a good approximation to estimate the parity, the always-on dispersive interaction during the $\pi/2$ pulses imparts substantial error to the measurements, with the effective Hamiltonian being $\hat{H} = (\Omega/2)\hat{\sigma}_y - \chi\hat{n}|e\rangle\langle e|$. When the cavity has $n > 0$ excitations, the qubit rotations happen along a slanted axis, given by $\hat{r} = (\Omega\hat{y} + \chi n\hat{z})/\sqrt{\Omega^2 + (\chi n)^2}$. As a result, the qubit state prematurely accumulates phase during the first $\pi/2$ pulse and it does not lie on the equator at the end of the rotation, yielding a distorted parity approximation that degrades dramatically with increasing n, as illustrated in Fig. 2(b).

A standard technique to mitigate parity-measurement errors from the skewed qubit rotation is to calibrate for a shorter waiting time $t_w < \pi/\chi$, which maximizes the contrast. Another method to mitigate parity-mapping errors is to perform two parity measurements in succession for each point in phase space. In this scheme, the second $\pi/2$ rotation in the sequence is enacted with opposite phases such that the mapping between the parity and the qubit state is flipped. Finally, the corrected parity is computed by taking the difference between these two different mappings to remove the spurious phase accumulation. However, this correction demands twice the number of measurements and still results in a scaling error of the resulting parity,

$$P_{\text{cor}} = (P - P_{\text{inv}})/2 = \eta P_{\text{id}}, \tag{3}$$

where P and P_{inv} are the standard and inverted parity, $P_{\text{id}} = \text{Tr}\left(\rho e^{i\pi\hat{n}}\right)$ is the ideal parity, and η is the scaling factor (Appendix B 2).

Experimentally, this degradation of parity mapping with increasing excitation number in the CV state can be readily observed. We measure the parity of a series of Fock states with the standard procedure described above, using 16-ns $\pi/2$ pulses and a 284-ns waiting time. The results, shown in Fig. 2(a), show a larger discrepancy between the experimentally measured $|P_{\text{corr}}|$ and its ideal value of 1 as the Fock-state number of the cavity $|n\rangle$ increases. This distorted mapping makes parity a nonideal observable for state reconstruction, highlighting the importance of exploring the excitation number as a more reliable observable.

Apart from its resilience against coherent errors stemming from the Hamiltonian itself, ORENS is also robust against qubit decoherence during the process of mapping the excitation number onto the qubit state. In general, qubit decoherence is described by the two independent mechanisms of energy decay and dephasing, which are characterized by the coherence times T_1 and T_ϕ, respectively. While standard cQED setups can reliably achieve a T_1 in the range of several tens to hundreds of microseconds

[35], ensuring a consistent T_ϕ proves to be a challenging task [36]. This challenge is particularly pronounced in the case of flux-tunable qubits, where T_ϕ can be as short as a few microseconds [37]. Considering a dephasing rate of the qubit $\Gamma_\phi = 1/T_\phi$, by solving the Lindblad master equation, the excitation number decays as

$$p_n' \equiv p_e'(\Delta = \chi n) \approx \rho_{nn} \times \frac{1}{2}(1 + e^{-\frac{t_\pi}{2T_\phi}}) \tag{4}$$

(Appendix C 1), where we have used a prime notation to denote the case under decoherence. Equation (4) shows that p_n' contains a constant term (half of the magnitude of the ideal observable, $\rho_{nn} = \langle n|\rho|n\rangle$) and another that decays exponentially with the dephasing rate. This means that a finite proportion of the contrast is always preserved in measurement.

We perform an experiment to measure p_0 of the vacuum cavity state at various engineered qubit dephasing times T_ϕ to observe its impact on the observable mapping (Fig. 3). The T_ϕ of the qubit is shortened on demand by using excitation-induced dephasing [38], where a weak coherent tone continuously drives the readout resonator (Appendix A 6). To effectively isolate the error due to dephasing, the dephasing tone is only applied during the mapping of the observable, after the state-preparation pulses, and with a calibrated buffer time to allow the resonator to depopulate before the readout tone. The duration t_π of the π pulse and the dispersive coupling χ are both fixed.

Our results show that excitation-number mapping is partially preserved under qubit dephasing. At $T_\phi = 4\ \mu\text{s} = 4t_\pi$, we observe a $p_0 > 0.8$, and for as low as $T_\phi = 0.1\ \mu\text{s} = t_\pi/10$, the expectation value remains above 0.4. To benchmark this result, we repeat the same protocol with parity mapping. From the Lindblad master equation, we find that the whole parity observable exponentially decays to 0 (i.e., total loss of information) as T_ϕ goes to 0. Analytically, this is described by

$$P' \equiv 2p_e' - 1 \approx P_{\text{id}} \times e^{-\frac{t_w}{T_\phi}}, \tag{5}$$

where t_w denotes the variable waiting time between the two $\pi/2$ pulses (Appendix C 1). Since parity mapping relies on the qubit state acquiring a deterministic phase, low dephasing times $T_\phi < 0.1\ \mu\text{s}$ completely erase this mapping, as shown in Fig. 3(a). At high T_ϕ, we expect both the excitation number and the parity of the vacuum state to saturate the ideal value in the absence of coherent errors. However, errors due to parasitic dynamics become dominant in the parity-mapping process for states with more excitations, as shown in Fig. 2. This leads to more pronounced reconstruction imperfections. Thus, it is apparent that excitation-number mapping offers a more robust performance than the standard parity measurements against both coherent errors and decoherence.

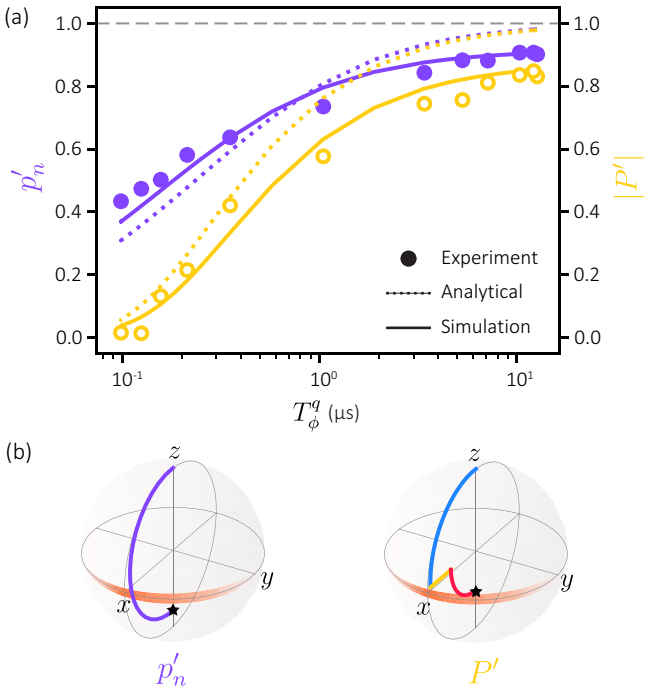


FIG. 3. The excitation number and parity of the vacuum cavity state under qubit dephasing. (a) The measurement outcomes of the excitation-number p_n' (purple) and the absolute value of the corrected parity $|P'|$ (yellow) mappings for the vacuum state are plotted versus the varying T_ϕ of the qubit. They show good agreement with analytical [Eq. (4) and (5)] and simulated trends based on real device parameters. With roughly half the contrast preserved, excitation-number measurements are robust against qubit dephasing in cQED, whereas parity mapping suffers significantly more degradation. (b) Intuitive illustrations for the resilience shown by excitation-number mapping. The qubit state spends less time near the equator (orange) in mapping p_n' than P' and is thus less susceptible to dephasing.

IV. ORENS STATE RECONSTRUCTION

To fully demonstrate the efficacy of ORENS, we use it to perform full CV state reconstruction for states of different sizes in our bosonic cQED system. For a given Hilbert space of dimension D, we first obtain the set of optimized displacements by sweeping over the excitation number $n \in [1, D-1]$ and implementing a gradient-descent algorithm to find $\{\alpha_k\}_{k=1}^{N_{\text{obs}}^*}$, where the cost function to minimize is the condition number of the measurement matrix (Appendixes D 1 and D 2). This set of measurements allows us to reconstruct any arbitrary state bounded by dimension D. To benchmark the performance of ORENS comprehensively, we prepare all Fock states $|k\rangle$ and their superpositions $|j\rangle + e^{i\phi}|k\rangle/\sqrt{2}$ with $j < k = 0, \ldots, D-1$ and $\phi = \{0, \pi/2\}$ (D^2 different states in total) within the same truncation dimension. Experimentally, each state is prepared by playing 2-μs numerical pulses after a qubit preselection measurement, which eliminates the residual thermal population (approximately 2%). Next, we displace

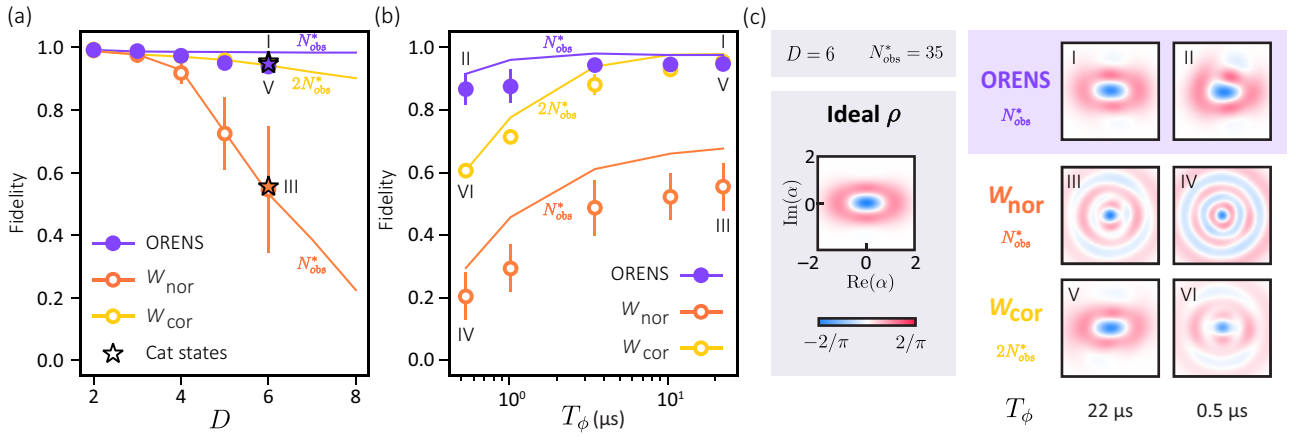


FIG. 4. State reconstruction for ORENS ($N^*_{\rm obs}$), normal Wigner ($N^*_{\rm obs}$), and corrected Wigner ($2N^*_{\rm obs}$). (a) The average reconstruction fidelities for D^2 different states (Fock states and their superpositions) versus the cutoff dimension D. They show good agreement with simulated trends (solid curves) based on real device parameters. The star markers are average fidelities for reconstructing four small cat states, well contained within $D = 6$. Beyond $D = 6$, there are no meaningful experimental points, as the readout was distorted by cross-Kerr between the readout resonator and cavity. The qubit dephasing time here is $T_\phi = 15.3$ μs. (b) The average reconstruction fidelities of the four small cat states for different qubit dephasing times T_ϕ. (c) Wigner functions of the reconstructed odd cat states ($|\alpha\rangle - |-\alpha\rangle$) with ORENS, normal Wigner, and corrected Wigner techniques at two different T_ϕ points. With strong retention of important cat-state properties, ORENS outperforms both Wigner-reconstruction techniques with the fewest measurements, particularly at high dimensions and under qubit dephasing.

the cavity state in phase space with the set of optimized displacements, $\hat{D}(\alpha_k)$, via a drive resonant at the cavity frequency with a Gaussian envelope. As a concrete example, the optimal displacements for $D = 6$ are 35 unique $\{\alpha_k\}_{k=1}^{35}$, each followed by a measurement of excitation number $n = 5$ in this example. This is enacted by exciting the qubit with a pulse resonant at the qubit frequency conditioned on the cavity having n excitations, $\omega_q - n\chi$, where $n = 5$. The resulting qubit state is obtained through standard single-shot dispersive readout. These measurement outcomes are processed to estimate the cavity state, first by inverting Born's rule to obtain the least-squares estimator $\rho_{\rm LS}$, and then by using Bayesian inference [27] to obtain $\rho_{\rm BME}$ as the final reconstructed estimator. The Bayesian method treats uncertainty in meaningful ways and utilizes all available information optimally. As a result, this approach affords the most faithful estimator for the state [39,40], particularly in comparison to the typical maximum-likelihood-estimation approach [41]. More details on the Bayesian interference technique used in this study are presented in Appendix D 3.

To evaluate the quality of the state reconstruction, we compute the fidelity between the estimated density matrix and the target density matrix $\rho_{\rm tar}$, generated by simulating the GRAPE pulses under decoherence,

$$F = \left(\mathrm{Tr}\sqrt{\sqrt{\rho_{\rm tar}}\rho_{\rm BME}\sqrt{\rho_{\rm tar}}}\right)^2. \tag{6}$$

The average fidelity over the D^2 different states for each dimension is plotted in Fig. 4(a), up to $D = 6$. Across all dimensions, the reconstruction fidelity using ORENS exceeds 95%. Beyond $D = 6$, the readout was distorted due to the dominant cross interaction between the cavity and the readout resonator present for larger bosonic states, resulting in the absence of meaningful experimental points. This is not a fundamental limitation of the technique but a rather device-specific artifact. The dominant error mechanisms impacting the reconstruction fidelity are the thermal population of the qubit and imperfect single-shot readout discrimination (Appendixes A 7 and C 2), which could be significantly reduced by improving device thermalization and using a quantum limited amplifier, respectively.

Having tested the reconstruction technique on Fock states and their superpositions, we evaluate the protocol for cat states, a versatile backbone of CV information processing protocols [42–46]. We repeat the same reconstruction protocol for $D = 6$ with four small cat states: $|\alpha\rangle \pm |-\alpha\rangle$ and $|\alpha\rangle \pm i|-\alpha\rangle$ (normalization implied), with $\alpha = 1$. The averaged fidelity matches the Fock-state superposition fidelities (see the purple star in Fig. 4).

We benchmark the reconstruction performance of both Fock and cat states against two versions of the standard Wigner protocol, normal and corrected. For both variations, we fix the measurement observable to parity and optimize the corresponding $N^*_{\rm obs}$ displacements. For corrected Wigner, the distortion due to the always-on dispersive interaction is partially removed at the cost of increasing the number of independent measurements to $2N^*_{\rm obs}$, which is twice the theoretical minimum.

With a focus on using only the minimal number of measurements $N^*_{\rm obs}$, ORENS demonstrates accurate

reconstruction across dimensions, significantly surpassing the performance of the normal Wigner protocol. While the corrected Wigner tomography affords state-reconstruction fidelity comparable to that of ORENS, it is more resource-intensive, requiring double the number of independent measurements. Compared to prior works using the generalized Q function measured with a grid of unoptimized displacements and all the Fock-state projections up to the maximum photon number [13], ORENS affords a 50-fold reduction in the number of measurements.

To further verify the performance of ORENS under qubit dephasing, we prepare the same four cat states and reconstruct them using ORENS for several independently calibrated qubit T_ϕ points, as shown in Fig. 4(b). Remarkably, we note only a slight deterioration of the average reconstruction fidelity, with fidelities exceeding $> 86\%$. This demonstrates that even when using a conditional π-pulse duration that far exceeds the qubit T_ϕ, the excitation-number observable still maps sufficient information to accurately reconstruct states.

The robustness under dephasing attests to the versatility of the ORENS in regimes of low χ between the cavity and the qubit. To maintain an equivalent frequency selectivity, a smaller χ demands a longer t_π (Appendix B 1). Considering our experimental (simulated) fidelity of 88% (92%) with $\chi/2\pi = 1.4$ MHz and $T_\pi/T_\phi = 1\,\mu\text{s}/0.5\,\mu\text{s} = 2$, we can expect an equivalent fidelity with $\chi/2\pi = 35$ kHz, $T_\pi/T_\phi = 40\,\mu\text{s}/20\,\mu\text{s} = 2$. This was verified with a simulation to reconstruct the cat states, with an average fidelity of 94%.

V. CONCLUSION AND DISCUSSION

Through the above analytical and experimental results, we have demonstrated: (1) a powerful technique for optimized reconstruction with excitation-number sampling (ORENS), with minimal measurements, that relies only on displacements and excitation counting and can be readily applied across CV experimental platforms; (2) clear evidence that excitation-number mapping in bosonic cQED is an ideal and convenient observable for state reconstruction that can be directly implemented on standard devices without any tailored operations or parameters; (3) the robustness of excitation-number mapping even under severe qubit dephasing in cQED; and (4) the ability of ORENS to reliably reconstruct arbitrary states of all dimensions in the presence of pronounced coherent and incoherent errors.

For each of the experiments, ORENS outperforms the state-of-the-art Wigner reconstruction with the fewest measurements. Although the fidelities obtained with ORENS are nearly matched by the corrected Wigner strategy, our method uses half the number of measurements and scales more favorably with state dimensionality. The primary drawback of ORENS in cQED is the high sensitivity to undesired residual excitations of the qubit (Appendix C 2). However, the reconstruction fidelity can be readily improved with good thermalization of the qubit as well as standard preselection measurements.

Looking beyond, the ORENS technique can be readily implemented for multimode systems. For example, for a two-mode system with modes A and B, each with dimension D, we would apply the displacements $\{D_A(\alpha_i) \otimes \mathbb{1}_B, \mathbb{1}_A \otimes D_B(\alpha_j), D_A(\alpha_i) \otimes D_B(\alpha_j)\}$, where $\{\alpha_i\}$ is the set of optimized single-mode displacement points, before measuring the joint excitation numbers. In cQED, the joint excitation numbers can be directly extracted by either using a single ancillary qubit coupled to both cavities via a selective π pulse tuned to $\omega_q - n_A\chi_A - n_B\chi_B$ or by using two single-cavity ancillary qubits and correlating their outcomes. Not only would this multimode reconstruction approach likely outperform the standard joint Wigner tomography for similar arguments of excitation-number observable robustness, but the measurement of the joint excitation number is significantly more convenient than that of joint parity. Joint-parity measurements are challenging, as they require either designing $\chi_A = \chi_B$ or utilizing higher levels of the transmon with concatenated single-mode conditional phase gates [18]. While the generalized Wigner function [16] helps overcome these practical challenges, the arbitrary relative phase of the modified Ramsey sequence tends to reduce the contrast of the measurement outcomes and renders the reconstruction less robust.

Overall, we have developed and demonstrated a versatile technique to efficiently and robustly estimate arbitrary CV states, accessible across different bosonic hardware platforms. Our results bring us one step closer to scalable and reliable characterization and verification across CV quantum applications.

ACKNOWLEDGMENTS

We acknowledge the funding support of the National Research Foundation (NRF) Grant No. NRF2020-NRF-ISF004-3540 and the Ministry of Education, Singapore Grant No. MOE-T2EP50121-0020. T.K. thanks Tomasz Paterek for discussions during the earlier stages of this project.

DATA AND CODE AVAILABILITY

All data and codes needed to evaluate the conclusions of the paper are available at ref. [47].

APPENDIX A: EXPERIMENTAL DEVICE

1. Design and tools

The experimental device used in this work is a standard bosonic cQED system [48,49] in the strong-dispersive-coupling regime. It consists of a superconducting microwave cavity, dispersively coupled to a transmon

qubit for controllability and readout, which is also dispersively coupled to a planar readout resonator. We simulate the electromagnetic fields of the device using the ANSYS finite-element high-frequency simulation software (HFSS) and obtain the Hamiltonian parameters using the energy-participation-ratio (EPR) approach [50]. The key system properties, such as the frequency of each circuit and the pairwise nonlinear couplings between them, are iteratively refined to meet the target parameters. In this section, we describe the details of the design considerations, the resulting properties of the main elements in the device, and the main calibration procedures for the different experimental parameters.

2. Package and chip fabrication

The storage cavity is a 3D high-Q coaxial $\lambda/4$-resonator with a cutoff frequency $f_{\text{cut}} \sim 600$ MHz. The cavity and the coaxial waveguide that hosts the qubit and the resonator are machined out of high-purity (4N) aluminum, where the external layer (0.15 mm) has been removed with chemical etching to reduce fabrication imperfections. The ancillary transmon qubit and the planar readout resonator are fabricated by evaporating aluminum on a sapphire substrate. The design is patterned using a Raith electron-beam-lithography machine, on a HEMEX sapphire substrate (i.e. sapphire substrate grown by the heat exchanger method graded based on superior optical properties) cleaned with 2:1 piranha solution for 20 min and coated with 800 nm of methyl methacrylate (MMA) and 250 nm of poly(methyl methacrylate) (PMMA) resist. The pattern is then developed with a mixture of deionized water and isopropanol at a 3:1 ratio. Using a PLASSYS double-angle evaporator, we deposit the two aluminum layers of 20 nm and 30 nm thickness at -25 and $+25$ degrees, respectively, separated by an oxidation step with a mixture of 85% O_2 and 15% Ar at 10 mBar for 10 min. The chip is finally diced on an Accretech machine and inserted in the waveguide, with an aluminum clamp where we use indium wire to improve thermalization.

3. Intrinsic Purcell filtering

The design of the Hamiltonian parameters considers the ORENS protocol requirements and the versatility to explore different decoherence regimes. We design the dispersive interaction between the cavity and the qubit to be approximately 1.4 MHz. Hence, our selective π pulses need to be approximately 1 μs long and the Ramsey revival time is approximately $\pi/\chi \approx 3.14$ μs.

To achieve long-enough coherence times for the qubit, we mitigate its resonator-mediated Purcell decay by designing an intrinsic Purcell-filter structure [51], see Fig. 5. This is done by optimizing the position of the coupled transmission line. Simulations show that the optimal position aligns with the voltage node of the qubit field at approximately $\lambda_{\text{qubit}}/4$ away from the end of the resonator. In this position, the qubit field is very weakly coupled to the transmission line, while the readout resonator is significantly coupled for fast readout.

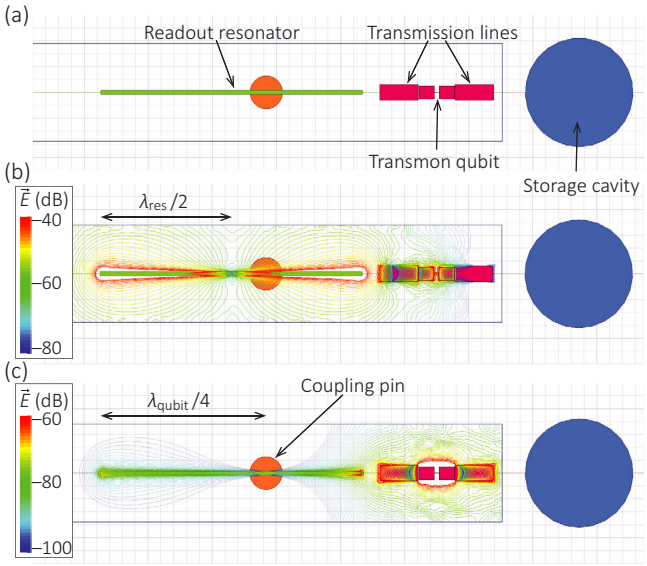


FIG. 5. The chip design. (a) The simulated readout resonator, qubit, and cavity modes in ANSYS HFSS. (b),(c) The complex magnitude of the electric field of the (b) resonator and (c) qubit modes. The position of the coupling pin (orange circle) is at a distance $\lambda_{\text{qubit}}/4$ from the end of the resonator, where the qubit field is a minimum, while the resonator field is nonzero.

For ease of fabrication and to preserve cavity coherence, the transmission line is kept at an appropriate distance. To satisfy this requirement, as well as the strong dispersive coupling and the intrinsic Purcell-filtering condition at the same time, we add two planar strip-line-like structures on both ends of the qubit pads. The strips are short enough not to introduce any mode below 8 GHz, while effectively guiding the transmon field to the cavity and to readout-resonator modes as desired.

4. Hamiltonian parameters and coherence times

Expanding the cosine term of the Josephson junction up to the fourth order, we can write the full Hamiltonian of our system as

$$\begin{aligned}\frac{\hat{H}}{\hbar} &= \omega_{\text{c}}\hat{c}^\dagger\hat{c} + \omega_{\text{q}}\hat{q}^\dagger\hat{q} + \omega_{\text{r}}\hat{r}^\dagger\hat{r} \\ &\quad - \frac{\chi_{\text{cc}}}{2}\hat{c}^\dagger\hat{c}^\dagger\hat{c}\hat{c} - \frac{\chi_{\text{qq}}}{2}\hat{q}^\dagger\hat{q}^\dagger\hat{q}\hat{q} - \frac{\chi_{\text{rr}}}{2}\hat{r}^\dagger\hat{r}^\dagger\hat{r}\hat{r} \\ &\quad - \chi_{\text{cq}}\hat{c}^\dagger\hat{c}\hat{q}^\dagger\hat{q} - \chi_{\text{qr}}\hat{q}^\dagger\hat{q}\hat{r}^\dagger\hat{r} - \chi_{\text{cr}}\hat{c}^\dagger\hat{c}\hat{r}^\dagger\hat{r}, \end{aligned} \tag{A1}$$

where ω_{i} and $\hat{\imath}$, respectively, denote the angular frequency and annihilation operator of the system, with $i = c, q$, and r corresponding to cavity, qubit, and resonator. The χ_{ij} of the second and third lines correspond to the self-Kerr and

TABLE I. The frequency and coherence times for the three modes of the device.

	$\omega/2\pi$(GHz)	T_1(μs)	T_2(μs)	T_2^{echo}(μs)
Qubit	5.277	85–113	14–22	44–48
Cavity	4.587	992	...	...
Resonator	7.617	2.08	...	...

cross-Kerr interactions between modes, respectively. The value of the experimentally calibrated parameters can be seen in Tables I and II. The high self-Kerr of the transmon allows effective treatment of it as a qubit with two energy levels, $|g\rangle$ and $|e\rangle$, where $\hat{q}$ ($\hat{q}^\dagger$) can be replaced with $|g\rangle\langle e|$ ($|e\rangle\langle g|$).

The next major contribution to Eq. (A1), stemming from the sixth-order term of the cosine expansion, corresponds to the second-order dispersive interaction between the cavity and the qubit, $-\chi'_{cq} q^\dagger q c^\dagger c^\dagger c c$. By fitting the resonance frequencies of the qubit to second order on the number of excitations in the cavity, we find $\chi'_{cq}/2\pi \approx 16$ kHz.

5. Microwave wiring

The radio-frequency (rf) pulses to drive the readout resonator, qubit, and cavity are created by in phase and quadrature (IQ) mixing of the local oscillator (LO) signal from a Vaunix Lab Brick microwave resonator with the intermediate-frequency (IF) $\mathcal{I}$ and $\mathcal{Q}$ quadratures generated by the digital-to-analog converter (DAC) port of a Quantum Machines fast field-programmable gate array (FPGA). For the readout, we measure the reflected signal from the resonator, which is amplified in a high-electron-mobility transistor (HEMT) amplifier and a room-temperature ZVA-183S+ amplifier before being down-converted to 50 MHz with a Marki IR mixer. The signal is finally amplified with a Stanford Research Systems SR445A room-temperature amplifier before being sampled in the analog-to-digital converter (ADC) block of the FPGA. The schematic wiring setup can be in seen Fig. 6.

6. Engineering qubit dephasing

To demonstrate the robustness of ORENS under qubit dephasing, we engineer the dephasing time T_ϕ of the qubit. This is achieved by driving the dispersively coupled readout resonator to a steady-state photon population that induces dephasing via photon-shot noise. The dephasing rate is controlled by varying the average excitation number in the resonator [38].

TABLE II. The Kerr interactions. The diagonal elements correspond to the self-Kerr interactions of each mode and the off-diagonal terms correspond to the cross-Kerr interactions between different modes.

	Cavity	Qubit	Resonator
Cavity	4–6 kHz	1.423 MHz	2 kHz
Qubit	1.423 MHz	175.3 MHz	0.64 MHz
Resonator	2 kHz	0.64 MHz	...

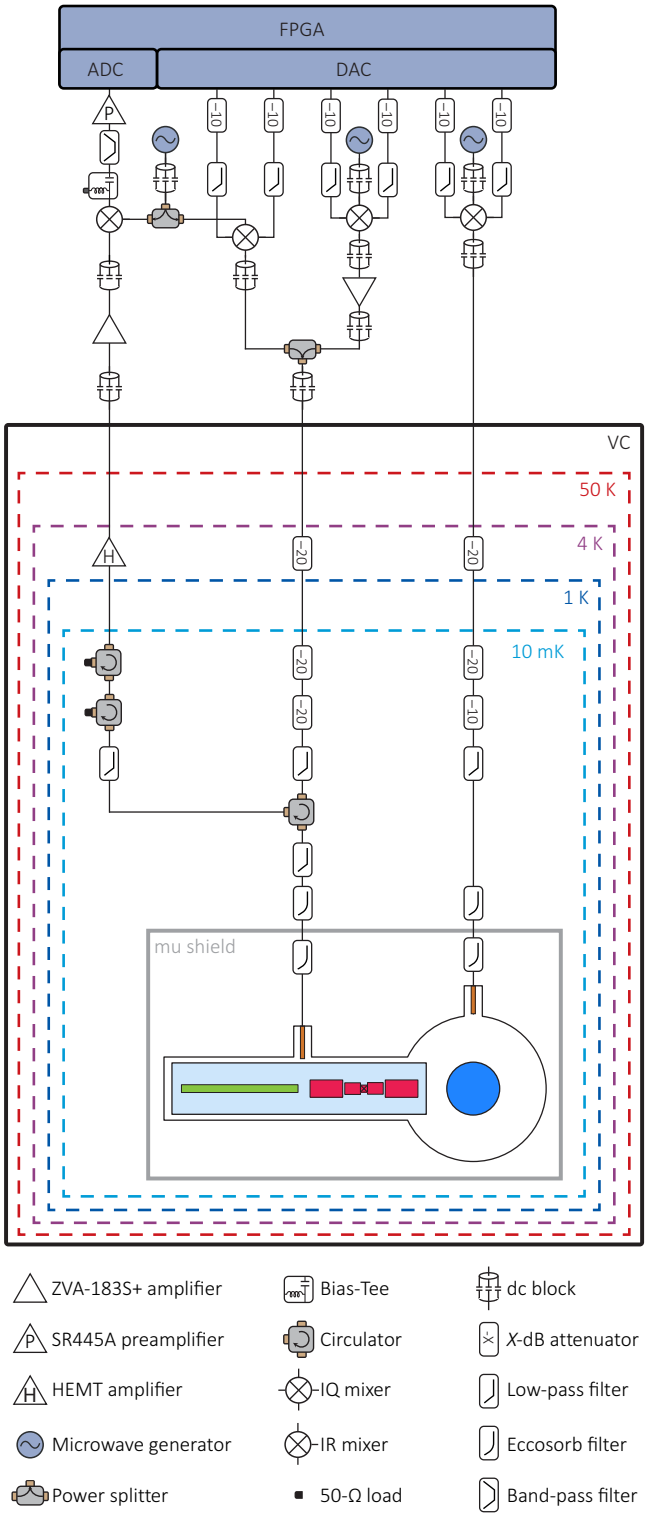


FIG. 6. The experimental setup. A schematic of the rf components and connections at room temperature and inside the Bluefors dilution refrigerator.

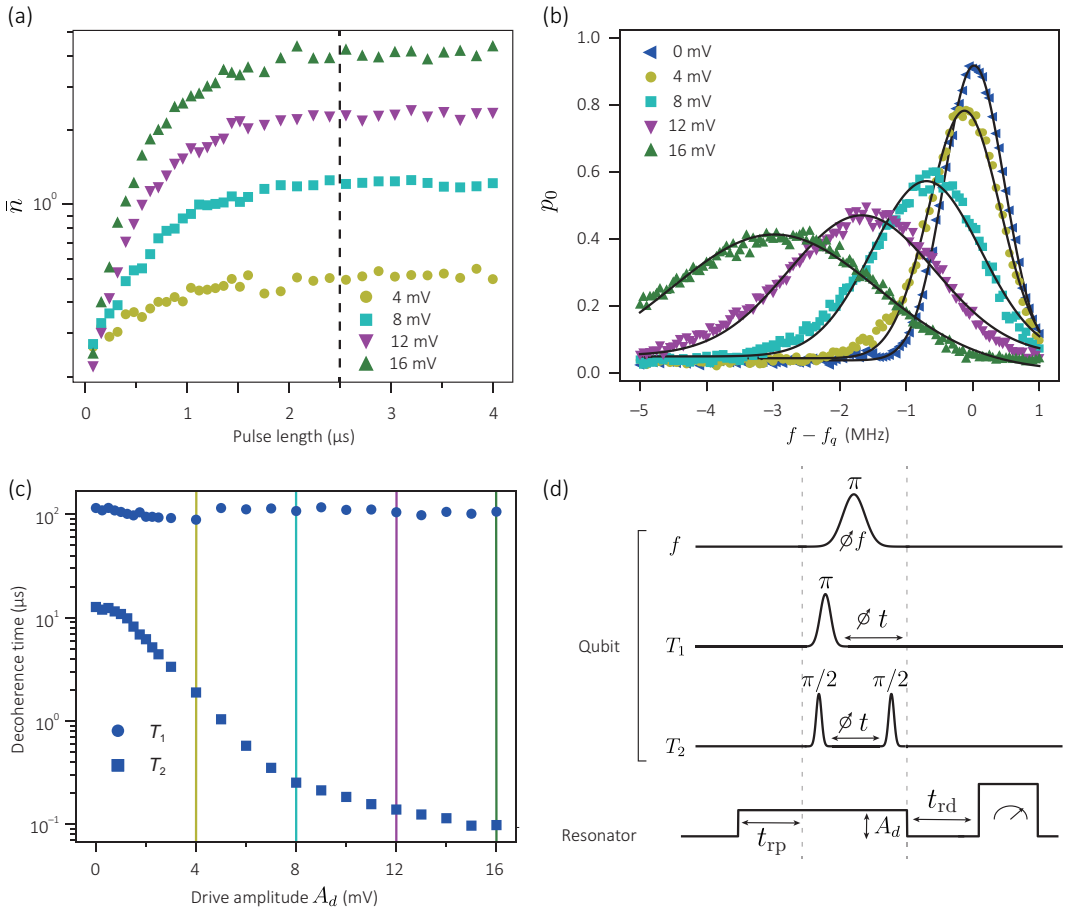


FIG. 7. Engineering qubit dephasing. (a) The average excitation number $\bar{n}$ in the resonator as a function of the various pulse lengths for several drive amplitudes A_d. The resonator population reaches a steady state after a ring-up time t_{rp} for all A_d, which is taken as 2.5 μs (black dotted line). (b) Qubit spectra after driving the resonator to steady state with square pulses of varying drive amplitudes. The black lines represent Gaussian fits. The measurement pulse sequence is shown in (d), where the ring-up time t_{rp} and the ring-down time t_{rd} are both 2.5 μs and the length of the π pulse is 1 μs. (c) The qubit energy relaxation time T_1 and dephasing time T_2 obtained with different resonator drive amplitudes. The vertical lines correspond to the drive amplitudes used in (a) and (b). (d) The qubit T_1 and T_2 times in the presence of photons in the readout resonator are obtained with the pulse sequences shown.

To calibrate the average excitation number $\bar{n}$ in the resonator as a function of the drive amplitude, we populate the resonator with a square pulse and conduct a Ramsey experiment. Subsequently, we fit the modulated Ramsey oscillations using the free parameter $\bar{n}$ [52]. For a given drive amplitude A_d—expressed as the voltage of the DAC output—we can extract $\bar{n}$ as a function of the pulse length [see Fig. 7(a)]. For all drive amplitudes, the resonator reaches a steady state after 2.5 μs, which is set as the ring-up time $\tau_{rp} = 2.5$ μs. Following a similar calibration, we choose the ring-down time τ_{rd} for resonator decay as 2.5 μs.

Driving the resonator not only induces qubit dephasing but also shifts the qubit frequency due to its dispersive interaction. As shown in the qubit-spectroscopy plot in Fig. 7(b), measured with the sequence in Fig. 7(d), the qubit peak broadens and shifts with increasing drive amplitude A_d. To isolate the dephasing effect from the frequency-shifting effect in the experiments that follow, we perform the observable-mapping protocols using the shifted frequency corresponding to the resonator drive amplitude.

To characterize the pure dephasing time $T_\phi = 1/(1/T_2 - 0.5/T_1)$ of the qubit for each resonator drive, we measure the qubit energy relaxation time T_1 and qubit dephasing time T_2 after driving the resonator to steady state and updating the qubit frequency [pulse sequences in Fig. 7(d)]. As shown in Fig. 7(c), T_1 stays relatively constant with increasing resonator drive amplitude, while T_2 decreases smoothly. This indicates that the photon-shot noise in the resonator only induces qubit dephasing, not qubit energy relaxation.

Having calibrated how to engineer T_ϕ, we now study how the measurement observables—the excitation number

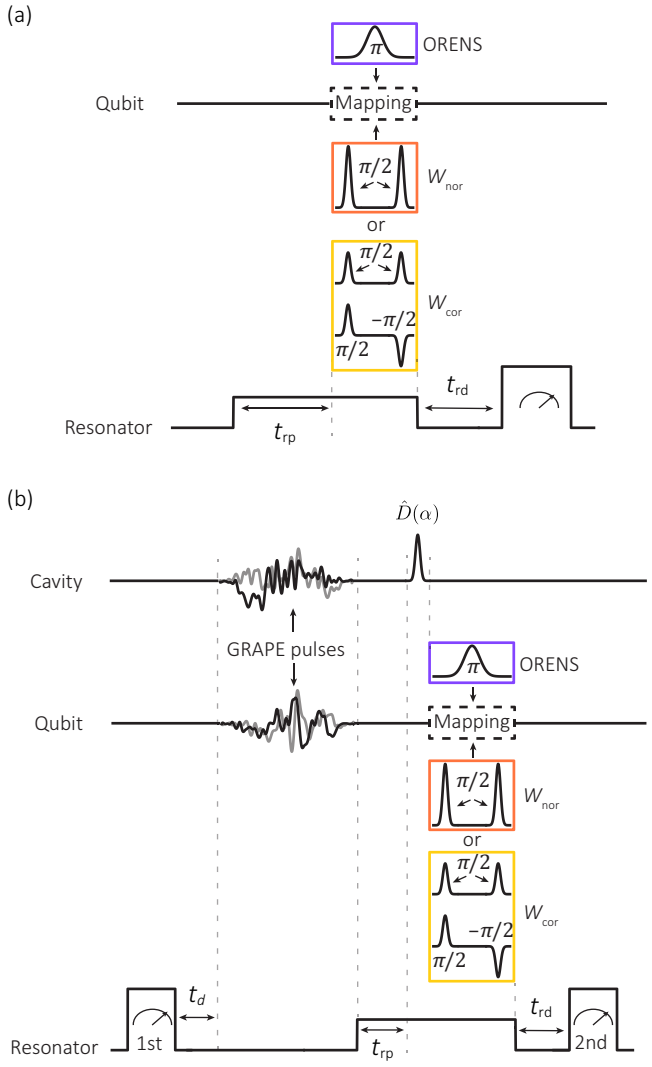


FIG. 8. Observable-mapping sequences with engineered qubit dephasing. (a) The pulse sequence for ORENS and Wigner observable mapping (Fig. 3). The ring-up time t_{rp} and the ring-down time t_{rd} are both 2.5 μs, the lengths of the π pulse and the $\pi/2$ pulse are 1 μs and 64 ns, respectively, and the delay time τ for parity measurement is 284 ns. (b) The pulse sequence for ORENS and Wigner-state reconstruction [Figs. 4(b) and 4(c)]. We apply an additional measurement pulse to the resonator at the start for postprocessing selection. The delay time t_d for the resonator decay after the first measurement, the ring-up time t_{rp}, and the ring-down time t_{rd} are both 2.5 μs. The length of the GRAPE pulses is 2 μs, the length of the displacement pulse $D(\alpha)$ is 240 ns, the lengths of the π pulse and the $\pi/2$ pulse are 1 μs and 64 ns, respectively, and the delay time τ for parity measurement is 284 ns.

p_n for ORENS and the parity P for Wigner—behave when increasing the qubit dephasing rate while keeping the cavity in the vacuum state (see Fig. 3). The pulse sequence is shown in Fig. 8(a). For a certain calibrated T_ϕ, we drive the resonator to steady state, update the qubit frequency, and conduct the standard p_n and P observable-mapping

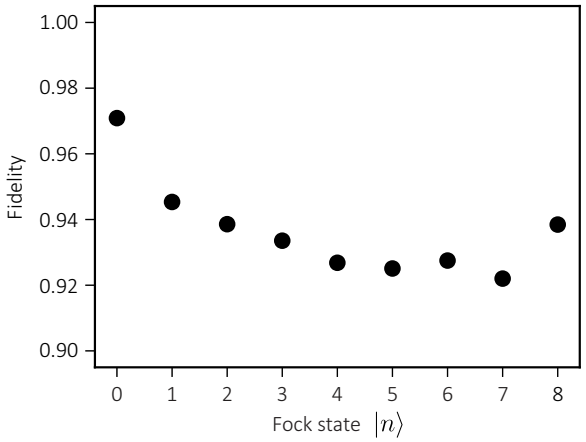


FIG. 9. The quality of the state preparation. The fidelities of the Fock states created using numerically optimized pulses.

protocols. Then, after a ring-down time t_{rd} for the resonator to depopulate, the measurement pulse is applied.

We evaluate ORENS and Wigner-state reconstruction under dephasing (see Fig. 4) using the pulse sequence in Fig. 8(b). To mitigate the error due to the qubit thermal population, we first measure the qubit state to later postselect the data. After waiting a delay time t_d for the resonator to decay, we prepare the small cat states using GRAPE pulses with a length of 2 μs. After state preparation, we drive the resonator to steady state to induce a reduced T_ϕ, apply a displacement pulse to the cavity, and perform either the p_n or P mapping. Then, we turn off the resonator drive and wait a ring-down time t_{rd} for the resonator decay, before applying a final measurement pulse.

7. Error budgeting

We use a standard square low-power readout pulse at the resonance frequency of the readout resonator when the qubit is in the ground state. The length of the readout pulse is 1.5 μs and the reflected signal is acquired for a total time of 2.4 μs. We measure a readout fidelity $F_{RO} = 1 - (p(e|g) + p(g|e))/2 = 96.5\%$, of which we estimate an infidelity of 1.9% due to thermal population, 1.4% due to readout discrimination error (overlap), and 0.4% due to qubit decay during the readout.

The cavity has been measured to have a 3% residual population, i.e., $p_1 = 3\%$. The cavity states have been prepared with numerical pulses optimized with the GRAPE algorithm applied simultaneously at the qubit and the cavity, with a fixed total duration of 2 μs, for all states. The average residual thermal population of the qubit after the GRAPE pulses has been measured to be 4.9% and is further attenuated down to 3% by performing a 1.5-μs-long readout before the numerical pulses and preselecting only those runs where the qubit was measured in the ground state. We have calibrated a 2.5-μs-long buffer time

between the readout pulse and the GRAPE pulses for the resonator to depopulate. We test the quality of the state preparation by simulating the effect of the numerical pulses with the whole qubit-cavity Hamiltonian accounting for decoherence and thermal populations. The fidelity in relation to the ideal pulses can be seen in Fig. 9. The 3% infidelity for preparing the cavity in vacuum (Fock 0 in the figure) can be explained by the previously mentioned thermal population of the cavity, and it matches the discrepancy between the experimental data and the ideal value in Fig. 2. In addition, the experimental p_n data show a slow systematic decay as a function of the Fock state, which can be explained by state-preparation errors of the GRAPE pulses (see Fig. 9).

The additional loss of contrast in Fig. 3 compared to Fig. 2 is caused by qubit decoherence during the additional ring-up and ring-down times, as explained in Sec. A 6.

APPENDIX B: COHERENT ERROR IN OBSERVABLES

1. Excitation-number mapping

Here, we derive the expression for the qubit probability following the excitation-number-mapping protocol, show how it maps to excitation number of the cavity state, and discuss its selectivity. We set $\hbar = 1$ henceforth.

First, we note that given the initial state of the qubit $|g\rangle$ and its evolution governed by a Hamiltonian $\hat{H} = \Delta|e\rangle\langle e| + (\Omega/2)\hat{\sigma}_y$, the probability of finding the qubit in the excited state at a time t is

$$p_e(t) = \frac{\Omega^2}{\Omega^2+\Delta^2}\sin^2(\sqrt{\Omega^2+\Delta^2}\,\frac{t}{2}). \tag{B1}$$

At a time $t_\pi \equiv \pi/\Omega$, the probability is at a maximum ($p_e = 1$) for $\Delta = 0$ and it decays as the detuning $|\Delta|$ increases.

To map the excitation number of the cavity onto the qubit, the sequence starts with the qubit in $|g\rangle$ and the cavity in an arbitrary state ρ, which evolve under the Hamiltonian

$$\hat{H} = \Delta|e\rangle\langle e| + \frac{\Omega}{2}\hat{\sigma}_y - \chi|e\rangle\langle e|\otimes\hat{n}, \tag{B2}$$

where $\hat{n} = \hat{c}^\dagger\hat{c}$.

The qubit state at time t follows

$$\begin{aligned}\rho_q(t) &= \mathrm{Tr}_{\mathrm{cav}}(e^{-i\hat{H}t}(|g\rangle\langle g|\otimes\rho)\,e^{i\hat{H}t})\\ &= \sum_n \langle n|e^{-i\hat{H}t}(|g\rangle\langle g|\otimes\rho)\,e^{i\hat{H}t}|n\rangle\\ &= \sum_n \rho_{nn}\,e^{-i\hat{H}_n t}|g\rangle\langle g|\,e^{i\hat{H}_n t},\end{aligned} \tag{B3}$$

where $\mathrm{Tr}_{\mathrm{cav}}$ is the partial trace with respect to the cavity state, $\rho_{nn} \equiv \langle n|\rho|n\rangle$ denotes the diagonal elements (excitation number) of ρ, and $\hat{H}_n \equiv \Delta_n|e\rangle\langle e| + (\Omega/2)\hat{\sigma}_y$ now acts only on the qubit with $\Delta_n \equiv \Delta - \chi n$.

The probability of the qubit being in the excited state at a time $t_\pi = \pi/\Omega$ is

$$\begin{aligned}p_e &= \sum_n \rho_{nn}|\langle e|\,e^{-i\hat{H}_n t_\pi}|g\rangle|^2,\\ &= \sum_n \rho_{nn}\frac{\Omega^2}{\Omega^2+\Delta_n^2}\sin^2\left(\sqrt{\Omega^2+\Delta_n^2}\,\frac{t_\pi}{2}\right),\end{aligned} \tag{B4}$$

where we have made use of Eq. (B1). Each component in the summation of Eq. (B4) features a function that peaks at $\Delta_n = \Delta - \chi n = 0$ to a value ρ_{nn}. For each component, the maximum with respect to the detuning Δ is different and, if the peaks are sharp enough ($\chi/\Omega \gg 1$), different peaks do not overlap with each other and we say that the excitation-number sampling is selective. The qubit probability can then be approximated as

$$p_e(\Delta = \chi n, t_\pi) \approx \rho_{nn}, \tag{B5}$$

which is the basis for the mapping of the excitation number.

Equation (B4) can also be written as

$$p_e = \sum_n \rho_{nn}\frac{1}{1+\eta^2}\sin^2\left(\sqrt{1+\eta^2}\,\frac{\pi}{2}\right), \tag{B6}$$

by defining the parameter $\eta \equiv \Delta_n/\Omega = (\Delta - \chi n)t_\pi/\pi$. If we rescale the detuning $\nu \equiv \Delta t_\pi$ (simply a scaling in the function with respect to Δ by t_π; the shape and, hence, the selectivity is the same), $\eta = (\nu - n\chi t_\pi)/\pi$. Consequently, it is clear that Eq. (B4), up to a scaling, is determined simply by χt_π. For instance, this means that having larger χ (more selective) is equivalent to having the duration of the protocol t_π longer. Also, a low χ (less selective) can be compensated by a longer t_π.

2. Parity mapping

In this section, we provide the analogous expressions for the qubit probability for the case of mapping the parity of the cavity state. We show that this mapping is inaccurate, which introduces a scaling and offset corrections to the ideal parity.

Parity mapping is done via a standard Ramsey-spectroscopy sequence ($\pi/2$ pulse - wait - $\pi/2$ pulse). The ideal evolutions during the $\pi/2$ pulses and the waiting time are governed by the Hamiltonians $\hat{H}_p = (\Omega/2)\hat{\sigma}_y$ applied for a time $t_{\pi/2} = \pi/(2\Omega)$ and $\hat{H}_d = -\chi|e\rangle\langle e|\otimes\hat{n}$ applied for a time $t_w = \pi/\chi$, respectively. Considering the qubit starting in $|g\rangle$ and the cavity in an arbitrary state ρ, both $\pi/2$ pulses realize a $\pi/2$ rotation on the qubit (along the y axis on the Bloch sphere) with a conditional phase gate $\hat{C}_\pi = |g\rangle\langle g|\otimes\mathbb{1} + |e\rangle\langle e|\otimes\hat{P}$, with $\hat{P} = e^{i\pi\hat{n}}$, implemented in between. This results in an ideal mapping of the

parity onto the qubit, the probability of finding the qubit in the excited state being

$$p_e = \frac{1}{2}(1 + P_{\mathrm{id}}), \tag{B7}$$

where $P_{\mathrm{id}} \equiv \mathrm{Tr}(\rho \hat{P})$ denotes the parity of the cavity state ρ. This can also be expressed as $P_{\mathrm{id}} = -\langle \hat{\sigma}_z \rangle$, since even- (odd-) parity states are mapped to the south (north) pole of the Bloch sphere of the qubit.

In a real scenario, however, the always-on dispersive coupling is also present during the $\pi/2$ pulses, which results in an actual Hamiltonian $\hat{H}_{pd} = (\Omega/2)\hat{\sigma}_y - \chi|e\rangle\langle e| \otimes \hat{n}$, which introduces a coherent error for the parity mapping. A standard technique to partially counter this is to reduce the duration during the wait $t_w < \pi/\chi$. For simplicity, considering the initial cavity state to be a pure Fock state $\rho = |n\rangle\langle n|$, the actual dynamics of the Ramsey spectroscopy yield a parity mapping

$$\begin{aligned} P' &\equiv -\langle \hat{\sigma}_z \rangle \\ &= f_1(\xi)\cos(\chi n t_w) - f_2(\xi)\sin(\chi n t_w) - f_3(\xi), \end{aligned} \tag{B8}$$

where $\xi \equiv \chi n/\Omega$ and

$$\begin{aligned} f_1 &= \frac{1}{(1+\xi^2)^2}\left(\sin^2\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right) + 2\xi^2\cos\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right)\left(1-\cos\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right)\right)\right), \\ f_2 &= \frac{2\xi}{(1+\xi^2)^{3/2}}\sin\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right)\left(1-\cos\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right)\right), \\ f_3 &= \frac{\left(\xi^2+\cos\left(\frac{\pi}{2}\sqrt{1+\xi^2}\right)\right)^2}{(1+\xi^2)^2}, \end{aligned} \tag{B9}$$

for a general waiting time t_w. In the limit of infinitely short $\pi/2$ pulses ($\xi \to 0$) and $t_w = \pi/\chi$, this equation reduces to the ideal parity $P' = \cos(n\pi) = (-1)^n = P_{\mathrm{id}}$. The f_1 term contains most of the ideal parity contrast, the f_2 term results in a correction from the optimal waiting time $t_w < \pi/\chi$, and the f_3 term yields a positive offset to the parity (regardless of its sign). To illustrate the values of f_1, f_2, and f_3 for higher Fock states, we plot these coefficients versus ξ in Fig. 10. By identifying the roles of these coefficients, we can conveniently define the parity errors

$$P' = P_{\mathrm{id}}\,\eta - \zeta, \tag{B10}$$

where η is a scaling error that incorporates the f_1 and f_2 terms and ζ is an offset error coming from f_3.

By reversing the second $\pi/2$ pulse in the sequence, i.e., using $\hat{H}_{pd} = -(\Omega/2)\hat{\sigma}_y - \chi|e\rangle\langle e| \otimes \hat{n}$, the parity now reads

$$\begin{aligned} P'_{\mathrm{rev}} &\equiv -\langle \hat{\sigma}_z \rangle \\ &= -f_1(\xi)\cos(\chi n t_w) + f_2(\xi)\sin(\chi n t_w) - f_3(\xi). \end{aligned} \tag{B11}$$

By using both P' and P'_{rev}, we can correct the offset error by computing

$$\begin{aligned} P'_{\mathrm{cor}} &\equiv \frac{P' - P'_{\mathrm{rev}}}{2} \\ &= f_1(\xi)\cos(\chi n t_w) - f_2(\xi)\sin(\chi n t_w) \\ &= P_{\mathrm{id}}\,\eta. \end{aligned} \tag{B12}$$

However, we can see that the scaling error remains.

We experimentally observe P' and the impact of the scaling and offset errors by measuring the parity of a series of Fock states with the Ramsey protocol, using 16-ns $\pi/2$ pulses and a 284-ns waiting time [see Fig. 11(a)]. Scaling error shows up as a decrease in the overall contrast

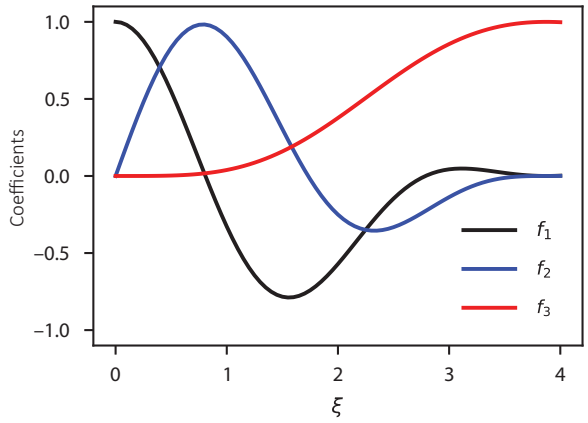


FIG. 10. The parity-mapping coefficients. The coefficients for the parity computed using Eq. (B9) as a function of $\xi = \chi n/\Omega$.

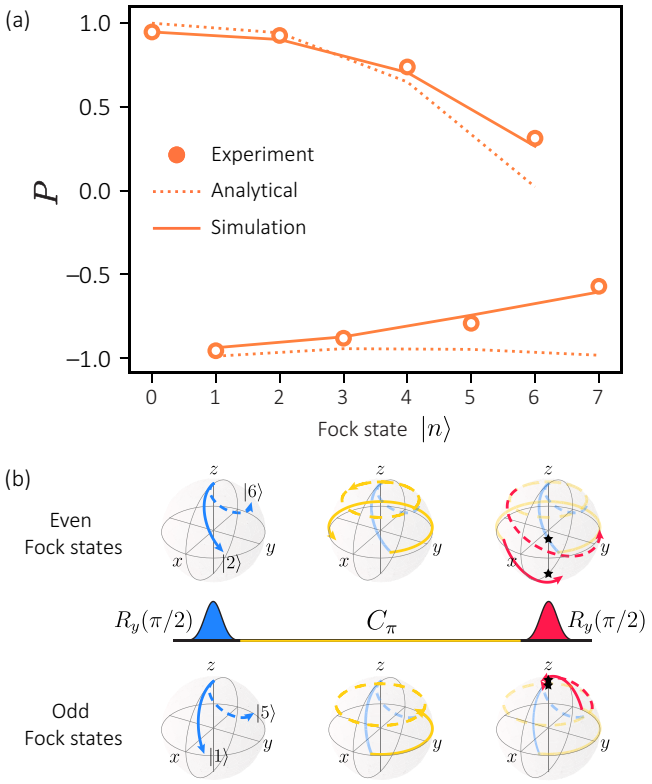


FIG. 11. Parity coherent errors.(a) The measurement outcomes of parity P under scaling and offset errors as a function of the Fock state $|n\rangle$ prepared in the cavity. They show good agreement with the analytical [Eq. (B10)] and simulated trends based on real device parameters. (b) The Bloch spheres of the qubit state at each step of the parity mapping for even and odd Fock states. The mapping of the even states suffers more due to the offset error. The black stars represent the measured outcomes (projections on the z axis).

and offset error appears as the skew between the even and odd Fock states. Figure 11(b) is an illustration of how the measurements of even Fock states are degraded more significantly than odd ones as n increases. This is due to the tilted rotation axis asymmetrically impacting the Ramsey evolutions of the evens and the odds. If we were to flip the mapping of the evens and odds to the qubit measurement outcomes, then odd Fock states would degrade more significantly.

APPENDIX C: ANALYSIS OF INCOHERENT ERRORS IN OBSERVABLES

1. Qubit dephasing T_ϕ

The primary limitation of excitation-number sampling is qubit decoherence. This stems from the frequency selectivity imposed by the finite dispersive coupling χ. To be precise, the coherence time imposes a limit on the maximum π-pulse duration before the qubit decoheres. The maximum pulse duration sets the frequency bandwidth, and thus the selectivity, of the pulse. Because the qubit frequency shifts by χn for each excitation n of the cavity, the maximum selectivity of the pulse then sets the minimum dispersive frequency shift χ necessary to resolve the excitation number. For the qubit, there are two loss channels to consider: energy decay and dephasing, which are characterized by their respective coherence times, T_1 and T_ϕ. While standard cQED setups can reliably achieve T_1s in the range of several tens to hundreds of microseconds [35], ensuring a consistent T_ϕ proves to be a challenging task [36].

In what follows, we will analyze the excitation-number and parity mapping only under qubit dephasing. For the case of the excitation number, Eq. (B3) in Appendix B 1 can be adapted to account for qubit dephasing by splitting the dynamics into small time intervals $\Delta t = t_\pi/N$ with $N \to \infty$, in which we repeatedly apply the loss channel $\varepsilon(\cdot) \equiv \sum_j E_j \cdot E_j^\dagger$, where the E_j are the Kraus operators and the unitary evolution $\hat{U}_{\Delta t} \equiv e^{-i\hat{H}\Delta t}$,

$$\begin{aligned}
\rho_q(t_\pi) &= \sum_n \langle n| \cdots \hat{U}_{\Delta t}\varepsilon(\hat{U}_{\Delta t}\varepsilon(|g\rangle\langle g| \otimes \rho)\hat{U}_{\Delta t}^\dagger)\hat{U}_{\Delta t}^\dagger \cdots |n\rangle \\
&= \sum_n \langle n|\rho|n\rangle \left(\cdots \hat{U}_{n,\Delta t}\varepsilon(\hat{U}_{n,\Delta t}\varepsilon(|g\rangle\langle g|)\hat{U}_{n,\Delta t}^\dagger)\hat{U}_{n,\Delta t}^\dagger \cdots\right), \\
p_e'(t_\pi) &= \sum_n \rho_{nn}\langle e| \left(\cdots \hat{U}_{n,\Delta t}\varepsilon(\hat{U}_{n,\Delta t}\varepsilon(|g\rangle\langle g|)\hat{U}_{n,\Delta t}^\dagger)\hat{U}_{n,\Delta t}^\dagger \cdots\right)|e\rangle. \qquad \text{(C1)}
\end{aligned}$$

Note that the error channel (ε) only acts on the qubit, while the Hamiltonian of $\hat{U}_{\Delta t}$ acts on the Fock state $|n\rangle$ of the cavity. This allows for the simplification where a specific $\hat{U}_{n,\Delta t} \equiv e^{-i\hat{H}_n\Delta t}$ with $\hat{H}_n = \Delta_n|e\rangle\langle e| + (\Omega/2)\hat{\sigma}_y$ only acts on the qubit. The third line then presents the qubit-probability and cavity-excitation mapping, analogous to Eq. (B4) in Appendix B 1. This last expression can be written as $p_e'(t_\pi) = \sum_n \rho_{nn} w_n$, where the weight

w_n is the probability of the qubit being in the excited state after starting in $|g\rangle$. With the same selectivity assumption ($\chi/\Omega \gg 1$) as used in Appendix B 1, we arrive at

$$p_e'(\Delta_n = 0) \approx \rho_{nn} w, \tag{C2}$$

where the weight w is obtained in the same way as w_n but with a Hamiltonian $\hat{H} = (\Omega/2)\hat{\sigma}_y$ under dephasing. In this way, qubit dephasing scales the probability of all excitation numbers ρ_{nn} with the same magnitude. The weight w can be computed by solving the qubit dynamics under decoherence with the Lindblad master equation

$$\dot{\rho} = -i[H, \rho] + \hat{J}\rho\hat{J}^\dagger - \frac{1}{2}\{\rho, \hat{J}^\dagger\hat{J}\}, \tag{C3}$$

where $\hat{J} \equiv \sqrt{2/T_\phi}|e\rangle\langle e|$ is the jump operator for qubit dephasing, which translates to solving a second-order differential equation. After a tedious but straightforward calculation, we obtain a weight

$$\begin{aligned} w &= \frac{1}{2}\left(1 - e^{-\gamma\pi}\left(\cos\left(\sqrt{1-\gamma^2}\pi\right) + \frac{\gamma}{\sqrt{1-\gamma^2}}\sin\left(\sqrt{1-\gamma^2}\pi\right)\right)\right), \quad \text{for } \gamma < 1, \\ w &= \frac{1}{2}\left(1 - e^{-\gamma\pi}\left(\frac{\gamma+\sqrt{\gamma^2-1}}{2\sqrt{\gamma^2-1}}e^{\sqrt{\gamma^2-1}\pi} + \frac{-\gamma+\sqrt{\gamma^2-1}}{2\sqrt{\gamma^2-1}}e^{-\sqrt{\gamma^2-1}\pi}\right)\right), \quad \text{for } \gamma > 1, \end{aligned} \tag{C4}$$

where $\gamma \equiv 1/(2T_\phi\Omega)$ and the first (second) line represents the small (over-) dephasing case. In most cases, we have small dephasing such that (up to the second order in γ) the weight can be approximated as $w \approx (1 + e^{-\gamma\pi})/2$, and the mapping

$$p_n' \equiv p_e'(\Delta_n = 0) \approx \rho_{nn} \times \frac{1}{2}(1 + e^{-\frac{t_\pi}{2T_\phi}}). \tag{C5}$$

For the case of parity mapping, we simplify the calculation by assuming that dephasing is only present during the waiting time, since the $\pi/2$ pulses are much shorter in time. Hence, only the off-diagonal elements of the qubit density matrix are degraded by a factor e^{-t_w/T_ϕ}. The parity mapping consequently yields

$$P' \equiv 2p_e' - 1 = P_{\text{id}} \times e^{-\frac{t_w}{T_\phi}}. \tag{C6}$$

2. Qubit thermal population

Both excitation-number and parity mapping rely on the qubit initialized in the ground state $|g\rangle$. In reality, the qubit might be in a mixed state, with some probability being in the excited state before the mapping protocol is performed. Experimentally, the residual excited probability is mainly due to imperfect state preparation from the GRAPE pulses. To analyze the effect of this, we assume an initial state $\rho_{\text{qth}} \otimes \rho$, where $\rho_{\text{qth}} = (1-\lambda)|g\rangle\langle g| + \lambda|e\rangle\langle e|$, where λ is the probability of the qubit being in the excited state.

For excitation-number mapping, following the derivation in Eq. (B3) results in

$$p_e'(t) = \sum_n \rho_{nn}\langle e|e^{-i\hat{H}_n t}((1-\lambda)|g\rangle\langle g| + \lambda|e\rangle\langle e|) \otimes \rho e^{i\hat{H}_n t}|e\rangle. \tag{C7}$$

When evaluating the term $\langle e|\cdots|e\rangle$ in the summation at $t_\pi = \pi/\Omega$, we note two cases: (i) $\Delta_n = 0$ and (ii) $\Delta_n \gg \Omega$. For (i), the unitary is $e^{-i\hat{H}_n t_\pi} = e^{-i\hat{\sigma}_y\pi/2}$, which is a π rotation of the qubit. In this case, we have $\langle e|\cdots|e\rangle = 1 - \lambda$. For (ii), the unitary is $e^{-i\hat{H}_n t_\pi} \approx e^{-i\Delta_n t|e\rangle\langle e|}$. In this case, we have $\langle e|\cdots|e\rangle = \lambda$. Both (i) and (ii) are essential for excitation-number mapping:

$$p_n' \equiv p_e'(\Delta_n = 0) = (1-\lambda)\rho_{nn} + \lambda\sum_{j\neq n}\rho_{jj}, \tag{C8}$$

where the biggest contribution comes from ρ_{nn} with weight $1 - \lambda$ accompanied by small contributions from other excitation probabilities ρ_{jj} ($j \neq n$), each with weight λ as they are far away from $\Delta_n = 0$ (we assume $\Delta_n \gg \Omega$). By noting that $\sum_{j\neq n}\rho_{jj} = (1 - \rho_{nn})$ ($\text{Tr}(\rho) = 1$), we have

$$p_n' = \rho_{nn}(1 - 2\lambda) + \lambda, \tag{C9}$$

where we see that the thermal population of the qubit introduces a scaling and offset error to the ideal excitation number. However, if λ is known, the excitation-number measurement can be corrected as $\rho_{nn} = (p_n' - \lambda)/(1 - 2\lambda)$, where p_n' is the measured value.

For parity mapping, solving the dynamics starting with the qubit in a mixed state can be done by separating the two possible initial states. When the qubit starts in $|g\rangle$ state, with a probability $1-\lambda$, the parity is given by Eq. (B7). On the other hand, when the qubit starts in $|e\rangle$ state, with a probability λ, we analogously arrive at $p_e = (1-P_{\rm id})/2$. Combining both contributions,

$$p_e' = (1-\lambda)\frac{1+P_{\rm id}}{2} + \lambda\frac{1-P_{\rm id}}{2}, \tag{C10}$$

which leads to a parity

$$P' \equiv 2p_e' - 1 = (1-2\lambda)P_{\rm id}. \tag{C11}$$

The thermal population of the qubit only results in a scaling error that can be corrected as $P_{\rm id} = P'/(1-2\lambda)$, where P' is the measured value and λ has been previously characterized.

APPENDIX D: ESTIMATOR FOR ρ

1. Linear inversion

The oldest and simplest procedure to build an estimator for ρ is called linear inversion. This method consists of interpreting the relative frequencies of measurement outcomes as probabilities and then inverting Born's rule through a least-squares (LS) inversion to obtain a $\rho_{\rm LS}$ that predicts these probabilities.

Born's rule relates the outcome probability p_k of a certain measurement observable $\hat{E}_k$ to ρ as

$$p_k = {\rm Tr}(\rho\hat{E}_k). \tag{D1}$$

Upon many measurement repetitions, we build a histogram and approximate each p_k with the corresponding relative frequency of the outcome k. For ORENS, each measurement observable is defined by a displacement $\hat{D}(\alpha_k)$ and an excitation number n, and can be written as

$$\hat{E}_{n,\alpha_k} = \hat{D}_{\alpha_k}|n\rangle\langle n|\hat{D}_{-\alpha_k}, \tag{D2}$$

with the corresponding probability $p_n = {\rm Tr}(\hat{E}_{n,\alpha_k}\rho)$. Let us define a $(D^2-1)\times D^2$ measurement matrix M to describe the set of ORENS measurements as

$$M = \begin{pmatrix} \vec{E}_{n,\alpha_1} \\ \vec{E}_{n,\alpha_2} \\ \vdots \\ \vec{E}_{n,\alpha_{D^2-1}} \end{pmatrix}, \tag{D3}$$

where D is the cutoff dimension of the cavity state and $\vec{E}$ is the row-wise vectorized form of $\hat{E}$ (truncated to $D\times D$). Vectorizing ρ column-wise to obtain $\vec{\rho}$ of length D^2 and writing the outcome probabilities as a vector of length D^2-1, $\vec{p}$, we can then write the matrix equation

$$M\vec{\rho} = \vec{p}. \tag{D4}$$

Linear inversion corresponds to inverting this system using the observed relative frequencies $\vec{p}$ to derive $\vec{\rho}$. Because M is not a square matrix, the system is solved using the Moore-Penrose pseudoinverse as

$$\vec{\rho} = (M^\dagger M)^{-1}M^\dagger\vec{p} = M^+\vec{p}. \tag{D5}$$

However, $M^\dagger M$ is not invertible, since $\det(M^\dagger M) = 0$, as M has fewer rows than columns [53]. Consequently, in our work, we parametrize and vectorize ρ using the real and imaginary components of the upper triangular off-diagonal elements of ρ, as well as the diagonal elements except for the last one. This parametrization yields a vector $\vec{Y}$ of length D^2-1, which completely characterizes the state.

Different parametrizations of ρ are related linearly, i.e., $\vec{\rho} = K\vec{Y} + \Theta$. Using Eq. (D4) and defining

$$\begin{aligned} \mathcal{M} &\equiv MK, \\ V &\equiv M\Theta, \end{aligned} \tag{D6}$$

we obtain the modified linear equation and its inverse (linear inversion),

$$\begin{aligned} \vec{p} &= \mathcal{M}\vec{Y} + V, \\ \vec{Y}_{\rm est} &= \mathcal{M}^{-1}(\vec{p} - V), \end{aligned} \tag{D7}$$

where the matrix $\mathcal{M}$ is now a square matrix. The density matrix corresponding to $\vec{Y}_{\rm est}$ is the least-squares estimate $\rho_{\rm LS}$.

2. Optimizing the set of measurements

For a given truncation dimension D, we then perform a classical optimization algorithm to obtain the set of measurement observables. The optimization involves determining the set of displacement points $\{\alpha_k\}_{k=1}^{D^2-1}$ and the corresponding excitation number n to measure. This set enables the reconstruction of *any* state that can be truncated within dimension D, without any further assumption.

The optimization is executed using a gradient-descent method, where the cost function is the condition number (CN) of the matrix $\mathcal{M}$ in Eq. (D7), defined as

$$||\mathcal{M}||\ ||\mathcal{M}^{-1}||, \tag{D8}$$

where $||\cdot||$ denotes the Euclidean norm. The CN quantifies the worst-case error amplification from the observables to the estimated state, i.e., when solving for $\vec{Y}_{est}$ in Eq. (D7) given the observables $\vec{p}$ [30]. With Eqs. (D2), (D3), and

(D6), the CN can be computed numerically given the choice of photon number n for measurement and the set of displacement points.

Our gradient-descent algorithm begins with randomly selected initial displacement points and a fixed photon number n for the excitation measurement. The algorithm iteratively updates the displacement points based on the gradient of the cost function, ultimately providing an optimized set of displacement points $\{\alpha_k\}$ corresponding to the minimum CN, for a given n. We repeat this gradient descent for all photon numbers $n \in [0, D-1]$ to find the n that yields the smallest CN. We consistently find that, for a given dimension D, the optimal CN is obtained for $n = D-1$. Lastly, we examine the scenario in which each of the D^2-1 displacements employs a distinct excitation-number measurement n. However, this approach does not significantly reduce the CN, although it does introduce increased experimental complexity.

The sets of displacement points that yield a low CN are not unique, as exemplified in Fig. 12(a), where we show two different sets of optimized displacement points for ORENS given $D = 6$ with a CN of approximately 3.1. Furthermore, when used for state reconstruction, simulations [similar to the one in Fig. 4(a)] show that both sets offer an average fidelity > 0.98. The fact that the displacement sets are not unique facilitates the incorporation of hardware constraints into the optimization algorithm. For instance, to mitigate readout distortion resulting from the residual cross-Kerr interaction when the cavity contains a high photon number, a constraint of $|\alpha_k| \leq 2$ can be readily integrated into the optimization.

For benchmarking purposes, we have also obtained the optimized displacement points for the Husimi-Q and Wigner functions. The Husimi-Q function corresponds to ORENS in the case in which the excitation number is constrained to measuring only zero excitations, $n = 0$. For the Wigner function, the photon-number projection operator $|n\rangle\langle n|$ in Eq. (D2) is simply replaced by the parity operator $\hat{P}$.

The CN achieved by ORENS across dimensions is comparable to that of Wigner, indicating that they have near-equivalent theoretical reconstruction capabilities [see Fig. 12(b)]. However, the optimal CN for the Husimi-Q function scales unfavorably beyond $D > 2$, illustrating that robust reconstruction is infeasible for large states.

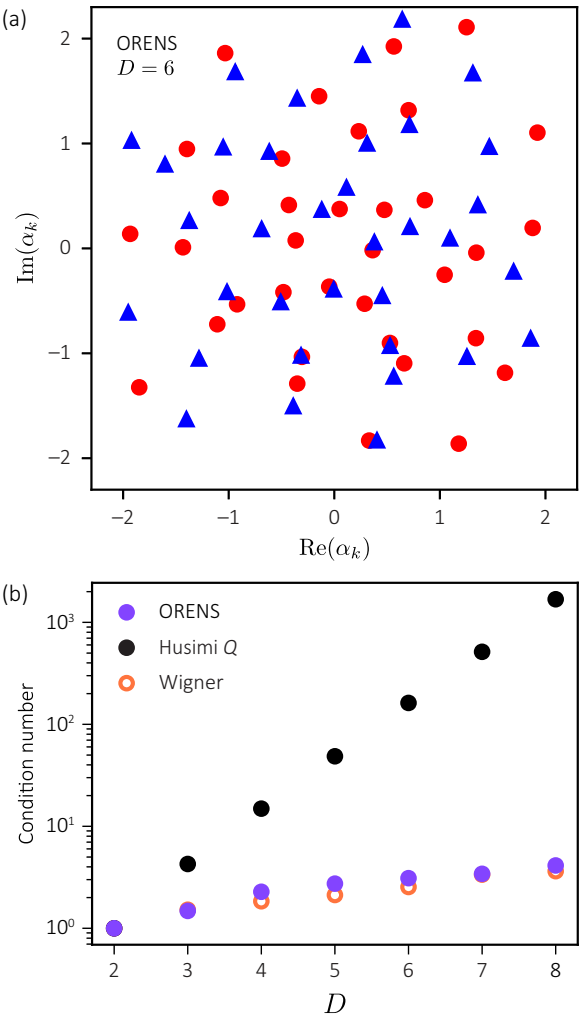


FIG. 12. (a) The optimized displacement points of ORENS for a given truncation dimension $D = 6$, all with excitation-number $n = 5$ measurement. Two different exemplary sets (blue triangles and red circles) are displayed. Both of these sets give condition numbers of approximately 3.1 and comparable reconstruction fidelities. (b) The optimized condition numbers of ORENS, Wigner, and Husimi Q versus the truncation dimension of ρ.

3. Bayesian inference

Accurately inferring the quantum state of a system from measurement outcomes is a crucial task in quantum state reconstruction. In this section, we will motivate the use of Bayesian inference to process measurement outcomes and build the optimal estimator for ρ. For a deeper analysis, see Ref. [40], and for details on the specific methodology used in our work, see Ref. [27].

The most notable limitation of linear inversion is that the estimated ρ_{LS} frequently has negative eigenvalues, indicating that it cannot represent a physical state. Maximum-likelihood estimation (MLE) has been adopted as a convenient way to impose physicality on ρ_{LS}, and has been the dominant approach to quantum state reconstruction in recent years. Intuitively, it returns a single non-negative state ρ_{MLE} that fits the observed data $\mathcal{D}$ as precisely as possible by maximizing the likelihood function,

$$\rho_{\mathrm{MLE}} = \arg\max_{\rho} L_{\mathcal{D}}(\rho), \tag{D9}$$

where $L_{\mathcal{D}}(\rho) \propto p(\mathcal{D}|\rho)$. An efficient way to obtain ρ_{MLE} is described in Ref. [41]. However, the MLE method does

TABLE III. The average ORENS reconstruction fidelity using BME and MLE across all Fock-state superpositions for a given cutoff dimension, as plotted in Fig. 4(a).

D	$\bar{F}_{\mathrm{BME}}$	$\bar{F}_{\mathrm{MLE}}$	$\Delta = \bar{F}_{\mathrm{BME}} - \bar{F}_{\mathrm{MLE}}$
2	0.992	0.987	0.005
3	0.988	0.979	0.009
4	0.973	0.958	0.015
5	0.950	0.933	0.017
6	0.939	0.918	0.021

TABLE IV. The average ORENS reconstruction fidelity using BME and MLE across all four cat states for different engineered qubit T_ϕ, as plotted in Fig. 4(b).

$T_\phi(\mu\mathrm{s})$	$\bar{F}_{\mathrm{BME}}$	$\bar{F}_{\mathrm{MLE}}$	$\Delta = \bar{F}_{\mathrm{BME}} - \bar{F}_{\mathrm{MLE}}$
22.4	0.947	0.932	0.015
10.4	0.946	0.927	0.018
3.48	0.944	0.931	0.013
1.02	0.875	0.856	0.019
0.535	0.867	0.851	0.016

not quantify the level of uncertainty of the result and, most critically, ρ_{MLE} often has zero eigenvalues. Consequently, it predicts exactly zero probability for every measurement outcome $|\psi\rangle\langle\psi|$ such that $\langle\psi|\rho|\psi\rangle = 0$. This implication of absolute certainty that a certain outcome will not be observed cannot reasonably be justified by a finite amount of data. The underlying flaw is that maximizing the likelihood is frequentist by nature; it interprets the observed relative frequencies of the measurement outcomes as probabilities and then seeks to fit the probabilities as precisely as possible. However, the goal of state estimation extends beyond explaining the data to predicting future evolutions and states. Thus, estimation should involve the knowledge of the system being estimated, especially its uncertainty.

In our work, we have employed a Bayesian-inference technique stemming from a different perspective on statistics that (1) considers many of the possible ρ, (2) accounts for experimental uncertainty explicitly through Bayes' rule, and (3) guarantees the most accurate estimate of the true ρ that can be made from the data [40,54,55]. Parametrizing $\rho(\mathbf{x})$ by some vector $\mathbf{x}$, such that any value of $\mathbf{x}$ within its support returns a physical ρ, Bayes' theorem states that the posterior probability distribution of $\mathbf{x}$ follows as

$$\pi(\mathbf{x}) = \frac{1}{Z} L_{\mathcal{D}}(\mathbf{x})\pi_0(\mathbf{x}), \tag{D10}$$

where $L_{\mathcal{D}}(\mathbf{x})$ is the same likelihood as in MLE, $\pi_0(\mathbf{x})$ is the prior distribution that encapsulates any knowledge or beliefs about ρ before the experiment, and Z is a normalizing constant. This posterior distribution gives us access to the expectation value of any function ϕ of ρ via

$$\langle\phi(\rho)\rangle = \int d\mathbf{x}\pi(\mathbf{x})\phi(\rho(\mathbf{x})). \tag{D11}$$

Evaluating integrals of this form is numerically challenging due to the high dimensionality and complicated features. We overcome this challenge by employing the efficient Bayesian-inference strategy [27], which is computationally practical and straightforward to implement through a combination of well-chosen parametrization of ρ and likelihood, and the Markov-chain Monte Carlo (MCMC) sampling algorithm. Intuitively, the algorithm draws random samples of possible ρ from a distribution across all physical states. These states are weighted by a pseudolikelihood function that scales inversely with the distance between the sample and ρ_{LS}. These samples allow us to estimate any function of ρ via

$$\langle\phi(\rho)\rangle \approx \frac{1}{R}\sum_{r=1}^{R}\phi(\rho_r), \tag{D12}$$

where R is the total number of MCMC samples. In detail, we have chosen the following parameters for Bayesian inference: $\alpha = 1$ for a uniform prior on all possible physical density matrices; $\sigma = 1/N$, with $N = 1000(D^2 - 1)$, as the variance for the pseudolikelihood function around ρ_{LS}; and 2^{10} MCMC samples with thinning parameter 2^7 to reduce serial correlation in the chain.

All simulated and experimental fidelities in this work have been calculated with the Bayesian mean estimator (BME), defined as

$$\rho_{\mathrm{BME}} = \frac{1}{R}\sum_{r=1}^{R}\rho_r, \tag{D13}$$

which stands as the most accurate estimator of the true ρ. The error bars in the main text represent the standard deviation across ρ_{BME} fidelities of all the set of reconstructed states (either Fock states or cat states). We demonstrate how the performance of BME surpasses that of MLE, as expected, in Tables III and IV for ORENS reconstruction across both dimensions and decoherence regimes, respectively.

[1] A. Copetudo, C. Y. Fontaine, F. Valadares, and Y. Y. Gao, Shaping photons: Quantum information processing with bosonic cQED, Appl. Phys. Lett. **124**, 080502 (2024).

[2] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, Gaussian quantum information, Rev. Mod. Phys. **84**, 621 (2012), publisher: American Physical Society .

[3] X. Pan, P. Song, and Y. Y. Gao, Continuous-variable quantum computation in circuit QED, Chinese Phys. Lett. **40**,

110303 (2023), publisher: Chinese Physical Society and IOP Publishing Ltd.

[4] A. Joshi, K. Noh, and Y. Y. Gao, Quantum information processing with bosonic qubits in circuit QED, Quantum Sci. Technol. **6**, 033001 (2021), publisher: IOP Publishing.

[5] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret, Quantum error correction of a qubit encoded in grid states of an oscillator, Nature **584**, 368 (2020), publisher: Nature Publishing Group.

[6] Y. Y. Gao, B. J. Lester, K. S. Chou, L. Frunzio, M. H. Devoret, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, Entanglement of bosonic modes through an engineered exchange interaction, Nature **566**, 509 (2019), publisher: Nature Publishing Group.

[7] W. Wang, Y. Wu, Y. Ma, W. Cai, L. Hu, X. Mu, Y. Xu, Z.-J. Chen, H. Wang, Y. P. Song, H. Yuan, C.-L. Zou, L.-M. Duan, and L. Sun, Heisenberg-limited single-mode quantum metrology in a superconducting circuit, Nat. Commun. **10**, 4382 (2019), publisher: Nature Publishing Group.

[8] K. Liu, Y. Xu, W. Wang, S.-B. Zheng, T. Roy, S. Kundu, M. Chand, A. Ranadive, R. Vijay, Y. Song, L. Duan, and L. Sun, A twofold quantum delayed-choice experiment in a superconducting circuit, Sci. Adv. **3**, e1603159 (2017), publisher: American Association for the Advancement of Science.

[9] L. G. Lutterbach and L. Davidovich, Method for direct measurement of the Wigner function in cavity QED and ion traps, Phys. Rev. Lett. **78**, 2547 (1997), publisher: American Physical Society.

[10] P. Bertet, A. Auffeves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. M. Raimond, and S. Haroche, Direct measurement of the Wigner function of a one-photon Fock state in a cavity, Phys. Rev. Lett. **89**, 200402 (2002), publisher: American Physical Society.

[11] K. Husimi, Some formal properties of the density matrix, Proceedings of the Physico-Mathematical Society of Japan. 3rd Series **22**, 264 (1940).

[12] T. Opatrný and D.-G. Welsch, Density-matrix reconstruction by unbalanced homodyning, Phys. Rev. A **55**, 1462 (1997), publisher: American Physical Society.

[13] G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Observation of quantum state collapse and revival due to the single-photon Kerr effect, Nature **495**, 205 (2013), number: 7440 Publisher: Nature Publishing Group.

[14] O. Landon-Cardinal, L. C. G. Govia, and A. A. Clerk, Quantitative tomography for continuous variable quantum systems, Phys. Rev. Lett. **120**, 090501 (2018), publisher: American Physical Society.

[15] K. He, M. Yuan, Y. Wong, S. Chakram, A. Seif, L. Jiang, and D. I. Schuster, Efficient multimode Wigner tomography, Nat. Commun. **15**, 4138 (2024).

[16] S. Chakram, K. He, A. V. Dixit, A. E. Oriani, R. K. Naik, N. Leung, H. Kwon, W.-L. Ma, L. Jiang, and D. I. Schuster, Multimode photon blockade, Nat. Phys. **18**, 879 (2022), number: 8 Publisher: Nature Publishing Group.

[17] C. Shen, R. W. Heeres, P. Reinhold, L. Jiang, Y.-K. Liu, R. J. Schoelkopf, and L. Jiang, Optimized tomography of continuous variable systems using excitation counting, Phys. Rev. A **94**, 052327 (2016).

[18] C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. M. Sliwa, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, A Schrödinger cat living in two boxes, Science **352**, 1087 (2016), publisher: American Association for the Advancement of Science.

[19] B. Calkins, P. L. Mennea, A. E. Lita, B. J. Metcalf, W. Steven Kolthammer, A. Lamas-Linares, J. B. Spring, P. C. Humphreys, R. P. Mirin, J. C. Gates, P. G. R. Smith, I. A. Walmsley, T. Gerrits, and S. W. Nam, High quantum-efficiency photon-number-resolving detector for photonic on-chip information processing, Opt. Express **21**, 22657 (2013), publisher: Optica Publishing Group.

[20] M. Brune, F. Schmidt-Kaler, A. Maali, J. Dreyer, E. Hagley, J. M. Raimond, and S. Haroche, Quantum Rabi oscillation: A direct test of field quantization in a cavity, Phys. Rev. Lett. **76**, 1800 (1996), publisher: American Physical Society.

[21] C. Guerlin, J. Bernu, S. Deléglise, C. Sayrin, S. Gleyzes, S. Kuhr, M. Brune, J.-M. Raimond, and S. Haroche, Progressive field-state collapse and quantum non-demolition photon counting, Nature **448**, 889 (2007), number: 7156 Publisher: Nature Publishing Group.

[22] 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, S. M. Girvin, and R. J. Schoelkopf, Resolving photon number states in a superconducting circuit, Nature **445**, 515 (2007), publisher: Nature Publishing Group.

[23] H. Wang, M. Hofheinz, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, M. Weides, J. Wenner, A. N. Cleland, and J. M. Martinis, Decoherence dynamics of complex photon states in a superconducting circuit, Phys. Rev. Lett. **103**, 200404 (2009), publisher: American Physical Society.

[24] D. Leibfried, D. M. Meekhof, B. E. King, C. Monroe, W. M. Itano, and D. J. Wineland, Experimental determination of the motional quantum state of a trapped atom, Phys. Rev. Lett. **77**, 4281 (1996), publisher: American Physical Society.

[25] S. An, J.-N. Zhang, M. Um, D. Lv, Y. Lu, J. Zhang, Z.-Q. Yin, H. T. Quan, and K. Kim, Experimental test of the quantum Jarzynski equality with a trapped-ion system, Nat. Phys. **11**, 193 (2015), number: 2 Publisher: Nature Publishing Group.

[26] H.-Y. Lo, D. Kienzler, L. de Clercq, M. Marinelli, V. Negnevitsky, B. C. Keitch, and J. P. Home, Spin-motion entanglement and state diagnosis with squeezed oscillator wavepackets, Nature **521**, 336 (2015), number: 7552 Publisher: Nature Publishing Group.

[27] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, A practical and efficient approach for Bayesian quantum state estimation, New J. Phys. **22**, 063038 (2020).

[28] A. I. Lvovsky and M. G. Raymer, Continuous-variable optical quantum-state tomography, Rev. Mod. Phys. **81**, 299 (2009).

[29] S. Deleglise, I. Dotsenko, C. Sayrin, J. Bernu, M. Brune, J.-M. Raimond, and S. Haroche, Reconstruction of non-classical cavity field states with snapshots of their decoherence, Nature **455**, 510 (2008).

[30] R. Bhatia, *Matrix Analysis*, Graduate Texts in Mathematics (Springer, New York, 1997), Vol. 169.

[31] J. S. Lundeen, A. Feito, H. Coldenstrodt-Ronge, K. L. Pregnell, Ch. Silberhorn, T. C. Ralph, J. Eisert, M. B. Plenio, and I. A. Walmsley, Tomography of quantum detectors, Nat. Phys. **5**, 27 (2009), number: 1 Publisher: Nature Publishing Group.

[32] L. Zhang, H. B. Coldenstrodt-Ronge, A. Datta, G. Puentes, J. S. Lundeen, X.-M. Jin, B. J. Smith, M. B. Plenio, and I. A. Walmsley, Mapping coherence in measurement via full quantum tomography of a hybrid optical detector, Nat. Photonics **6**, 364 (2012), number: 6 Publisher: Nature Publishing Group.

[33] T. B. Smith, Generalized Q-functions, J. Phys. A: Math. Gen. **39**, 13747 (2006).

[34] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, Nat. Commun. **8**, 94 (2017), number: 1 Publisher: Nature Publishing Group.

[35] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, Annu. Rev. Condens. Matter Phys. **11**, 369 (2020).

[36] O. Gargiulo, S. Oleschko, J. Prat-Camps, M. Zanner, and G. Kirchmair, Fast flux control of 3D transmon qubits using a magnetic hose, Appl. Phys. Lett. **118**, 012601 (2021).

[37] M. D. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. L. T. Plourde, Tunable superconducting qubits with flux-independent coherence, Phys. Rev. Appl. **8**, 044003 (2017), publisher: American Physical Society.

[38] A. P. Sears, A. Petrenko, G. Catelani, L. Sun, H. Paik, G. Kirchmair, L. Frunzio, L. I. Glazman, S. M. Girvin, and R. J. Schoelkopf, Photon shot noise dephasing in the strong-dispersive limit of circuit QED, Phys. Rev. B **86**, 180504 (2012), publisher: American Physical Society.

[39] R. Blume-Kohout and P. Hayden, Accurate quantum state estimation via "Keeping the experimentalist honest," arXiv:quant-ph/0603116.

[40] R. Blume-Kohout, Optimal, reliable estimation of quantum states, New J. Phys. **12**, 043034 (2010).

[41] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient method for computing the maximum-likelihood quantum state from measurements with additive Gaussian noise, Phys. Rev. Lett. **108**, 070502 (2012), publisher: American Physical Society.

[42] J. Joo, W. J. Munro, and T. P. Spiller, Quantum metrology with entangled coherent states, Phys. Rev. Lett. **107**, 083601 (2011), publisher: American Physical Society.

[43] T. C. Ralph, A. Gilchrist, G. J. Milburn, W. J. Munro, and S. Glancy, Quantum computation with optical coherent states, Phys. Rev. A **68**, 042319 (2003), publisher: American Physical Society.

[44] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, A. Keller, G. Refael, J. Preskill, L. Jiang, A. H. Safavi-Naeini, O. Painter, and F. G. S. L. Brandão, Building a fault-tolerant quantum computer using concatenated cat codes, PRX Quantum **3**, 010329 (2022), publisher: American Physical Society.

[45] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Extending the lifetime of a quantum bit with error correction in superconducting circuits, Nature **536**, 441 (2016), publisher: Nature Publishing Group.

[46] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia, and S. M. Girvin, Bias-preserving gates with stabilized cat qubits, Sci. Adv. **6**, eaay5901 (2020), publisher: American Association for the Advancement of Science.

[47] https://github.com/clarayfontaine/ORENS_bosonic_state_reconstruction

[48] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, Phys. Rev. A **69**, 062320 (2004), publisher: American Physical Society.

[49] S. M. Girvin, in *Quantum Machines: Measurement and Control of Engineered Quantum Systems*, edited by Michel Devoret, Benjamin Huard, Robert Schoelkopf, and Leticia F. Cugliandolo (Oxford University Press, Oxford, 2014), 1st ed., p. 113.

[50] Z. K. Minev, Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, and M. H. Devoret, Energy-participation quantization of Josephson circuits, npj Quantum Inf. **7**, 1 (2021), number: 1 Publisher: Nature Publishing Group.

[51] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic Purcell filter, Phys. Rev. Appl. **17**, 044016 (2022).

[52] 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, Phys. Rev. Appl. **5**, 011001 (2016), publisher: American Physical Society.

[53] T. Krisnanda, H. Xu, S. Ghosh, and T. C. H. Liew, Tomographic completeness and robustness of quantum reservoir networks, Phys. Rev. A **107**, 042402 (2023).

[54] B. P. Williams and P. Lougovski, Quantum state estimation when qubits are lost: A no-data-left-behind approach, New J. Phys. **19**, 043003 (2017), publisher: IOP Publishing.

[55] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, New J. Phys. **18**, 033024 (2016), publisher: IOP Publishing.