


# Reservoir Computing Approach to Quantum State Measurement

Gerasimos Angelatos<sup>✉,\*</sup>, Saeed A. Khan<sup>✉</sup>, and Hakan E. Türeci<sup>✉</sup>

*Department of Electrical Engineering, Princeton University, Princeton, New Jersey 08544, USA*

 (Received 3 December 2020; revised 4 August 2021; accepted 28 October 2021; published 29 December 2021)

Efficient quantum state measurement is important for maximizing the extracted information from a quantum system. For multiqubit quantum processors, in particular, the development of a scalable architecture for rapid and high-fidelity readout remains a critical unresolved problem. Here we propose reservoir computing as a resource-efficient solution to quantum measurement of superconducting multiqubit systems. We consider a small network of Josephson parametric oscillators, which can be implemented with minimal device overhead and in the same platform as the measured quantum system. We theoretically analyze the operation of such a device as a reservoir computer to classify stochastic time-dependent signals subject to quantum statistical features. We apply this reservoir computer to the task of multinomial classification of measurement trajectories from joint multiqubit readout. For a 2-qubit dispersive measurement under realistic conditions we demonstrate a classification fidelity reliably exceeding that of an optimal linear filter using only 2–5 reservoir nodes, while simultaneously requiring far less calibration data—as little as a few shots per state. We understand this remarkable performance through an analysis of the network dynamics and develop an intuitive picture of reservoir processing generally. Finally, we demonstrate how to operate this device to perform 2-qubit state tomography and continuous parity monitoring with equal effectiveness and ease of calibration. This reservoir processor avoids computationally intensive training common to other machine learning frameworks and can be implemented as an integrated cryogenic superconducting device for low-latency processing of quantum signals on the computational edge.

DOI: [10.1103/PhysRevX.11.041062](https://doi.org/10.1103/PhysRevX.11.041062)

Subject Areas: Complex Systems, Quantum Information

## I. INTRODUCTION

Rapid and high-fidelity single-shot readout is a fundamental element of the manipulation and processing of quantum information. In superconducting circuit quantum processors [1,2], this requires a careful calibration of the entire measurement chain, including cryogenic and room-temperature amplifiers, circulators, attenuators, and room-temperature electronics. This calibration becomes particularly resource intensive for readout systems attached to multiqubit quantum processors. The optimization of quantum state readout has therefore been the focus of considerable ongoing research [3–12], involving a delicate balance of competing requirements: fidelity and speed.

For single-qubit readout, optimal filtering approaches [3] and hardware architectures have been developed and implemented to achieve fast and high-fidelity measurements

without affecting qubit coherence [13]. More recently, recognizing that the quantum measurement problem in its very essence is the classification of time-dependent voltage signals acquired at the end of a measurement chain, machine learning solutions have been investigated [14–17], and have shown an increase in single-qubit state discrimination by a few percent with respect to these conventional approaches [16,17]. For measurement in multiqubit systems, however, the optimization and calibration of a readout system presents a difficult hardware design as well as a computationally intensive signal processing problem [5,6]; measurement cross talk, in particular, imposes significant limitations on device scaling [7,17]. Here we propose a novel hardware-efficient approach to high-fidelity multiqubit readout based on reservoir computing and theoretically analyze its training-based operation.

Reservoir computing is a machine learning framework for the processing of time-dependant data [18–21]. It is founded on the idea that any sufficiently complex and high-dimensional dynamical system, where only the linear output layer is optimized, can have the same computational capacity as a recurrent neural network in approximating arbitrary functions or functionals [22,23]. Vastly different physical systems have been employed as reservoir computers (RCs) for applications such as forecasting and classification [24–30].

\*Corresponding author.  
ga4@princeton.edu

*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.*

The field of reservoir computing has recently expanded to include quantum systems [31–36]. However, the application of reservoir processing to the problem of quantum measurement has not yet been explored.

Our goal in this paper is threefold: (1) describe a reservoir computing approach to quantum measurement that utilizes a physical system with recurrent connections, (2) analyze its efficacy for fast and high-fidelity readout and monitoring of multiple qubits simultaneously, and (3) propose a superconducting preprocessor based on a network of Josephson parametric oscillators (JPOs) to enable hardware-efficient and low-latency multiqubit measurement. While we discuss this approach for a multiqubit superconducting platform and corresponding Josephson-junction-based superconducting reservoir, we anticipate that the techniques are general enough to also be applicable for a broader class of quantum systems, measurement tasks, and reservoirs.

Conventional RC wisdom suggests that very high-dimensional dynamical systems are necessary for strong computational performance, with  $10^2$ – $10^3$  nodes typically used [24,37] in software or time-delay architectures (where there is less overhead associated with increasing the size of the network). Here, we show that a small physical RC (2–5 nodes) is able to classify 2-qubit measurement trajectories with a fidelity higher than is achievable under the same conditions with optimal conventional filtering approaches. Equally strong performance is seen across a variety of quantum systems and measurement tasks, without requiring any modification of the RC.

This non-von Neumann-architecture computer can be implemented in the same hardware platform as the target quantum system with minimal overhead, providing a uniquely low-latency approach to quantum measurement. An important conclusion of our study is that the physical RC we consider requires dramatically less training data than a readout system that is calibrated using an optimal matched filter. Our results indicate that a cryogenic readout device should provide a rapid, robust, and autonomous preprocessor for quantum state measurement. Such analog processors are capable of operating on timescales orders of magnitude faster than digital processors in head-to-head comparisons on the same computational task [38], and enable signal processing on the “computational edge” [21,29], significantly reducing computational costs.

In addition, we demonstrate that the RC provides a model-independent approach to quantum state measurement, ideal for multiqubit systems whose readout chains are projected to become increasingly more complex. The readout problem we consider here is one of retrodicting certain features of the initial state of a measured quantum system (QS), based on information obtained after a specific quantum process, such as the scattering of a probe pulse off of the QS. Within superconducting circuit implementations, the most widely employed readout setup is that of

quantum nondemolition (QND) dispersive measurement [3–8]; however, actual hardware implementations exhibit non-QND effects and experimental imperfections such as drift and cross talk. Such nonidealities are difficult to optimize in hardware and require several calibration experiments to characterize, making precise knowledge of the implemented physical model difficult to acquire. This lack of an accurate physical model generally rules out a description of the measurement chain via a stochastic master equation (SME), integration of which would predict *precisely* the measurement signal obtained given any initial state of the QS.

The difficulty of extracting the implemented physical model increases the appeal of model-independent approaches, such as linear filtering of the experimental data (discussed in Sec II); however, this typically requires a large amount of training data and is susceptible to errors from quantum jumps and qubit decay. Similarly, current machine learning approaches to readout are likely to be limited by computing capacity and training time when applied to systems more complex than a single qubit [14,17]. In this paper we apply a RC to a scenario in which constraints on QNDness and cross talk in multiqubit readout are relaxed, thus simulating quantum measurement with unoptimized hardware where the complexity of readout is relegated to the processing of acquired signals. We find that the RC is able to perform said processing with high fidelity and minimal computational cost, enabling a powerful model-free approach to readout. Specifically, we consider a situation where two qubits are measured simultaneously through a common resonator, without dedicated readout cavities and Purcell filters; our objective is not to propose this particular measurement scheme, but rather emphasize the reduced hardware and optimization overhead, and thus increased potential scalability, enabled through our reservoir processing approach generally.

We begin in Sec. II with a description of the joint dispersive readout task we consider in this work. We then give a high-level overview of our proposed RC quantum measurement system in Sec. III A, followed by a detailed description of the Kerr network RC model in Sec. III B. The dynamics and performance of a typical Kerr RC are presented in Sec. IV A. We then demonstrate the ability of a Kerr RC to quickly learn a quantum measurement task in Sec. IV B, enabling rapid readout chain calibration. In Sec. IV C, we use an analysis of the Kerr RC phase space dynamics to explain the strong performance of reservoirs with as few as two Kerr nodes and develop an intuitive picture of RC processing. In Sec. IV D we explore how behavior varies with system hyperparameters and present basic principles for the optimization of a hardware RC. Finally, in Sec. V we demonstrate how one can operate this reservoir processor to perform two additional important quantum information tasks: 2-qubit state tomography and continuous parity monitoring.

## II. QUANTUM STATE READOUT OF MULTIPLE QUBITS

Multiqubit readout presents a sufficiently difficult problem to quantitatively assess the advantage provided by more sophisticated signal processing techniques, which we therefore discuss in some detail below. Extensions of single-qubit quantum readout approaches to larger multiqubit systems through various multiplexing techniques have been extensively investigated [5–7,9]. A majority of these schemes rely on the premise of quantum nondemolition measurement through the dispersive readout technique. In the single-qubit variant, the binary state of the qubit ( $|0\rangle, |1\rangle$ ) is encoded in the amplitude and phase of a microwave pulse scattered from a readout resonator that is dispersively coupled to the qubit [4].

Cross talk, which we take here to be any effect of the measurement process on parts of the system one is not trying to measure, is a significant problem in multiqubit devices which scales very unfavorably with system size. Reducing readout errors generally requires precise calibration of readout pulses, carefully designed Purcell filters, and a chip layout that minimizes cross talk [6,7]. Such optimized calibration is difficult for a practical multiqubit quantum processor due to drifts in system parameters, and reducing cross talk imposes severe limitations on readout spectral bandwidth. Here we consider the joint dispersive readout scenario [5,9] where all qubits are coupled to the same mode of a common readout resonator, with the goal of measuring the combined state of all qubits in a single shot.

Our starting point is the multiqubit Jaynes-Cummings (JC) Hamiltonian:

$$\hat{\mathcal{H}}_{\text{JC}} = \Delta_c \hat{d}^\dagger \hat{d} + \epsilon(t)(\hat{d} + \hat{d}^\dagger) + \sum_j \frac{\Delta_{q,j}}{2} \hat{\sigma}_{z,j} + g_j(\hat{\sigma}_{+,j} \hat{d} + \hat{\sigma}_{-,j} \hat{d}^\dagger). \quad (1)$$

Here,  $\hat{d}$  and  $\hat{\sigma}_j$  are cavity field and qubit Pauli operators, respectively, and  $\epsilon(t)$  describes the amplitude of a coherent drive applied at carrier frequency  $\omega_d$ . We are in a frame rotating with  $\omega_d$ :  $\Delta_c = \omega_c - \omega_d$  and  $\Delta_{q,j} = \omega_{q,j} - \omega_d$  are the cavity and qubit detunings, respectively. The cavity qubit coupling, with strength  $g_j$ , is treated in the rotating wave approximation.

The JC Hamiltonian is accurate for weak readout pulses  $\epsilon(t)$ , where the role of other qubit energy levels can be ignored [39–41]. If the qubit-resonator detuning  $\delta_j = \omega_{q,j} - \omega_c$  is large, the cavity population remains below a critical photon number  $\langle \hat{d}^\dagger \hat{d} \rangle \ll \min(|\delta_j|/2g_j)^2$  for a sufficiently weak drive. Equation (1) can then be perturbatively transformed [42], yielding what we refer to here as the dispersive model:

$$\hat{\mathcal{H}}_D = \Delta_c \hat{d}^\dagger \hat{d} + \epsilon(t)(\hat{d} + \hat{d}^\dagger) + \sum_j \frac{\tilde{\Delta}_{q,j}}{2} \hat{\sigma}_{z,j} + \chi_j \hat{\sigma}_{z,j} \hat{d}^\dagger \hat{d} + \sum_{jk} J_{jk} \hat{\sigma}_{-,j} \hat{\sigma}_{+,k}, \quad (2)$$

valid to second order in  $g_j/\delta_j$ . Here  $\chi_j = g_j^2/\delta_j$  describes the dispersive shift of the cavity frequency due to the state of qubit  $j$ , and  $\tilde{\Delta}_{q,j} = \Delta_{q,j} + \chi_j$  is the renormalized qubit detuning. The effective coupling between qubits via their shared cavity is  $J_{jk} = g_j g_k (\delta_j + \delta_k)/2\delta_j \delta_k$ , a manifestation of cavity-mediated cross-coupling.

We denote the multiqubit state  $|\psi(t)\rangle = \sum_z c_z(t)|z\rangle$ , where  $|z\rangle = |z_j\rangle^{\otimes N_q}$  represents the  $z$ -basis state of each qubit as a binary digit ( $\hat{\sigma}_z|0/1\rangle = \mp|0/1\rangle$ ). We consider the standard measurement process here, where the cavity is initially in the vacuum state and a coherent drive is applied at  $t=0$ :  $\epsilon(t) = \epsilon_0 \Theta(t)$ . The  $X$  quadrature of the cavity follows a unique trajectory for each multiqubit state, and by measuring  $\langle \hat{d} + \hat{d}^\dagger \rangle(t)$  one seeks to determine  $|\psi(0)\rangle$ . In this work, we specialize to the case of 2-qubit readout ( $z = \{00, 01, 10, 11\}$ ), which is later seen to be a nontrivial classification task. We consider both the dispersive and JC models with following parameters in units of the cavity decay rate  $\kappa$ ,  $\Delta_c = 0$ ,  $\epsilon_0 = 2$ ,  $\chi_1 = 1.8$ ,  $\chi_2 = 1.3$ ,  $\Delta_{q,1} = 180$ ,  $\Delta_{q,2} = 130$ ,  $g_j/\delta_j = 10^{-1}$ , and include additional qubit decay with rate  $\gamma_h = 10^{-2}$ . These are all physically plausible parameters for current superconducting circuit implementations of this system.

A thorough discussion of the joint dispersive measurement process is contained in Appendix B; in the remainder of this section we summarize the salient results. Figure 1(a) depicts the expected readout cavity evolution  $\langle \hat{d} + \hat{d}^\dagger \rangle(t)$  for each initial qubit state  $|z\rangle$  in the measurement basis under the dispersive ( $\hat{\mathcal{H}}_D$ ) and JC models ( $\hat{\mathcal{H}}_{\text{JC}}$ ). In both cases, the cavity evolves to distinct steady states over a timescale set by  $\kappa$ . The difference between these models is manifest in the qubit evolution shown in Fig. 1(b). Here we plot the expected probability that the system will be measured to be in the multiqubit state it was prepared in:  $|c_z(t)|^2 = |\langle z|\psi(t)\rangle|^2$ , for  $|\psi(0)\rangle = |z\rangle$ . The decay of initially excited states can be seen to be significantly faster for the JC model. In the dispersive model the qubit state evolution is due to  $J_{12}$  and  $\gamma_h$ , and the corresponding timescales are taken to be slow relative to the system dynamics. As  $J_{12}/\kappa, \gamma_h/\kappa \rightarrow 0$ , the measurement process becomes QND since the qubit state is conserved, where in a perfect QND measurement  $|c_z(t)|^2 = 1$ . In contrast, the JC interaction  $\propto g_j$  does not commute with  $\hat{\sigma}_{z,j}$ , and this fast Hamiltonian evolution causes information about the initial qubit state to be lost more quickly during measurement.

We consider the situation where only the  $X$  quadrature of the output cavity field is continuously monitored via

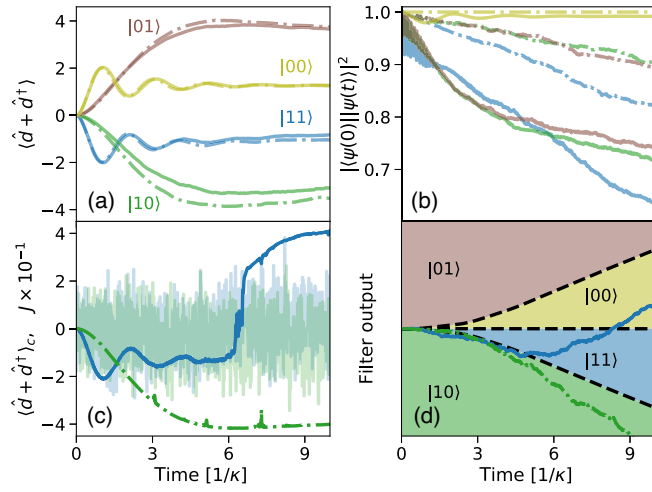


FIG. 1. (a) Ensemble-averaged cavity field quadrature during readout, for each initial qubit state  $|\psi(0)\rangle = |z\rangle$ . (b) Ensemble-averaged decay of the initial qubit state occupation during measurement:  $|\langle \psi(0) | \psi(t) \rangle|^2$ . In all plots, results for the dispersive and JC models are shown in dash-dotted and solid lines, respectively, and  $|\psi(0)\rangle$  for each curve is indicated via the colors labeled in (a). Panels (c) and (d) demonstrate the matched filtering of dispersive measurement currents. In (c) we show sample readout cavity quadratures for  $\hat{\mathcal{H}}_{JC}$  initialized in  $|11\rangle$  (blue) and  $\hat{\mathcal{H}}_D$  in  $|10\rangle$  (green). The corresponding measurement currents  $J$  are also plotted, reduced by an order of magnitude for visibility. The subsequent filtered readout signals  $y(t)$  are shown in (d), as computed via an ideal  $Q \rightarrow \infty$  matched filter. The filtered output is classified according to which expected bin it falls in, which are labeled with their associated quantum state. All figures are for the parameters in text.

homodyne detection, with efficiency  $\eta = 1$  unless specified. The QS ( $\hat{\mathcal{H}}_S \in \{\hat{\mathcal{H}}_D, \hat{\mathcal{H}}_{JC}\}$ ) then evolves under the stochastic master equation [4,43]:

$$\dot{\hat{\rho}} = -i[\hat{\mathcal{H}}_S, \hat{\rho}] + \gamma_h \sum_j \mathcal{D}[\hat{\sigma}_{-,j}] \hat{\rho} + \kappa \mathcal{D}[\hat{d}] \hat{\rho} + \sqrt{\kappa\eta} \mathcal{M}[\hat{d}] \hat{\rho} \xi(t). \quad (3)$$

In the above, the dissipative and measurement superoperators are, respectively,  $\mathcal{D}[\hat{O}] \hat{\rho} = \hat{O} \hat{\rho} \hat{O}^\dagger - \frac{1}{2} \{\hat{O}^\dagger \hat{O}, \hat{\rho}\}$ ,  $\mathcal{M}[\hat{O}] \hat{\rho} = \hat{O} \hat{\rho} + \hat{\rho} \hat{O}^\dagger - \langle \hat{O} + \hat{O}^\dagger \rangle_c \hat{\rho}$ . The SME describes the evolution of the QS  $\hat{\rho}$  conditioned on the observed measurement outcome  $J(t)$ . This measurement outcome is the continuous classical current,

$$J(t) = \sqrt{\kappa\eta} \langle \hat{d} + \hat{d}^\dagger \rangle_c(t) + \xi(t), \quad (4)$$

where  $\xi(t)$  is white noise, arising from fundamental quantum uncertainty in the cavity state:  $\langle \xi(t) \rangle = 0$ ,  $\langle \xi(t) \xi(t') \rangle = \delta(t - t')$ . In the above, the subscript  $c$  denotes expectation values taken with respect to the conditional state  $\hat{\rho}$ .

These quantities evolve stochastically during individual measurements: samples of  $J(t)$  and  $\langle \hat{d} + \hat{d}^\dagger \rangle_c(t)$  are depicted in Fig. 1(c). The measurement signals  $J(t)$  are dominated by noise  $\xi(t)$ , and the measurement process produces backaction on the quantum state, resulting in, for example, the sudden jump seen for the JC sample. By taking the ensemble average of many measurement records, however, one recovers the unconditional system dynamics of Fig. 1(a), described in Appendix B. Our quantum measurement data are constructed by numerically integrating the SME of Eq. (3) from initial states  $\hat{\rho}(0) = |0, z\rangle\langle 0, z|$  using QuTiP [44]. Each trajectory  $q$  has a unique noise record  $\xi^{(q)}(t)$  and thus conditional expectation values  $\langle \hat{O} \rangle_c^{(q)}$  and measurement signal  $J^{(q)}(t)$ .

The individual measurement currents  $J^{(q)}(t)$  have a small signal-to-noise ratio (SNR) due largely to the additive white noise term in Eq. (4), obscuring the relevant conditional evolution, particularly when the cavity photon number is kept low to keep the measurement in the QND regime. In particular, for the chosen parameters, including a sampling period of  $10^{-2}/\kappa$ , the steady-state measurement current SNRs are  $-6.8$  and  $-7.1$  dB for the dispersive and JC systems, respectively. Further signal processing is thus needed to extract the underlying initial qubit state; conventionally, this is done by constructing a matched filter (MF) from a large set of measurement currents for which the initial qubit state is known. Specifically, a MF is a linear filter with kernel  $h(\tau) = \langle J^*(t - \tau) \rangle_Q$ , the conjugated and time-reversed mean of a  $Q$  sample training set. For an input consisting of a signal plus white noise, the output approaches the autocorrelation function of the signal as  $Q \rightarrow \infty$ :  $y^{(q)}(t) = \int^t d\tau h(t - \tau) J^{(q)}(\tau) = \int^t d\tau \langle J^*(\tau) \rangle_Q J^{(q)}(\tau)$ . This maximizes the SNR; a MF is the optimal linear filter for distinguishing signals, such as we study here, with additive noise. For the 2-qubit readout system, the MF is constructed by summing the magnitude of the mean currents associated with each initial state:  $h(\tau) = \sum_z |\langle J^*(t - \tau) \rangle_{z,Q}|$ , resulting in a scalar kernel maximizing the overall fidelity with which currents associated with different states can be distinguished.

The filter is used to define an expected bin for each filtered signal as a function of time, and quantum states are classified according to which bin they fall into. The MF classification process is depicted in Fig. 1(d); the filter is seen to remove the white noise from the homodyne signals, and the sample from the dispersive system falls into the correct bin reasonably quickly. The readout process has a more significant influence on the qubit state in the JC system, which in this case causes the first qubit to decay at  $t \sim 6/\kappa$ . This results in the cavity state suddenly jumping as well, and the filtered output falls into the wrong bin at later times as a result. This loss of initial state information, at rates depicted in Fig. 1(b), places a limit on maximum

fidelity with which signals can be classified, since the task is to learn  $|\psi(0)\rangle$ , not  $|\psi(t)\rangle$ .

### III. SUPERCONDUCTING RESERVOIR COMPUTER

#### A. Proposal and overview

Having discussed the conventional treatment of quantum state measurement, we now describe our proposed RC approach, depicted schematically in Fig. 2. Generally, the RC's task is to classify the state  $|\psi(0)\rangle$  of the target QS, based on a quantum measurement that results in a noisy readout signal containing information about this initial state. We specialize to the QS described in the previous section and the classification of its homodyne measurement current Eq. (4) in this paper, but the same approach can be applied to entirely different systems and measurement modalities.

Instead of conventional processing with room-temperature electronics and a software backend [14,16,17], the readout current is routed to a superconducting reservoir. We assume a generic architecture for this device, modeled by a coupled network of Kerr oscillators coupled to this signal via a linear input matrix  $\mathbf{W}_I$ . The reservoir nodes then evolve according to this input and an internal structure defined by a connectivity matrix  $\mathbf{W}_R$  and nonlinearity vector  $\Lambda$ . This results in a complex dynamical mapping: the physical state of the RC  $\mathbf{x}(t)$  is a high-dimensional nonlinear functional of the input history  $\mathbf{u}(\tau \leq t)$ . The output of the computation is a

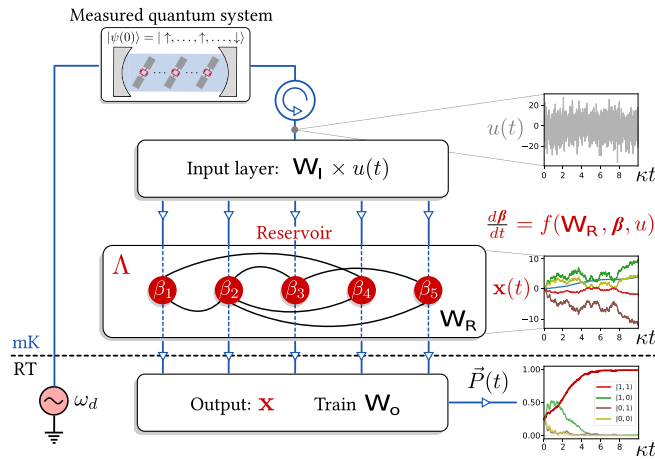


FIG. 2. A schematic of our proposal to use a hardware RC to process quantum measurement signals, here shown for  $K = 5$  Kerr nodes. The QS is interrogated, and the resulting measurement current  $u(t)$  is input to the reservoir. The input layer  $\mathbf{W}_I$  randomly couples this signal to each node of the Kerr network, whose subsequent dynamics are a nonlinear function of  $u(t)$  and the network itself via Eq. (8). The output layer performs classification by learning linear combinations of measured reservoir variables  $\mathbf{x}$  which extract the probability that the input history corresponds to a given underlying quantum state.

linear combination of the RC nodes  $\mathbf{y}(t) = \mathbf{W}_o \mathbf{x}(t)$ , where the matrix  $\mathbf{W}_o$  can be trained such that  $\mathbf{y}(t)$  approximates a desired function or functional of the input  $\mathbf{F}(\mathbf{u}(\tau \leq t))$ . This linear output layer is the only RC property which is varied, and in practice can be implemented at room temperature in an field-programmable gate array after simply extracting all available reservoir degrees of freedom.

Because only  $\mathbf{W}_o$  is optimized, training is a simple convex optimization problem with a limited number of parameters and is guaranteed to converge [19,26]. Physical RC thus avoids the “vanishing gradient problem” that plagues training of other neural networks [45], as well as the “simulation-reality gap” [46]. Although it may seem like training only output weights  $\mathbf{W}_o$  would lead to inferior results for time-series processing, head-to-head comparisons between state-of-the-art recurrent neural networks and RCs show surprisingly similar performance [47–49], despite RC training protocols being  $10^3$ – $10^6$  times quicker. The internal structure ( $\mathbf{W}_I$ ,  $\mathbf{W}_R$ ,  $\Lambda$ ) is not optimized; RCs can thus be quickly retrained for different tasks, rendering them powerful and generalizable analog neural networks [21,24,30,38].

The reservoir we propose may be realized in a superconducting circuit platform via a network of coupled Josephson parametric oscillators (JPOs). These JPOs can be either single Josephson junctions or composite elements such as superconducting nonlinear asymmetric inductive elements (SNAILs) [50]. The recurrent connections can be flexibly generated by coupling the JPOs to a common electromagnetic resonator. Variants of such networks have been considered as hardware for superconducting quantum annealers [51–53] and for stabilization of multiqubit entanglement [54,55]. Such a RC then could be integrated with the QS to be measured, sharing a cryogenic environment. For the present work we require that the JPOs are in the weakly nonlinear regime, that their scale of nonlinearity is much smaller than their dissipation. This is the regime of Josephson parametric amplifiers [56]. Another particularly interesting platform to realize this hardware RC is an optical Kerr network, which would then be well suited to the readout of optical QSs. An important advantage of either of these realizations of Fig. 2 is that the RCs will operate at the timescales of the measured quantum system, faster than conventional field-programmable gate array-based electronics and potentially allowing for real-time analog processing.

#### B. Kerr network reservoir computer model

We model the reservoir as a generic network of coupled Kerr-nonlinear oscillators, described by the master equation,

$$\dot{\hat{\rho}}_{RC} = -i[\hat{\mathcal{H}}_{RC}, \hat{\rho}_{RC}] + \sum_k \gamma_k \mathcal{D}[\hat{b}_k] \hat{\rho}_{RC}, \quad (5)$$

where the governing Hamiltonian  $\hat{\mathcal{H}}_{RC}$  takes the form

$$\begin{aligned} \hat{\mathcal{H}}_{\text{RC}} = & \sum_k \Delta_k \hat{b}_k^\dagger \hat{b}_k - \frac{\lambda_k}{2} \hat{b}_k^{\dagger 2} \hat{b}_k^2 + \sum_{kl} g_{kl} \hat{b}_k^\dagger \hat{b}_l \\ & + \sum_{km} i \varepsilon_{km} (u_m(t) e^{-i\Delta_{d,m} t} \hat{b}_k^\dagger - u_m^\dagger(t) e^{i\Delta_{d,m} t} \hat{b}_k). \end{aligned} \quad (6)$$

The input to the RC is a collection of signals  $\mathbf{u}(t)$  where each  $u_m(t)$  has a unique carrier frequency  $\omega_{d,m}$ . Each nonlinear oscillator is described by a field operator  $\hat{b}_k$ , with detuning  $\Delta_k$  with respect to a common reference  $\omega_{d,1}$ , and Kerr nonlinearity  $\lambda_k$ .  $g_{kl}$  and  $\varepsilon_{km}$  are the linear couplings between oscillators and to the input, respectively, and  $\gamma_k$  is the energy decay rate. Interoscillator couplings can be generated by cavity-mediated interactions [54,55] or through parametric means [53].

The evolution of the field amplitude from each node is given by the Heisenberg equation of motion:

$$\begin{aligned} \dot{\langle \hat{b}_k \rangle} = & - \left( i\Delta_k + \frac{\gamma_k}{2} \right) \langle \hat{b}_k \rangle + i\lambda_k \langle \hat{b}_k^\dagger \hat{b}_k^2 \rangle \\ & - i \sum_l g_{kl} \langle \hat{b}_l \rangle + \sum_m \varepsilon_{km} u_m(t) e^{-i\Delta_{d,m} t}. \end{aligned} \quad (7)$$

This general form of reservoir evolution is suitable for frequency-multiplexed readout of several qubits [6,7,17]. Henceforth we simplify to the joint dispersive readout scenario introduced in Sec. II, with a single carrier frequency  $\omega_d$  and scalar input  $u(t)$ .

We will consider Kerr networks where the input signals and corresponding field amplitudes are sufficiently large that they are in the classical regime. For  $c \in \mathbb{R}^+$ , defining scaled drive strengths  $\tilde{\varepsilon}_{km} = \sqrt{c} \varepsilon_{km}$ , nonlinearity  $\tilde{\lambda}_k = \lambda_k/c$ , and introducing  $\beta_k \equiv \sqrt{c} \langle \hat{b}_k \rangle$ , it can be shown that  $\langle \hat{b}_k^\dagger \hat{b}_k^2 \rangle = |\beta_k|^2 \beta_k + O(1/c)$ . Heuristically, this indicates that for  $c \rightarrow \infty$ , where the nonlinearity grows weaker and the ‘‘classical’’ occupation  $|\beta_k|^2 = c |\langle \hat{b}_k \rangle|^2$  becomes simultaneously larger, quantum correlations captured in higher-order moments can be neglected. Note that now  $\mathbf{u}$  are classical currents as well, whose potential quantum-limited amplification is absorbed in  $c$  [56,57].

In this case, Eq. (7) becomes

$$2\dot{\beta}_k/\gamma = -\beta_k + i\Lambda_k |\beta_k|^2 \beta_j - i(\mathbf{W}_R \cdot \boldsymbol{\beta})_k + (\mathbf{W}_I \cdot \mathbf{u})_k. \quad (8)$$

Without loss of generality, we have chosen the nodal decay rates to be identical,  $\gamma_k = \gamma$ , to define the dimensionless parameters, familiar to the RC framework:

$$\begin{aligned} W_{R,kk} &= 2\Delta_k/\gamma, & W_{R,kl} &= 2g_{kl}/\gamma, \\ W_{I,kl} &= 2\tilde{\varepsilon}_{kl}/\gamma, & \Lambda_k &= 2\tilde{\lambda}_k/\gamma. \end{aligned} \quad (9)$$

Equations (8) govern the  $K$ -node RC response to a given input signal  $\mathbf{u}(t)$ . This description applies both to the high power limit of  $\mathbf{u}(t)$  from measured QS on chip (Fig. 2) and

to the injection of any classical signal directly to the reservoir for processing.

The philosophical underpinning of reservoir computing holds that if one has a sufficiently complex and high-dimensional system, there is no need to optimize its many internal parameters for a specific computational task. As such, here we consider random Kerr networks, whose internal structure and dynamics, specified by  $\mathbf{W}_I$ ,  $\mathbf{W}_R$ ,  $\boldsymbol{\Lambda}$ , and  $\boldsymbol{\gamma}$  via Eqs. (8), are set randomly and *not individually optimized*. Instead, the RC properties are controlled by the scale-independent hyperparameters  $\{\gamma, \alpha, \bar{\Lambda}, \mu\}$ , which define the limits  $[a, b]$  of uniform distributions which reservoir internal parameters are randomly sampled from:

$$\begin{aligned} W_{I,kl} &\in [-\mu, \mu], & \Lambda_k &\in [0, 2\bar{\Lambda}], \\ W_{R,kl} &\propto [-1, 1], & \text{such that } \alpha &= \lambda_{\max}(\mathbf{W}_R), \end{aligned} \quad (10)$$

where  $\lambda_{\max}(\mathbf{W}_R)$  refers to the maximum singular value of the connectivity matrix  $\mathbf{W}_R$ . We constrain the ranges of these hyperparameters to be compatible with the proposed hardware realizations, while also importantly enabling desired RC properties of fading memory, separability and nonlinearity; their selection is discussed in more detail in Sec. IV D. We will see throughout this work that the generic behavior of a RC is well quantified by its hyperparameters, and that performance is robust to both network structure and variations in these values. We have also introduced significant variation into node decay rates  $\gamma_k$  and observed that the RC performance is again unchanged.

We describe the state of a  $K$ -node Kerr RC in terms of the  $2K$ -dimensional vector of complex node amplitudes:  $\mathbf{x}(t) = \sqrt{2}(\text{Re}\{\beta_1\}, \text{Im}\{\beta_1\}, \dots, \text{Re}\{\beta_K\}, \text{Im}\{\beta_K\})^T$ . The learned RC output can always be expressed as a linear combination of these amplitudes  $\mathbf{y}(t) = \mathbf{W}_o \mathbf{x}(t)$ , where  $\mathbf{W}_o$  is a  $(C \times 2K)$ -dimensional matrix of output weights. If each Kerr oscillator is subject to a heterodyne measurement as is standard in the superconducting realization we envision, one extracts both orthogonal  $(\text{Re}\{\beta_k\}, \text{Im}\{\beta_k\})$  quadratures, to which the weights  $\mathbf{W}_o$  can be applied externally [58].

To facilitate our analysis of reservoir processing, and accommodate realizations where only a single quadrature from each oscillator is measured, we specify that a single quadrature for each node  $x_k^{\phi_k}$ , at angle  $\phi_k$ , contributes to the output:

$$\mathbf{y}(t) = \mathbf{W}_o \mathbf{x}(t) = \tilde{\mathbf{W}}_o \mathbf{C}(\boldsymbol{\phi}) \mathbf{x}(t) \equiv \tilde{\mathbf{W}}_o \mathbf{x}^\phi(t). \quad (11)$$

Here  $\tilde{\mathbf{W}}_o$  is a  $(C \times K)$ -dimensional matrix of output weights over nodes, while  $\mathbf{C}(\boldsymbol{\phi})$  is a  $K \times 2K$  projection matrix, parametrized by  $K$  angles  $\boldsymbol{\phi} = (\phi_1 \dots \phi_K)$ , with its only nonzero entries given by

$$C_{k,2k-1}(\boldsymbol{\phi}) = \cos \phi_k, \quad C_{k,2k}(\boldsymbol{\phi}) = \sin \phi_k. \quad (12)$$

This choice reduces the linear output layer from  $2CK$  to  $(C+1)K$  parameters to optimize in RC training. More significantly,  $\mathbf{x}^\phi$  becomes the RC variables which contain the most information for the learned task, through which we develop an intuitive picture of reservoir information processing in Sec. IV C. Further, we find that this effective reduction in the number of output variables does not reduce performance, particularly with the small training sets considered in this work. However, we emphasize that this projection—simply a constraint on the weight matrix  $\tilde{W}_o$  implemented in software—is not necessary, and using both quadratures per node in general tasks can only improve performance.

Finally, specific to the task of classifying  $C$  states of the QS, it is appealing to map the RC output to a probability that the input to the RC corresponds to the QS initialized in  $|\psi(0)\rangle = |z\rangle$ . This is achieved by applying the “softmax” function (Boltzmann distribution) to the trained RC output:

$$P_j(t) = e^{y_j(t)} / \sum_k e^{y_k(t)}, \quad (13)$$

which expresses  $y_j$  as mutually exclusive probabilities that sum to unity.

Details of and background on RC training are discussed in Appendix C; we briefly summarize some of the salient aspects here. We first construct a training dataset consisting of  $Q$  measurement currents  $U(t) = \{J^{(q)}(t)\}$  for a measurement period  $\tau_m$  for associated known initial states  $\mathbf{y}^{*(q)} = z^{(q)}$  ( $q = 1, \dots, Q$ ).  $M$  measurements are conducted for each initial state by integrating the SME of Eqs. (3) and (4). The RC node trajectories are simulated by solving Eqs. (8), and the resultant states are used to minimize the multinomial cross entropy cost function  $\mathcal{L}_x$  (Appendix C), optimizing  $\tilde{W}_o$  and  $\boldsymbol{\phi}$  for the training set. We consider small reservoirs ( $K = 2\text{--}5$ ) and training sets ( $Q < 100$ ). Therefore, the optimization of the cross entropy loss function on a digital computer is a quick convex optimization problem with a small set of output parameters.

It is worth reiterating here that reservoir internal parameters are not optimized, and this output layer can be implemented via a simple linear combination of the reservoir node quadratures *externally*: there is no need to tune any aspect of the hardware during training or operation. We further stress that this network size is orders of magnitude smaller than typical reservoirs, a choice made for hardware realizability. After training, the direct RC output  $\max_z \{y_z^{(q)}(t)\}$  from Eq. (11) indicates the initial state  $z^{(q)}$  most likely to correspond to the observed record  $J^{(q)}(t)$  up to the current time, which we express for

interpretive convenience as a continuous probability  $P_z^{(q)}(t)$  via Eq. (13).

#### IV. RESERVOIR PROCESSING OF QUANTUM MEASUREMENT

We now analyze the performance of the proposed Kerr network RC on the 2-qubit readout task described above. We focus on the ability of a RC trained with a small labeled training set to classify a much larger set of unknown test signals (quantum states) from either the dispersive or JC system:  $\hat{\mathcal{H}}_S \in \{\hat{\mathcal{H}}_D, \hat{\mathcal{H}}_{JC}\}$  in Eq. (3), respectively. A transparent metric to evaluate the performance is the “classification accuracy”  $C_{\mathcal{F}}(t)$ , which refers to the fraction of test signals the RC correctly classifies at a given readout time. For short readout times, the cavity is still being populated and the signal is dominated by shot noise. However, classification accuracy does not increase monotonically with readout time; as indicated in Fig. 1, the initial qubit state stochastically evolves and can be lost during the measurement process, an outcome that is particularly likely for the JC system. After this point  $J(t)$  no longer faithfully provides information about  $|\psi(0)\rangle$ , imposing a steadily decreasing ceiling on  $C_{\mathcal{F}}(t)$  (as can be seen in Fig. 4).

In practice, one would simply stop the measurement process and RC computation at the time  $C_{\mathcal{F}}(t)$  peaks; this time  $\tau_m$  is consistent for a given QS being measured. Thus, we define the “classification fidelity”  $\mathcal{F}$  of a RC or filter to be this peak accuracy,

$$\mathcal{F} = \max_t \{C_{\mathcal{F}}(t)\}, \quad (14)$$

and we seek to optimize this metric when considering RC design in Sec. IV D. Since the maximum possible  $C_{\mathcal{F}}(t)$  decays with time,  $\mathcal{F}$  combines both speed and accuracy of classification.

We compare the RC performance against that of conventional filtering. In particular, we consider a boxcar filter (BF), which amounts to integrating  $J(t)$  to remove noise, a matched filter constructed using a training set of size  $Q$  of labeled homodyne currents, and a MF constructed using the analytic solution  $\langle \hat{a} + \hat{a}^\dagger \rangle_z$  to  $\hat{\mathcal{H}}_D$ . For the dispersive system, this corresponds to the  $Q \rightarrow \infty$  limit of the MF constructed from homodyne currents. In many realistic scenarios, the system model and parameters are not known exactly or are subject to experimental drift, and attempting to construct an analytic MF is impractical. Instead, either the BF or a finite  $Q$  MF is used because they (like the RC) do not require a model of the underlying system; for the MF this requires regularly producing large training sets every time the device is recalibrated.

##### A. Reservoir dynamics and classification fidelity

The response of a representative  $K = 5$  node RC to a quantum readout signal from the dispersive system is




FIG. 3. A demonstration of the RC classification process. Panel (a) depicts a sample quantum signal  $u(t) = J^{(q)}(t)$  generated by the joint dispersive readout under  $\hat{H}_D$  of the initial state  $|\psi(0)\rangle = |11\rangle$ . This drives a  $K = 5$  node RC, producing the node state dynamics  $x_k^{\phi}(t)$  shown in (b). The classification result is in (c): the RC, previously trained with  $Q = 40$  readout signals, continuously updates the probability that the initial state was a specific state in the measurement basis. These probabilities are shown in solid lines in the color-scheme of Fig. 1. The decomposition of the true quantum state conditional on the measurement record up to time  $t$  is depicted in lighter dashed lines; for this sample the RC output is correct after  $\sim 3/\kappa$ . Note that  $|c_{01/00}|^2(t)$  (and associated  $P$ ) are vanishingly small.

depicted in Fig. 3. Here (and in Fig. 4) the RC has hyperparameters  $\gamma = 0.7\kappa$ ,  $\alpha = 1.9$ ,  $\bar{\Lambda} = 5 \times 10^{-2}$ , and  $\mu = 5$ . The RC is driven with measurement signals  $u(t) = J^{(q)}(t)$  with a total measurement time  $10/\kappa$  and has previously been trained with  $Q = 40$  readout signals ( $M = 10$  samples for each state) to optimize  $\{\bar{W}_o, \phi\}$ . In Fig. 3(a) we show one such noisy signal for an unknown (to the RC) quantum state  $|\psi(0)\rangle = |11\rangle$ . Once the readout drive is turned on, the noisy input  $u(t)$  acquires a nonzero mean. This excites the network away from its rest state  $x = 0$  to a nontrivial trajectory  $x^{\phi}$  in its measured phase space, as shown in Fig. 3(b). Real-time classification is performed by reading out each node continuously: the outputs  $P_z^{(q)}(t)$  are shown in Fig. 3(c), where the RC quickly stabilizes with  $P_{11}^{(q)} > 0.5$ , thus correctly classifying this measurement record. The low measurement SNR while the cavity is being populated is responsible for the RC initially not distinguishing  $|10\rangle$  and  $|11\rangle$ , with  $P_{10}^{(q)} \sim P_{11}^{(q)}$ . The true quantum state  $|c_z^{(q)}(t)|^2 = \text{tr}\{\hat{\rho}(t)|z\rangle\langle z|\}$  is also depicted, and we see that the RC classification  $P_{11}^{(q)}(t)$  is robust to the significant probability amplitude fluctuations: the system almost jumps to the state  $|01\rangle$  for this specific trajectory, but  $P_{11}^{(q)}$  remains saturated.

The ability of this RC to classify unknown quantum states is quantified in Fig. 4, where we evaluate its performance on a test set of 1200 unknown quantum signals generated from the dispersive and JC SMEs, and compare with that of various linear filters. The RC is trained as above using a  $Q = 40$  measurement set for the corresponding QS, while the MF is constructed from a much larger  $Q = 1200$  set. We plot the classification accuracy on




FIG. 4. Classification accuracy versus readout time on a 1200 signal test set for the dispersive system in (a) and the JC system in (b). In both cases, the RC of Fig. 3 (solid blue line) is trained with  $Q = 40$  measurements and the optimal readout time is indicated. The average classification fidelity of 10 different random RCs sharing the same hyperparameters is also shown in light dash-dotted blue line. For comparison purposes we also show the classification accuracy of a BF (dashed green line), analytic MF (light red line), and  $Q = 1200$  MF (dash-dotted red line).

both datasets for the RC and the filters; it is apparent that the RC is able to rapidly and reliably extract the initial quantum state and thus perform the readout for both Qs.

For the dispersive system both the RC and the MFs have a classification fidelity  $\mathcal{F} > 0.96$ , achieved for an optimal readout time  $\tau_m \sim 6.7/\kappa$ . Interestingly, one can compute the mutual information between the generated reservoir state  $x^{\phi(q)}(t)$  and initial qubit states  $z^{(q)}$ , and find that it follows  $C_{\mathcal{F}}(t)$ , saturating at  $\tau_m$ .  $\mathcal{F} \sim 0.96$  corresponds to a mutual information of 1.9 bits (both limited by non-QND effects), which implies that we can expect the RC's state at  $\tau_m$  to yield, on average, 1.9 bits of information about the 2-bit initial state. A dip in the MF performance is seen around  $2.5/\kappa$ ; the expected filtered signals for  $|1(0)0\rangle$  and  $|1(0)1\rangle$  cross for this specific measurement time, meaning the MF has no information about the second qubit. This occurs in the BF as well at a slightly later time. The RC avoids such a problem by mapping the qubit state into its higher-dimensional phase space.

As described in Sec. II, initial state information is lost more quickly for the JC system due to non-QND Hamiltonian evolution. Despite this, the RC is able to accurately process these quantum measurements, with a classification fidelity of 0.92, exceeding that of any linear filter. The optimal measurement time is slightly later, at  $\sim 7/\kappa$ . This is due to the decreased SNR for the JC readout task:  $\sqrt{\kappa}\langle\hat{d} + \hat{d}^\dagger\rangle$  is reduced relative to the dispersive case and there is increased measurement backaction noise [see Figs. 1(a) and 1(c)]. We can decrease the SNR further by decreasing the measurement efficiency  $\eta \leq 1$ : for both systems the fidelity of the RC classification decreases steadily with increasing relative noise strength, but remains comparable with that of the MF.

This performance is in no way unique to the specific random realization of the RC network; the average



classification accuracy of 10 random  $K = 5$  RCs is also shown for both Qs in Fig. 4. These RCs have different structure ( $\mathbf{W}_I, \mathbf{W}_R, \mathbf{\Lambda}$ ) but the same hyperparameters and are trained on the same  $Q = 40$  training set. The peak classification accuracy occurs at different times for different RCs, resulting in the average curve being artificially flattened. The average classification fidelity of the 10 different RCs is 0.951 for the dispersive system and 0.905 for the JC system (for  $Q = 100$ , this increases to 0.915), indicating the robustness of the RC approach. A few of the random networks experience a sharper drop in classification accuracy after the optimal readout time: this is because the RC states corresponding to different input signals are less separated in phase space. In the next section we explore the role of the RC evolution in its phase space in more detail.

### B. Rapid training

The results of Fig. 4 suggest that a RC is capable of matching the performance of a MF constructed with a much larger training set. We find that RCs hold this training-cost advantage over a MF generally, and demonstrate its dramatic nature in Fig. 5. Here, we plot classification fidelity as a function of the number of training signals  $Q$  used to train 10 random  $K = 5$  RCs from Sec. IV A. Each specific network is indicated in a different color, with the




FIG. 5. Test set classification fidelity as a function of number of training trajectories for the MF and  $K = 5$  RCs described in the text. The dispersive and JC readout systems are in (a) and (b), respectively. Individual RCs are indicated with different colored circles, and the RC of Fig. 3 is emphasized with a black plus. The average RC fidelity is indicated with a dashed blue line, and can be compared with the MF performance, denoted with red crosses and a dashed line. The ( $Q$ -independent) performance of the BF and analytic MFs are denoted with horizontal dashed green and pale red lines, respectively.

network emphasized in Figs. 3 and 4 highlighted with a plus. The average classification fidelity across the 10 networks is also shown. The RCs perform remarkably well at low  $Q$ , with 4 readout signals (just one for each qubit state) being sufficient to achieve an average fidelity of 0.90. This average fidelity increases and the variance in performance across networks decreases as the training set grows, up to  $\sim 40$  for the dispersive system ( $Q \sim 80$  for the JC). Beyond this point, the small fluctuations in fidelity are due to the finite *test* set (1200 signals) used to calculate  $\mathcal{F}$ . The RC thus achieves optimal performance with a small  $Q = 40$ – $80$  training set for both the dispersive and JC quantum systems, demonstrating its efficacy for rapid measurement calibration. In contrast, the MF performs very poorly in this regime, needing  $O(10^3)$  training signals (note the horizontal axis scale) to converge to the average RC network classification fidelity.

This significant advantage is particularly relevant for readout of online quantum processors for which calibrations need to be done regularly: a RC-based measurement scheme would thus require far fewer initialization and measurement runs to train than are needed to construct a MF of comparable performance. This can enable a resource-efficient readout calibration system that can also be robustly automatized. The RC network can easily and quickly be retrained as conditions or even the target QS change, facilitated by the computationally inexpensive software component of training, particularly for small training sets.

We note that while widely used high-level metrics such as squared error or classification fidelity are ideal for quantifying the performance of the RC, they do not elucidate the fundamental source of its classification efficacy and fast learning ability. To address this, we carry out a detailed analysis of the phase space dynamics of the Kerr RC nodes in the following section, directly connecting the underlying reservoir to the observed performance.

### C. Phase space dynamics

Generally, the effectiveness of the RC approach is attributed to the expressive power of its high-dimensional state space [23,37,49]. In the present discussion, the RC under scrutiny transforms a scalar input signal into the  $2K$ -dimensional state space of the RC  $\{\beta_k, \beta_k^*\}$ . It is also well understood that the nonlinearity of the RC plays a crucial role [22,24,30], but the underlying mechanism behind how these and other RC properties impact the measurement task is not immediately clear. To perform a fundamental analysis of Kerr RC dynamics and gain unique insight into its previously described performance, we introduce the measured section (MS) of the reservoir. Generally, it is difficult to visualize the dynamics in the high-dimensional RC state space. Recalling, however, that we only project out a specific quadrature of the RC oscillators  $x^\phi$  [Eq. (11)], these “visible” RC nodes evolve in




FIG. 6. (a) Classification dynamics in the measured subspace for a  $K = 2$  RC with parameters  $\gamma = 0.2$ ,  $\alpha = 1.9$ ,  $\bar{\Lambda} = 5 \times 10^{-2}$ , and  $\mu = 5$ . In (b) we show the same network but with  $\bar{\Lambda} = 0$ . Regions in RC phase space are colored according to their learned classification output; e.g., the region in  $(x_1^{\phi_1}, x_2^{\phi_2})$  space for which  $P_{11}$  is largest is filled in blue. The separating lines constructed from the trained  $\tilde{W}_o$  are also shown. Dots indicate the final RC state generated by each signal in the test set, which are colored according to their ground truth  $|\psi(0)\rangle = |z\rangle$ . Therefore, when the colors of the dots and the underlying region are the same, the classification result is correct.

the MS, a  $K$ -dimensional subspace of the full  $2K$ -dimensional phase space. As the angles  $\phi_k$  have been optimized, this phase space projection contains the relevant information for the computational task.

With the aid of RC dynamics in the MS, we will show that  $K = 2$  degrees of freedom are sufficient to perform a four-state classification task, *provided* the RC is sufficiently nonlinear. For this two-node reservoir, the classification process can be conveniently visualized in the MS, as presented in Fig. 6. We consider the response of two RCs, both trained with  $Q = 40$  measurements, plotting the final reservoir state (at  $\tau_m = 6.7/\kappa$ ) in the MS for each signal in a test set of 1200 signals, color coded with its true  $|\psi(0)\rangle$ . The RC on the left is nonlinear with  $\bar{\Lambda} = 0.05$  and attains a  $\mathcal{F} = 0.96$  on the shown test set, while the RC on the right is linear ( $\bar{\Lambda} = 0$ ), with only  $\mathcal{F} = 0.75$ .

The training task of learning  $\{\tilde{W}_o, \phi\}$  is equivalent to finding hyperplanes in the full RC phase space which separate the regions in Fig. 6 in a manner that maximally distinguishes the RC states corresponding to different inputs. These hyperplanes for  $K = 2$  are simply lines in the MS, shown in Fig. 6 separating colored regions, and defined by  $(\tilde{W}_o \mathbf{x}^\phi)_j = (\tilde{W}_o \mathbf{x}^\phi)_k$ . Note that due to the symmetry of the dispersive readout classification problem, whereby  $\langle \hat{d} + \hat{d}^\dagger \rangle_{11/10} \simeq -\langle \hat{d} + \hat{d}^\dagger \rangle_{00/01}$ , two pairs of these four separating lines fall almost on top of each other.

We note that the dynamics of the Kerr RC demonstrate the four properties required of a reservoir, such that it forms an effective RC [19,22,23,29]: separation, approximation, fading memory, and nonlinearity. *Separation* is the requirement that different input classes map to distinct regions of phase space, while *approximation* ensures that input series

which are close generate RC states which are similarly close. Both these properties are manifest in Fig. 6, where the final reservoir states for each class fall into separated regions in the MS, and the final output is robust to noise in individual trajectories. These regions are statistical steady states: the ensemble average of the RC states for each input are in the vicinity of their fixed points, with limited diffusion resulting from individual stochastic trajectories. The *fading memory* property requires that the current RC state depends on the history of the input signal, with an increasing importance placed on more recent inputs. Here, the final RC state depends more strongly on the steady-state readout signal, where the SNR is higher and QS states can be more easily distinguished, aiding classification. At the same time, the RC state does not depend *entirely* on its most recent input, making the classification result resilient to sudden changes in the input signal, for instance those due to qubit decay.

The final requirement of a RC is of some degree of *nonlinearity* in its dynamics or output layer for nontrivial computation. For the linear RC, note that the classes approximately lie on a line in the MS. The linearity of the  $\bar{\Lambda} = 0$  RC enables its dynamics to be solved analytically:  $\beta_j(t) = \sum_k c_{jk} \int^t d\tau e^{(i\delta_k + \gamma/2)(\tau-t)} u(\tau)$ , where  $\delta_k$  are the real parts of the eigenvalues of  $\mathbf{W}_R$ , and  $c_{jk}$  is the product of the projections of eigenvector  $k$  onto node  $j$  and the input  $\mathbf{W}_I$ . In the steady state, the different nodes  $x_j$  are now effectively scaled and rotated copies of each other. The information capacity of classifying the final RC outputs is then no different from that of classifying the input signals themselves with a linear classifier.

By comparison, the role played by the Kerr nonlinearity to provide high fidelities for this task is evident from Fig. 6. The RC's nonlinearity “shears” the high-amplitude readout signals (associated with states  $|01\rangle$  and  $|10\rangle$ ) out of the line connecting the low-amplitude signals (states  $|00\rangle$  and  $|11\rangle$ ). The measured RC quadratures are then linearly independent and not trivially related to each other or the input signal, in contrast to the linear RC. The nonlinearity of the Kerr reservoir allows it to utilize its dimensionality, forming up to  $K$  linearly independent outputs, rather than being bound by the dimension of the input signal. Nonlinearity, together with the other dynamical RC properties satisfied by the Kerr RC, thus allow the four different classes of input signals and their resultant RC state distributions to be linearly separated via the output layer.

Finally, we discuss why the RC appears to require fewer training signals than an experimentally constructed MF to attain high-fidelity classification. In Fig. 6 we also plot the response to a *single* input trajectory for each class of input signal in black. Even though these are unlabeled, it is clear what their ground truth initial quantum state was; an effective RC is able to integrate out the noise in the readout signals, enabling efficient separation of the input signals into the corresponding (color-coded) phase space

distributions. These trajectories are information dense; each time point functions as a new point of training data. Hyperplanes drawn to separate these single trajectories are found to differ only slightly from those shown on the plot, constructed using  $Q = 40$  trajectories: training from a small set of measurement signals can be as effective as a much larger training set.

This is in stark contrast to the MF approach, which uses the noisy training data to construct a time-dependent kernel, as opposed to static hyperplanes. The MF needs to be able to separate signals at each point in time, and so a large number of training sets are needed to produce an estimate of the mean input at each time which is not dominated by noise. It is lastly reasonable to ask why these linear filters, which perform only a linear operation on the input signal, are able to perform the classification task with high fidelity, but a linear RC is not. The answer lies in the classification step: the continuous filtered signal  $y^{(q)}(t)$  is mapped to a discrete class label  $z$  by comparing which expected signal  $y^{(q)}(t)$  is closest to, via a “distance” calculation that introduces the necessary non-linearity. The RC approach is very different: the nonlinearity occurs in the dynamical signal processing, and the output classification step  $\mathbf{y} = \tilde{\mathbf{W}}_o \mathbf{x}^\phi$  is linear. The use of a nonlinear classification step would then supply the required non-linearity and enable the linear RC to perform comparably to a nonlinear RC.

#### D. Optimization of the Kerr reservoir

In the previous sections, we investigated the ability of specific Kerr RC networks to perform a quantum measurement task. We also demonstrated that this performance is not particularly dependent on that network structure by completing the same task with a set of random networks that share the same hyperparameters. We now explore the role of these hyperparameters in determining RC performance, presenting the dispersive system readout fidelity of random RC networks as a function of  $\gamma$ ,  $\bar{\Lambda}$ , and  $\mu$  in Fig. 7.

In Fig. 7(a), we plot  $\mathcal{F}$  for 10 random RCs with  $K = 2, 5, 10$  (blue, green, and brown, respectively), sharing the same hyperparameters. We vary *only*  $\gamma$ , which sets the rate at which the RC nodes evolve and thus the timescale over which it samples the input signal. This system response time is typically only considered for hardware RCs; in software approaches a RC conventionally evolves under an update that is equivalent to  $\gamma = 1/\Delta t$  [19]. From Fig. 7, we see that it is important for  $\gamma$  to be approximately matched to the timescale of the input signal’s evolution [29]; for the quantum readout task, the signal  $\sqrt{\kappa}\langle\hat{d} + \hat{d}^\dagger\rangle_c$  evolves at rate  $\kappa$ . To understand this relationship, consider first a slow RC, with  $\gamma \leq 0.1\kappa$ : the RC then responds to the input signal averaged over a large window,  $\beta_j \propto \int^t d\tau e^{\gamma/2(\tau-t)} u(\tau)$ , and it is consequently more difficult to distinguish between different signals on the timescale over which the measurement is done. This slow evolution furthermore results in




FIG. 7. Performance as a function of various hyperparameters. In (a) we vary the RC timescale  $\gamma$  for RCs with  $K = 2, 5$ , and  $10$  nodes, with  $\mu = 5$  and  $\bar{\Lambda} = 5 \times 10^{-2}$ . Panel (b) shows fidelity as a function of input coupling strength  $\mu$  and nonlinearity  $\bar{\Lambda}$  for  $K = 5$  networks with  $\gamma = 0.5\kappa$ . In all plots  $\alpha = 1.9$  and the RCs are trained with  $Q = 80$  measurements. Ten different random networks are shown in (b), and 10 for each  $K$  in (a), with individual  $\mathcal{F}$  in colored circles, and the average fidelity as a dashed line. The fidelity of the boxcar and analytic MF are indicated in green and red dashed lines as in previous plots.

very little displacement from the initial RC state  $\beta = 0$  over this  $10/\kappa$  measurement window. Since the nonlinearity is  $\propto \beta_j^3$ , the low node amplitude also results in the RC being effectively linear.

There is a second timescale in the input signal: that of the noise  $\xi(t)$ , at the sampling rate of the quantum measurement  $\Delta t$ . This defines an upper limit on  $\gamma$ ; it is advantageous for the RC to respond to the much slower underlying quantum signal rather than the rapid white noise, and so one should have  $\gamma \ll 1/\Delta t$ . This allows the RC to average over some of the white noise, improving its performance. As seen in Fig. 7,  $\gamma \sim 0.2\text{--}0.8\kappa$  is the roughly optimal range for the dispersive readout task, allowing the RC to still respond to signal dynamics while integrating out much of the readout noise.

Quite generally, with increasing  $K$  the performance of the Kerr RCs becomes more robust to variation in hyperparameter values and to randomness in reservoir structure ( $\mathbf{W}_I, \mathbf{W}_R, \bar{\Lambda}$ ). In the optimal range of  $\gamma \sim 0.2\text{--}0.8\kappa$ , most of the RCs (30 for each  $\gamma$ ) achieve effectively the theoretical limit  $\mathcal{F} \sim 0.95$ , and all the poorly performing networks have  $K = 2$  nodes [Fig. 7(a), blue]. Recall that as little as two Kerr nodes are sufficient to perform the readout task; thus as  $K$  increases, so does the probability that some subset of the RC phase space forms a good network. If the other hyperparameters are within some reasonably optimal

window, there is a  $K$  for which the probability of finding a good subnetwork is high enough that performance saturates and any Kerr RC achieves high fidelity. In fact, note that the  $K = 5$  and  $K = 10$  sets of networks show almost equivalent performance, indicating that 5 nodes are sufficient for a random network to reliably complete the present readout task. Finally, we emphasize that the performance of larger reservoirs is more robust to hyperparameter values outside the optimal range.

In Fig. 7(b), we consider the role of nonlinearity and input scaling, plotting  $\mathcal{F}$  obtained using 10 random RCs, while varying  $\mu$  for various fixed values of  $\bar{\Lambda}$ . The nonlinear term in Eq. (8) scales as the node amplitude cubed, and as such the effective nonlinearity is related to both  $\bar{\Lambda}$  and  $\mu$ . To zeroth order,  $\beta \propto \mu$ , and so the effective nonlinearity in  $\beta$  and  $x^\phi$  can be roughly quantified by  $\bar{\Lambda}\mu^2$ . As we saw in Sec. IV C, some nonlinearity is necessary for computation, so  $\mathcal{F}$  initially increases as either  $\bar{\Lambda}$  or  $\mu$  is increased. We encounter an upper limit, however—if nonlinearity is too strong, the fixed points of the RC network for a given steady-state drive become less stable. The dynamics of the reservoir will then generically exhibit large oscillations and not settle near their steady states over the measurement timescale. This is exacerbated by the noise in the input signal; as these fixed points in phase space are less attractive, the strong white noise is able to generate larger excursions in phase space. Thus, as seen in Fig. 7(b),  $\mathcal{F}$  falls off sharply for  $\sqrt{\bar{\Lambda}\mu}$  above some upper limit; we find that  $\sqrt{\bar{\Lambda}\mu}$  should be in the range of 0.5 to 1 for optimal performance, a trend we verified for additional  $\bar{\Lambda}$  which are not shown. This is precisely the regime for a single Kerr oscillator where the nonlinearity begins to significantly influence its dynamics and steady state. RC intuition suggests that dynamics should be affected but not dominated by the nonlinearity; the observation of an optimal nonlinearity strength for a Kerr network appears consistent with this understanding. It should be noted that for this plot we have chosen  $\Lambda_j = \bar{\Lambda}$  to make the  $\bar{\Lambda} - \mu$  relationship clearer, but the results are qualitatively unchanged when randomness in the nonlinearity is reintroduced as well.

The final hyperparameter  $\alpha$  is the largest singular value of the Kerr network connectivity matrix  $\mathbf{W}_R$  and sets the strength of the node-node coupling. In reservoir computing literature, it is commonly stated that for an echo-state network (with a hyperbolic tangent nonlinearity), if the spectral radius of  $\mathbf{W}_R$  is much larger than 1, the RC steady state is not guaranteed to be stable and limit-cycle dynamics can emerge [19,26,29]. A spectral radius close to this limit (the so-called “edge of stability”) is often found to be optimal for tasks requiring significant memory (i.e., where previous states of the input are important) [22,29]. For the Kerr network RC of Eq. (8), the coupling matrix is symmetric, so  $\alpha$  is also the spectral radius of  $\mathbf{W}_R$ . Since the network has an explicit decay term, the linear network is stable for  $\alpha < 2$ . When the nonlinearity is included,

numerically we find a stable steady state and thus the fading memory property is always present for  $\alpha < 2$  for the nonlinearity parameters we consider. In agreement with other works, we have found  $\alpha \sim 1.5\text{--}2$  results in optimal performance, and thus have chosen to present results with  $\alpha = 1.9$ .

Overall, there is generally a broad range of Kerr RC hyperparameters resulting in high-fidelity classification, which we can summarize as  $\gamma \lesssim \kappa$ ,  $\sqrt{\bar{\Lambda}\mu} \sim 0.5\text{--}1$ , and  $\alpha \leq 2$ . This performance is independent of specific network structure, and is robust to moderate disorder in these structural parameters.

## V. QUANTUM INFORMATION APPLICATIONS

In this final section, we demonstrate a pair of relevant quantum information applications that can be implemented in this same reservoir computing system: 2-qubit state tomography and continuous qubit parity monitoring. In both cases the description of Fig. 2 applies: a Kerr network reservoir continually processes the measurement current from a joint dispersive readout system. This is not intended to be an exhaustive survey of potential applications, but to emphasize the ease of generalizing our reservoir processing approach.

### A. Multiqubit tomography

To this point, we have only evaluated the ability of the RC to classify measurement currents from quantum systems prepared in computational basis eigenstates. However, a reservoir processor trained using only these states  $|\psi(0)\rangle = |z\rangle$  can measure *arbitrary* joint qubit states with high fidelity. To be specific, when the target quantum system is interrogated by driving the readout cavity, backaction rapidly causes the joint qubit state to collapse to one of the measurement basis eigenstates, with probability  $\simeq |c_z(0)|^2$ . The RC will then faithfully return the current state of the target quantum system  $|\psi(\tau_m)\rangle$ . If this quantum state is repeatedly prepared and measured, the distribution of RC outputs will thus agree with that of the underlying state. Furthermore, since the measured cavity quadrature is a nonlinear function of the multiqubit operator  $\hat{\chi} = \sum_j \chi_j \hat{\sigma}_{z,j}$  (Appendix B), one can perform full tomography on the 2-qubit density matrix by simply preceding the measurement with a set of single-qubit rotations [5].

To demonstrate this capability, in Fig. 8 we compare the RC output with the true quantum state at the measurement time, for the dispersive system initialized in the indicated product or Bell states. The reservoir is that of Figs. 3 and 4, and in particular has the same  $\mathbf{W}_o$ : training was again done by simply preparing the qubits in each computational basis state 10 times, measuring for  $10/\kappa$ , and using that initial state as the target. Even though the qubit state can jump during this readout, training is still effective, and this approach should be robust to preparation errors as well. The quantum system was then initialized in each of the test superposition states 800 times, and the time-dependent




FIG. 8. Tomography demonstration. The dispersive system is prepared in the indicated initial state and measured 800 times. Measurement causes these superposition states to collapse; the true distribution in the computational basis at  $\tau_m = 7.5/\kappa$  is indicated in red. In blue we plot the distribution of classification outputs for the RC of Figs. 3 and 4, which has not been retrained.

RC output  $P_z(t)$  was compared against the current quantum state  $|c_z(t)\rangle^2$ . Despite this simple training with only easily accessible initial computational basis state labels, the RC is highly successful at producing this dynamical quantum variable, returning the current state  $|\psi(t)\rangle$  ( $\max_z\{P_z(t)\} = \max_z\{|c_z(t)\rangle^2\}$ ) with 98% fidelity across all test states.

Figure 8 compares the distribution of RC outputs with the average quantum state at  $\tau_m$ ; it is clear the RC accurately determines the underlying quantum distribution for the three states tested. For simplicity we have chosen a set of states where all the unique density matrix elements are diagonal in the measurement basis, so this single set of measurements is sufficient to distinguish states. As described in Ref. [5], full tomography on arbitrary states can be done by repeating this process after applying single-qubit rotations; this reservoir processor can thus be an effective tool for general tomography in addition to computational basis measurement.

## B. Joint parity monitoring

This reservoir processing approach is not limited to determining qubit states from measurement currents—generally, one can train a RC to return arbitrary dynamical observables given an appropriate measurement record. In this final example we describe the operation of the  $K = 5$  Kerr reservoir of Figs. 3 and 4 to measure multiqubit parity:  $\langle\hat{\sigma}_{z,1}\hat{\sigma}_{z,2}\rangle$ . Recall that for the dispersive readout quantum system model described previously, the RC was able to learn the quantum state and output  $|c_z(t)\rangle^2$  with very high fidelity: thus, one can trivially modify the output layer to instead return the expected parity via  $\mathbf{W}_o \rightarrow (1, -1, -1, 1) \cdot \mathbf{W}_o$ .

We instead consider a more interesting and relevant task for quantum information applications by modifying the dispersive readout system such that the observed quadrature does not distinguish between states in a given parity subspace, setting  $\chi_1 = -\chi_2 = \kappa$ ,  $\varepsilon_0 = 2i\kappa$ , and leaving all other parameters unchanged. In this situation, the measured quadrature will differ only for states of different parity, and be the same for states in the even  $\{|11\rangle, |00\rangle\}$  or odd parity subspaces  $\{|10\rangle, |01\rangle\}$  [Eq. (B3)]. As a result, when the readout cavity is measured there is no backaction on qubit states in a given parity subspace: this allows one to generate and maintain superposition states, or manipulate the joint qubit state within a parity subspace, a requirement for many quantum error correction protocols. This specific parity readout system has been explored both for Bell state generation and error syndrome monitoring [9,59,60]. In particular, this is a common error syndrome in quantum error correction, where the detection of a change in parity between two qubits is an unambiguous indicator that an error has occurred [60,61].

We operate this system as previously, interrogating the measurement cavity and inputting the resultant measurement current to the RC; this task is thus also described via Fig. 2, with the output now  $\langle\hat{\sigma}_{z,1}\hat{\sigma}_{z,2}\rangle$  instead of  $P_j$ . For training, the quantum system is again prepared in each of the computational basis states  $Q = 40$  times and a measurement is performed for  $10/\kappa$ , with the target being the initial parity,  $\pm 1$  for  $|\psi(0)\rangle = |11/00\rangle, |10/01\rangle$ , respectively. For the parameters chosen, the probability of a parity jump occurring during this window is small due to the lack of measurement backaction; thus this simple training procedure allows the RC to compute the current multiqubit parity. This is demonstrated in Fig. 9(a), where the trained RC continuously monitors the parity as the quantum system is measured over a much longer  $50/\kappa$  window. We plot both the evolution of the parity (ground truth obtained from the SME) as well as the RC output for three sample trajectories. It is seen that the RC output follows the true parity closely, and in particular, quickly switches after the parity jumps due to qubit decay processes. Over 800 test measurements (each of duration  $50/\kappa$ ), the RC and true parity agree at 93% of times. This is limited by the finite response time of the reservoir: on average the RC parity will flip  $2.3/\kappa \sim 1/\gamma$  after that of the quantum system.

For both error syndrome monitoring and Bell state generation, it is only necessary to detect these parity jumps, rather than exactly reproduce the evolution of the parity. In Fig. 9(a), we indicate the times of parity jumps in the quantum system and RC output with circles and stars, respectively. It is clear that for the examples shown, the RC accurately detects when a parity jump occurs, and does not predict one where there is no jump. In Fig. 9(b), we evaluate the ability of the RC to detect parity jumps: the quantum system is initialized in the 8 indicated states (of definite parity), and measured 100 times each over a

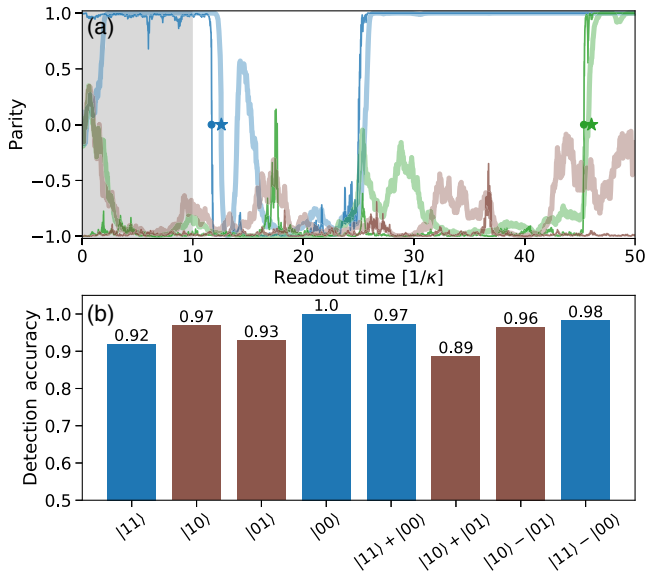


FIG. 9. (a) Ground truth parity trajectories in dark thin lines and corresponding RC parity estimate in light thick lines of the same colors. The times at which the qubit parity jumps are indicated with circles; the resultant RC parity change is denoted with a star. The RC of Figs. 3 and 4 is again used, trained with  $Q = 40$  measurements over the  $10/\kappa$  window shaded in gray. (b) Fraction of trajectories for which parity monitoring was successful over the entire  $50/\kappa$  monitoring window for each initial state. Failures occur when the RC either predicts a parity change where none occurred or fails to detect a parity change. Odd and even initial states are colored red and blue respectively.

duration of  $50/\kappa$ . Plotted is the fraction of measurements in which the RC successfully predicted a parity jump after one occurred in the quantum system, or did not predict a parity jump if none occurred. Overall, parity change detection accuracy was 95.5%, indicating that this RC could be a powerful tool for monitoring error syndromes or tracking the evolution of general observables.

## VI. CONCLUSIONS

In this work, we propose and analyze the operation of a hardware reservoir computer, carefully designed to be implementable in current superconducting platforms, to facilitate measurement of an integrated multiqubit system. We find that reservoirs as small as two oscillators can reliably perform the nontrivial task of joint dispersive quantum measurement of two qubits, achieving fidelity comparable to or exceeding an optimally matched filter while requiring orders of magnitude less training overhead. We further demonstrate this processor can be readily employed in important quantum information tasks such as multiqubit tomography and the continuous monitoring of observables such as parity, with similar fidelity and simplicity of calibration. The proposed approach offers significant efficiency improvements across several dimensions: (i) hardware resources via compatibility with

existing quantum platforms, (ii) computational resources via ease of training, and (iii) latency and data overhead via edge computing in a physical reservoir.

Through a first-principles consideration of the system dynamics, we explore the features and properties of the Kerr oscillator network that enable this performance and make it an attractive platform for reservoir computing more generally. We additionally develop an intuitive phase space picture which provides insight into exactly how RCs process information, and the important role of nonlinearity. Here we have considered the operation of a Kerr reservoir in the semiclassical regime where it is not entangled with the quantum system it measures. The extension of this work to the quantum limit—a challenging theoretical problem requiring a unified quantum description of *both* the measured system and measurement chain, including the reservoir—is addressed in our follow-up work, Ref. [62]. This accounts for the measurement backaction and nonlinearly transformed quantum noise of the quantum reservoir itself, and forms the foundation of a description of quantum measurement using an apparatus that is architecturally complex, nonlinear, and quantum-mechanical.

## ACKNOWLEDGMENTS

We thank Daniel J. Gauthier, Michael Hatridge, Peter L. McMahon, Shyam Shankar, and Nikolas Tezak for discussions. This work is supported by NSERC Grant No. 02547 and AFOSR Grant No. FA9550-20-1-0177.

*Note added.*—Recently, three papers were published applying software machine learning models to single-qubit characterization [63], multiplexed readout [17], and rapid three-state classification [64], respectively. Our work stands in contrast to these studies in that we describe a piece of hardware—a reservoir—whose continuous evolution performs the desired computation. The application of machine learning to quantum information is a rapidly developing field which holds much promise; we believe the integration of RCs with quantum information processing platforms can be of particular benefit.

## APPENDIX A: TABLE OF NOTATIONS

The notation used throughout this paper is summarized in Table I.

TABLE I. Descriptions of symbols, sorted by topic.

Measured quantum system	
$\hat{d}$	Cavity field operator
$\hat{\sigma}_j$	Qubit Pauli operator
$J^{(q)}$	Cavity measurement current, sample $q$
$\kappa$	Readout cavity decay rate
$\omega_d$	Frequency of cavity drive, measurement current

(Table continued)

TABLE I. (Continued)

Measured quantum system	
$\Delta_c$	Readout cavity—drive detuning: $\omega_c - \omega_d$
$\Delta_{q,j}$	Qubit $j$ —drive detuning: $\omega_{q,j} - \omega_d$
$\delta_j$	Qubit $j$ —cavity detuning: $\omega_{q,j} - \omega_c$
$g_j$	Qubit $j$ —cavity direct coupling
$\chi_j$	Qubit $j$ dispersive shift: $g_j^2/\delta_j$
$\gamma_h$	Qubit decay rate (non-Purcell)
$J_{12}$	Qubit-qubit coupling via cavity
$\tau_m$	Duration of measurement signal
Kerr network reservoir	
$\hat{b}_k$	Kerr oscillator $k$ field operator
$\beta_k$	Kerr oscillator $k$ field amplitude
$\mathbf{u}$	Input signals to Kerr network at $\omega_d$
$\Delta_k$	Kerr oscillator $k$ —input detuning: $\omega_k - \omega_d$
$\lambda_k$	Kerr oscillator $k$ nonlinearity
$g_{kl}$	Linear coupling between Kerr oscillators $k$ and $l$
$\varepsilon_{km}$	Input $m$ —Kerr oscillator $k$ coupling strength
$\gamma$	Kerr oscillator decay rate, network evolution rate
Reservoir computer	
$\mathbf{x}(t)$	Reservoir variables: $\sqrt{2}(\text{Re}\{\beta_1\}, \text{Im}\{\beta_1\}, \dots)^T$
$\mathbf{x}^\phi$	Optimized node quadratures: $(\beta_k e^{-i\phi_k} + \beta_k^* e^{i\phi_k})/\sqrt{2}$
$\mathbf{W}_R$	Dimensionless connectivity matrix: $2(\Delta_k \delta_{k,l} + g_{kl})/\gamma$
$\mathbf{W}_I$	Dimensionless input layer: $2\tilde{\varepsilon}_{kl}/\gamma$
$\Lambda$	Dimensionless node nonlinearity: $2\tilde{\lambda}_k/\gamma$
$K$	Number of Kerr oscillators
$\mu$	Input coupling strength: $W_{I,kl} \in [-\mu, \mu]$
$\bar{\Lambda}$	Average node nonlinearity: $\Lambda_k \in [0, 2\bar{\Lambda}]$
$\alpha$	Spectral radius of connectivity matrix: $\lambda_{\max}(\mathbf{W}_R)$
$\mathbf{C}(\phi)$	Matrix of node quadrature projection angles
$\tilde{\mathbf{W}}_o$	Projected quadrature output weights
$\mathbf{W}_o$	Full output layer: $\tilde{\mathbf{W}}_o \mathbf{C}(\phi)$
$Q$	Number of training samples
$\mathbf{y}^*$	$C$ -dimensional target output
$\mathbf{y}$	RC output: $\mathbf{W}_o \mathbf{x}$
$P_z$	RC output converted to probability: $e^{y_z(t)}/\sum_k e^{y_k(t)}$
$C_{\mathcal{F}}$	Fraction of test signals classified correctly
$\mathcal{F}$	Classification fidelity: $\max_t \{C_{\mathcal{F}}(t)\}$

## APPENDIX B: JOINT DISPERSIVE MEASUREMENT

The unconditional evolution of the multiqubit measurement system under either model is described by the master equation (ME) [5,9,42]:

$$\dot{\hat{\rho}} = \mathcal{L}\hat{\rho} = -i[\hat{\mathcal{H}}_S, \hat{\rho}] + \gamma_h \sum_j \mathcal{D}[\hat{\sigma}_{-j}]\hat{\rho} + \kappa \mathcal{D}[\hat{d}]\hat{\rho}, \quad (\text{B1})$$

where  $\kappa$  and  $\gamma_h$  describe cavity loss and qubit decay through coupling to the external environment.  $\hat{\mathcal{H}}_S = \{\hat{\mathcal{H}}_{\text{JC}}, \hat{\mathcal{H}}_D\}$ : for  $\hat{\mathcal{H}}_S = \hat{\mathcal{H}}_D$  the ME acquires an additional term to account for correlated qubit decay via the cavity:  $\mathcal{L} \rightarrow \mathcal{L} + \kappa \mathcal{D}[\sum_j (g_j/\delta_j)\hat{\sigma}_{-j}]$ , although it is  $O((g_j/\delta_j)^2)$

and thus weaker than other rates considered. We could also include pure dephasing as well but consider a regime where this environmental dephasing is much weaker than dephasing induced by measurement.

The utility behind this dispersive measurement system is that it allows one to perform a near-QND measurement of all the qubits simultaneously. We recall first the basic description of the measurement process from Ref. [4], based on the pointer-state formalism. With the cavity initially in the vacuum and a constant drive applied at  $t = 0$ , under  $\hat{\mathcal{H}}_S = \hat{\mathcal{H}}_D$  there is a cavity coherent state associated with each joint qubit state [i.e.,  $\hat{\rho} \propto \sum_{z,z'} p(t)_{z,z'} |\alpha_z(t), z\rangle \langle \alpha_{z'}(t), z'|$ ]. The coherent state amplitudes evolve as

$$\frac{d}{dt} \alpha_z(t) = -i\epsilon_0 - \left( i(\Delta_c + \langle \hat{\chi} \rangle_z) + \frac{\kappa}{2} \right) \alpha_z(t), \quad (\text{B2})$$

where we have defined the qubit operator  $\hat{\chi} = \sum_j \chi_j \hat{\sigma}_{z,j}$ . If  $\chi_j$  are distinct then for each of the  $2^{N_q}$   $\{|z\rangle\}$  qubit states,  $\hat{\chi}$  has a unique value and a different coherent state is generated in the measurement cavity. The corresponding steady states are

$$\alpha_{z,ss} = -\epsilon_0 \frac{\Delta_c + \langle \hat{\chi} \rangle_z + i\frac{\kappa}{2}}{(\Delta_c + \langle \hat{\chi} \rangle_z)^2 + (\frac{\kappa}{2})^2}, \quad (\text{B3})$$

and the measurement of this cavity field results in an effective measurement of the joint qubit state, i.e., all qubits in the  $z$ -basis simultaneously. This measurement becomes QND if  $\hat{\chi}$  commutes with all the operators in  $\mathcal{L}$ ; joint qubit eigenstates are then preserved by the measurement process. This can be achieved in practice in the dispersive regime if  $\kappa$  is larger than  $J_{jk}$  and the qubit decay rates, such that these terms can be ignored on the timescale over which measurement proceeds. Even away from the dispersive regime, where the system is better described by  $\hat{\mathcal{H}}_{\text{JC}}$ , the state of the cavity field still contains information about the joint qubit state, and a sufficiently rapid cavity measurement can be used to learn the initial joint qubit state  $|\psi(0)\rangle$ .

By continuously recording the  $X$  quadrature of the field radiating from the measurement cavity, one obtains the homodyne measurement current of Eq. (4), which provides information about  $\hat{\rho}(t)$ , the current state of the QS. The consequent evolution of an observer's knowledge of the QS, conditioned on observing  $J(t)$ , is found by appending the measurement superoperator to Eq. (B1), resulting in the SME of Eq. (3). Recall that  $\langle \xi(t) \rangle = 0$ , and so by taking the ensemble average of Eq. (3) we recover Eq. (B1) and  $\langle J(t) \rangle = \sqrt{\kappa\eta} \langle d + d^\dagger \rangle$ : density matrix evolution and operator expectation values converge to their unconditional results.

Equation (4) describes a measurement current and corresponding noise which is continuous; both due to finite sampling times in real measurements and for our numerical simulations, the homodyne current is actually sampled at discrete times  $t_n = n\Delta t$ , where we use  $\Delta t = 10^{-2}/\kappa$ . As a consequence,  $\xi(t) \rightarrow \xi(t_n)$  in both Eqs. (3) and (4):

$$\xi(t_n) = \mathcal{N}_n(0, 1)/\sqrt{\Delta t}, \quad (\text{B4})$$

where  $\mathcal{N}_n(0, 1)$  are samples drawn from the normal distribution with zero mean and unit variance. This converges to white noise in the continuum limit  $\Delta t \rightarrow 0$ , with  $\langle \xi(t)\xi(t') \rangle \rightarrow \delta(t-t')$ . The integrated noise power is importantly independent of  $\Delta t$ :  $\int^t d\tau d\tau' \langle \xi(\tau)\xi(\tau') \rangle = t$ , so the SNR of the integrated measurement current, and consequent filter performance, remains fixed for any sampling rate.

The noisy continuous signal described in Eq. (4) is what is seen by any classical system interacting with the readout cavity output field [43]. When we take  $J(t)$  as the input to a RC, there is no entanglement between QS being measured and the classical RC processing this measurement record. In this work we assume each Kerr oscillator has sufficiently large field amplitudes to be in a classical regime. In practice, passing the output of the QS through a quantum-limited amplifier or circulator, as depicted in Fig. 2, results in this description, which is also necessary to describe realizations where the RC and QS are built on different platforms.

### APPENDIX C: RESERVOIR COMPUTER TRAINING

Training amounts to choosing an optimal set of phase angles for each oscillator and linear weights to apply to each resultant node. For a  $K$ -node network and  $C$ -dimensional target output, these are encoded in the  $(C \times K)$  output matrix  $\tilde{\mathbf{W}}_o$  and the  $K$  element vector  $\boldsymbol{\phi} = (\phi_1 \dots \phi_K)$  of measurement angles, for a total of  $(C+1)K$  parameters to optimize. To train a RC, a set of labeled training data  $\{\mathbf{u}^{(q)}(t), \mathbf{y}^{*(q)}(t)\}$ , consisting of  $Q$  (generally multidimensional) input signals  $\mathbf{U} = \{\mathbf{u}^{(q)}(t)\}$  and their respective target outputs  $\mathbf{Y}^* = \{\mathbf{y}^{*(q)}(t)\}$  is constructed. The training data are fed into the RC, producing the dynamical response  $\boldsymbol{\beta}^{(q)}(t)$ , which we decompose into independent quadratures  $\mathbf{x}^{(q)}(t)$ . Importantly, both quadratures must be extracted during training, in order to optimize the projection angle. A loss function depending on the difference between the RC output and target  $\mathcal{L}(\tilde{\mathbf{W}}_o, \boldsymbol{\phi})$  is then constructed from this training set. By minimizing the loss function with respect to  $\tilde{\mathbf{W}}_o$  and  $\boldsymbol{\phi}$ , one hopes to obtain the RC output  $\mathbf{y}^{(q)}(t) = \tilde{\mathbf{W}}_o \mathbf{C}(\boldsymbol{\phi}) \mathbf{x}^{(q)}(t) \rightarrow \mathbf{y}^{*(q)}(t)$  and thus reproduce the desired target function.

The specific task we consider in this work is the retrodiction of the initial state that produced an observed measurement record, requiring that the computation returns the probability that the input is a homodyne record from a QS with  $|\psi(0)\rangle = |z\rangle$ . The training data labels are thus ground truth probabilities for each input signal  $\mathbf{P}^{*(q)} \in [0, 1]$ . The cross entropy loss function is then

$$\mathcal{L}_x(\tilde{\mathbf{W}}_o, \boldsymbol{\phi}) = -\frac{1}{QN} \sum_{q, t_n} \mathbf{P}^{*(q)}(t_n) \cdot \log[\mathbf{P}^{(q)}(t_n)], \quad (\text{C1})$$

where  $\mathbf{P}^{(q)}(t_n)$  is found from  $\mathbf{y}^{(q)}(t)$  via Eq. (13).

Throughout this work, training is done by minimizing the  $\mathcal{L}_x$  of Eq. (C1), with an added  $L_2$  regularization term  $\lambda_r |\mathbf{W}_o|^2$ , via gradient descent. We consider small  $K = 2-10$  RCs and small training sets, so this loss function minimization is a computationally easy task. In our physical RC framework, we envision this training is done using external software to the RC. The RC node trajectories are measured during training, and  $\{\mathbf{x}^{(q)}(t), \mathbf{y}^{*(q)}(t)\}$  are used to minimize  $\mathcal{L}_x$  to compute  $\tilde{\mathbf{W}}_o$  and  $\boldsymbol{\phi}$  for the task at hand.  $\boldsymbol{\phi}$  then sets the projection angles for subsequent RC processing and  $\tilde{\mathbf{W}}_o$  maps these measurements to the RC output. The description here is equally valid if one skips the projection step and just uses both quadratures directly to form the output:  $\mathbf{y}^{(q)} = \mathbf{W}_o \mathbf{x}^{(q)}$ .  $\mathcal{L}_x(\mathbf{W}_o)$  is then simply a function of the general  $(C \times 2K)$ -independent element output matrix, whose minimization remains convex. The RC output  $\mathbf{y}^{(q)}(t)$  is the linear combination of measured node quadratures that expresses the probability of the input measurement record being associated with each state. The softmax function is applied in software and is only necessary to construct  $\mathcal{L}_x$  during training; linear classification using  $\mathbf{y}^{(q)}(t)$  or  $\mathbf{P}^{(q)}(t)$  is equivalent. Thus, outside of training, the only necessary computation to apply to the RC output is the weighting and addition of measured reservoir variables, a simple task to accomplish externally in devices on equally rapid timescales, or on chip with dedicated superconducting electronics for feedback applications.

- 
- [1] J. M. Gambetta, J. M. Chow, and M. Steffen, *Building Logical Qubits in a Superconducting Quantum Computing System*, *npj Quantum Inf.* **3**, 2 (2017).
  - [2] 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).
  - [3] J. Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, *Protocols for Optimal Readout of Qubits Using a Continuous Quantum Nondemolition Measurement*, *Phys. Rev. A* **76**, 012325 (2007).
  - [4] J. Gambetta, A. Blais, M. Boissonneault, A. A. Houck, D. I. Schuster, and S. M. Girvin, *Quantum Trajectory Approach to Circuit QED: Quantum Jumps and the Zeno Effect*, *Phys. Rev. A* **77**, 012112 (2008).
  - [5] S. Filipp, P. Maurer, P. J. Leek, M. Baur, R. Bianchetti, J. M. Fink, M. Göppel, L. Steffen, J. M. Gambetta, A. Blais, and A. Wallraff, *Two-Qubit State Tomography Using a Joint Dispersive Readout*, *Phys. Rev. Lett.* **102**, 200402 (2009).
  - [6] E. Jeffrey, D. Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan,



- A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, *Fast Accurate State Measurement with Superconducting Qubits*, *Phys. Rev. Lett.* **112**, 190504 (2014).
- [7] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, and C. Eichler, *Rapid High-Fidelity Multiplexed Readout of Superconducting Qubits*, *Phys. Rev. Applied* **10**, 034040 (2018).
- [8] J. Ikonen, J. Goetz, J. Ilves, A. Keränen, A. M. Gunyho, M. Partanen, K. Y. Tan, D. Hazra, L. Grönberg, V. Vesterinen, S. Simbierowicz, J. Hassel, and M. Möttönen, *Qubit Measurement by Multichannel Driving*, *Phys. Rev. Lett.* **122**, 080503 (2019).
- [9] K. Lalumière, J. M. Gambetta, and A. Blais, *Tunable Joint Measurements in the Dispersive Regime of Cavity QED*, *Phys. Rev. A* **81**, 040301(R) (2010).
- [10] N. Boulant, G. Ithier, P. Meeson, F. Nguyen, D. Vion, D. Esteve, I. Siddiqi, R. Vijay, C. Rigetti, F. Pierre, and M. Devoret, *Quantum Nondemolition Readout Using a Josephson Bifurcation Amplifier*, *Phys. Rev. B* **76**, 014525 (2007).
- [11] M. Boissonneault, J. M. Gambetta, and A. Blais, *Improved Superconducting Qubit Readout by Qubit-Induced Nonlinearities*, *Phys. Rev. Lett.* **105**, 100504 (2010).
- [12] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, *High-Fidelity Readout in Circuit Quantum Electrodynamics Using the Jaynes-Cummings Nonlinearity*, *Phys. Rev. Lett.* **105**, 173601 (2010).
- [13] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, *Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits*, *Phys. Rev. Applied* **7**, 054020 (2017).
- [14] E. Flurin, L. S. Martin, S. Hacothen-Gourgy, and I. Siddiqi, *Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations*, *Phys. Rev. X* **10**, 011006 (2020).
- [15] A. M. Palmieri, E. Kovlakov, F. Bianchi, D. Yudin, S. Straupe, J. D. Biamonte, and S. Kulik, *Experimental Neural Network Enhanced Quantum Tomography*, *npj Quantum Inf.* **6**, 20 (2020).
- [16] E. Magesan, J. M. Gambetta, A. D. Córcoles, and J. M. Chow, *Machine Learning for Discriminating Quantum Measurement Trajectories and Improving Readout*, *Phys. Rev. Lett.* **114**, 200501 (2015).
- [17] B. Lienhard, A. Vepsilinen, L. C. G. Govia, C. R. Hoffer, J. Y. Qiu, D. Rist, M. Ware, D. Kim, R. Winik, A. Melville *et al.*, *Deep Neural Network Discrimination of Multiplexed Superconducting Qubit States*, *arXiv:2102.12481*.
- [18] H. Jaeger, *Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication*, *Science* **304**, 78 (2004).
- [19] M. Lukoševičius and H. Jaeger, *Reservoir Computing Approaches to Recurrent Neural Network Training*, *Comput. Sci. Rev.* **3**, 127149 (2009).
- [20] G. Van der Sande, D. Brunner, and M. C. Soriano, *Advances in Photonic Reservoir Computing*, *Nanophotonics* **6**, 561 (2017).
- [21] G. Tanaka, T. Yamane, J. B. Hroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, *Recent Advances in Physical Reservoir Computing: A Review*, *Neural Netw.* **115**, 100 (2019).
- [22] J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar, *Information Processing Capacity of Dynamical Systems*, *Sci. Rep.* **2**, 514 (2012).
- [23] L. Gonon, L. Grigoryeva, and J.-P. Ortega, *Approximation Bounds for Random Neural Networks and Reservoir Systems*, *arXiv:2002.05933*.
- [24] L. Appeltant, M. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. Mirasso, and I. Fischer, *Information Processing Using a Single Dynamical Node as Complex System*, *Nat. Commun.* **2**, 468 (2011).
- [25] N. D. Haynes, M. C. Soriano, D. P. Rosin, I. Fischer, and D. J. Gauthier, *Reservoir Computing with a Single Time-Delay Autonomous Boolean Node*, *Phys. Rev. E* **91**, 020801 (R) (2015).
- [26] L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, *High-Speed Photonic Reservoir Computing Using a Time-Delay-Based Architecture: Million Words per Second Classification*, *Phys. Rev. X* **7**, 011015 (2017).
- [27] C. Du, F. Cai, M. A. Zidan, W. Ma, S. H. Lee, and W. D. Lu, *Reservoir Computing Using Dynamic Memristors for Temporal Information Processing*, *Nat. Commun.* **8**, 2204 (2017).
- [28] J. C. Coulombe, M. C. A. York, and J. Sylvestre, *Computing with Networks of Nonlinear Mechanical Oscillators*, *PLoS One* **12**, e0178663 (2017).
- [29] D. Canaday, A. Griffith, and D. J. Gauthier, *Rapid Time Series Prediction with a Hardware-Based Reservoir Computer*, *Chaos* **28**, 123119 (2018).
- [30] A. Griffith, A. Pomerance, and D. J. Gauthier, *Forecasting Chaotic Systems with Very Low Connectivity Reservoir Computers*, *Chaos* **29**, 123108 (2019).
- [31] K. Fujii and K. Nakajima, *Harnessing Disordered-Ensemble Quantum Dynamics for Machine Learning*, *Phys. Rev. Applied* **8**, 024030 (2017).
- [32] S. Ghosh, T. Paterek, and T. C. H. Liew, *Quantum Neuro-morphic Platform for Quantum State Preparation*, *Phys. Rev. Lett.* **123**, 260404 (2019).
- [33] M. Schuld and N. Killoran, *Quantum Machine Learning in Feature Hilbert Spaces*, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [34] L. C. G. Govia, G. J. Ribeill, G. E. Rowlands, H. K. Krovi, and T. A. Ohki, *Quantum Reservoir Computing with a Single Nonlinear Oscillator*, *Phys. Rev. Research* **3**, 013077 (2021).
- [35] S. Ghosh, A. Opala, M. Matuszewski, T. Paterek, and T. C. H. Liew, *Reconstructing Quantum States with Quantum Reservoir Networks*, *IEEE Trans. Neural Net. Learn. Syst.* **32**, 3148 (2021).
- [36] J. Chen, H. I. Nurdin, and N. Yamamoto, *Temporal Information Processing on Noisy Quantum Computers*, *Phys. Rev. Applied* **14**, 024065 (2020).
- [37] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, *Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach*, *Phys. Rev. Lett.* **120**, 024102 (2018).
- [38] D. Canaday, A. Pomerance, and D. J. Gauthier, *Model-Free Control of Dynamical Systems with Deep Reservoir Computing*, *arXiv:2010.02285*.

- [39] L. C. G. Govia and F. K. Wilhelm, *Entanglement Generated by the Dispersive Interaction: The Dressed Coherent State*, *Phys. Rev. A* **93**, 012316 (2016).
- [40] M. Khezri, E. Mlinar, J. Dressel, and A. N. Korotkov, *Measuring a Transmon Qubit in Circuit QED: Dressed Squeezed States*, *Phys. Rev. A* **94**, 012347 (2016).
- [41] M. Malekakhlagh, A. Petrescu, and H. E. Türeci, *Lifetime Renormalization of Weakly Anharmonic Superconducting Qubits. I. Role of Number Nonconserving Terms*, *Phys. Rev. B* **101**, 134509 (2020).
- [42] C. L. Hutchison, J. M. Gambetta, A. Blais, and F. K. Wilhelm, *Quantum Trajectory Equation for Multiple Qubits in Circuit QED: Generating Entanglement by Measurement*, *Can. J. Phys.* **87**, 225 (2009).
- [43] K. Jacobs and D. A. Steck, *A Straightforward Introduction to Continuous Quantum Measurement*, *Contemp. Phys.* **47**, 279 (2006).
- [44] J. R. Johansson, P. D. Nation, and F. Nori, *QuTiP 2: A Python Framework for the Dynamics of Open Quantum Systems*, *Comput. Phys. Commun.* **184**, 1234 (2013).
- [45] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, <http://www.deeplearningbook.org>.
- [46] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, *Deep Physical Neural Networks Enabled by a Backpropagation Algorithm for Arbitrary Physical Systems*, [arXiv:2104.13386](https://arxiv.org/abs/2104.13386).
- [47] P. R. Vlachas, W. Byeon, Z. Y. Wan, T. P. Sapsis, and P. Koumoutsakos, *Data-Driven Forecasting of High-Dimensional Chaotic Systems with Long Short-Term Memory Networks*, *Proc. R. Soc. A* **474**, 20170844 (2018).
- [48] S. Bompas, B. Georgeot, and D. Gury-Odelin, *Accuracy of Neural Networks for the Simulation of Chaotic Dynamics: Precision of Training Data vs Precision of the Algorithm*, *Chaos* **30**, 113118 (2020).
- [49] E. Bollt, *On Explaining the Surprising Success of Reservoir Computing Forecaster of Chaos? The Universal Machine Learning Dynamical System with Contrast to VAR and DMD*, *Chaos* **31**, 013108 (2021).
- [50] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret, *3-Wave Mixing Josephson Dipole Element*, *Appl. Phys. Lett.*, **110**, 222603 (2017).
- [51] S. Puri, C. K. Andersen, A. L. Grimsmo, and A. Blais, *Quantum Annealing with All-to-All Connected Nonlinear Oscillators*, *Nat. Commun.* **8**, 15785 (2017).
- [52] S. E. Nigg, N. Lörch, and R. P. Tiwari, *Robust Quantum Optimizer with Full Connectivity*, *Sci. Adv.* **3**, e1602273 (2017).
- [53] T. Onodera, E. Ng, and P. L. McMahon, *A Quantum Annealer with Fully Programmable All-to-All Coupling via Floquet Engineering*, *npj Quantum Inf.* **6**, 48 (2020).
- [54] C. Aron, M. Kulkarni, and H. E. Türeci, *Photon-Mediated Interactions: A Scalable Tool to Create and Sustain Entangled States of  $N$  Atoms*, *Phys. Rev. X* **6**, 011032 (2016).
- [55] M. E. Kimchi-Schwartz, L. Martin, E. Flurin, C. Aron, M. Kulkarni, H. E. Türeci, and I. Siddiqi, *Stabilizing Entanglement via Symmetry-Selective Bath Engineering in Superconducting Qubits*, *Phys. Rev. Lett.* **116**, 240503 (2016).
- [56] A. Roy and M. Devoret, *Introduction to Parametric Amplification of Quantum Signals with Josephson Circuits*, *C.R. Phys.* **17**, 740755 (2016).
- [57] M. Silveri, E. Zaly-Geller, M. Hatridge, Z. Leghtas, M. H. Devoret, and S. M. Girvin, *Theory of Remote Entanglement via Quantum-Limited Phase-Preserving Amplification*, *Phys. Rev. A* **93**, 062310 (2016).
- [58] In practice, the microwave signal recorded at the room temperature electronics is expected in the ideal case to be a unitary transformation of these complex-valued quadratures. We do not make this distinction here.
- [59] D. Risté, M. Dukalski, C. A. Watson, G. de Lange, M. J. Tiggelman, Y. M. Blanter, K. W. Lehnert, R. N. Schouten, and L. DiCarlo, *Deterministic Entanglement of Superconducting Qubits by Parity Measurement and Feedback*, *Nature (London)* **502**, 350 (2013).
- [60] S. E. Nigg and S. M. Girvin, *Stabilizer Quantum Error Correction Toolbox for Superconducting Qubits*, *Phys. Rev. Lett.* **110**, 243604 (2013).
- [61] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).
- [62] S. A. Khan, F. Hu, G. Angelatos, and H. E. Türeci, *Physical Reservoir Computing Using Finitely-Sampled Quantum Systems*, [arXiv:2110.13849](https://arxiv.org/abs/2110.13849).
- [63] É. Genois, J. A. Gross, A. Di Paolo, N. J. Stevenson, G. Koolstra, A. Hashim, I. Siddiqi, and A. Blais, *Quantum-Tailored Machine-Learning Characterization of a Superconducting Qubit*, [arXiv:2106.13126](https://arxiv.org/abs/2106.13126).
- [64] R. Navarathna, T. Jones, T. Moghaddam, A. Kulikov, R. Beriwal, M. Jerger, P. Pakkiam, and A. Fedorov, *Neural Networks for On-the-Fly Single-Shot State Classification*, *Appl. Phys. Lett.* **119**, 114003 (2021).