SAND2015-2589J

Deterministic generation of remote entanglement with active feedback

Leigh Martin[1,2],* Felix Motzoi[1,3], Hanhan Li[1,2], Mohan Sarovar[4], and K. Birgitta Whaley[1,3]
[1] *Berkeley Center for Quantum Information and Computation, Berkeley, California 94720 USA*
[2] *Department of Physics, University of California, Berkeley, California 94720 USA*
[3] *Department of Chemistry, University of California, Berkeley, California 94720 USA and*
[4] *Department of Scalable and Secure Systems Research, Sandia National Laboratories, Livermore, California 94550 USA*
(Dated: March 25, 2015)

We develop and study protocols for deterministic remote entanglement generation using quantum feedback, without relying on an entangling Hamiltonian. In this context, we derive a semiclassical and a quantum feedback strategy. We show how our protocol, formulated first with discrete-time generalized measurements reduces to a direct feedback protocol in the continuous-time limit whose dynamics can be modeled by a Wiseman-Milburn feedback master equation. We find an analytic solution to this equation which yields a simple, closed-form expression for the optimal feedback. The ability to tune smoothly between continuous-time and discrete-time descriptions of the feedback protocols is a useful feature of our formalism that allows us to derive a superior combined protocol that switches between the quantum and semiclassical protocols as a way to counter measurement inefficiency. We conclude by simulating feedback using capabilities that have already been demonstrated with superconducting transmon qubits, and find that it is possible to deterministically entangle remote qubits using this approach. Our results show that this system is an interesting ground for studying quantum feedback and entanglement generation from both a theoretical and experimental point of view.

1. INTRODUCTION

Engineering of quantum devices requires careful optimization of two essentially contradictory requirements. On one hand, uniquely quantum states such as superposition and entanglement upon which these devices rely are fragile, and require careful isolation from external degrees of freedom. On the other hand, control and measurement of a system requires coupling to an external device, which often runs contrary to the necessity for decoupling from the environment. To balance this tradeoff, many quantum systems that are anticipated to have applications in quantum computing and sensing lack effectively instantaneous projective read-outs. A great deal of recent research has taken advantage of the fact that measurement occurs over a finite, resolvable time scale in such systems. The ability to perform continuous weak measurement enables direct observation of the continuous time evolution of a quantum system, (quantum trajectories) [1–3], and also permits operations on the system *during* the measurement process, including feedback and feed-forward control, which in turn permits the implementation of adaptive measurements.

In particular, quantum feedback control is expected to be broadly applicable to many problems in quantum information science. Quantum feedback was theoretically shown to increase the purification rate of a single qubit undergoing weak measurement [4]. This work has since been expanded in many directions, including analysis of imperfections [5–7], reconsideration of the feedback protocol goals [8] and extension to rapid purification of multi-qubit states [9]. Quantum feedback control has also been applied to develop schemes for quantum error correction [10, 11], transmission of quantum information through noisy channels [12] and implementation of adaptive measurements [13, 14].

Experimentally, quantum feedback has been mostly demonstrated using atomic, molecular, or optical systems, *e.g.*, [15, 16]. However, the recent advent of quantum-limited microwave amplifiers has enabled experimental realization of quantum feedback in superconducting circuits, where it has been used to stabilize the Rabi oscillations of a single qubit [17] and deterministically create local entanglement [18], although the latter worked in the projective measurement limit when considering feedback.

Several works have also suggested using quantum feedback to enhance generation of remote entanglement [19, 20]. These works considered the case in which the controller has access to a joint measurement on a pair of qubits that are too far apart to engineer a direct or photon-mediated interaction, an important scenario for quantum networks or large-scale quantum computers [21]. The remote aspect requires that feedback be restricted to local unitaries, which cannot on their own generate entanglement. A joint measurement alone cannot deterministically project a separable system into a fixed entangled state, but access to this measurement and local unitary feedback can.

In this work, we build upon these results to search for optimal protocols for remote entanglement, focusing on the joint measurement implemented in superconducting qubits [3]. Our approach enables us to closely examine questions of optimality in the goal of entangling remote qubits. *En route* to optimizing for this goal, we use the POVM formalism to continuously tune the measurement

*Electronic address: Leigh@Berkeley.edu

time, and this permits us to identify a protocol that surpasses the locally optimal strategy (as defined in section 3). We also explicitly consider non-ideal measurement efficiency, which turns out to dictate the form of the optimal feedback protocol. These considerations lead us to a semiclassical and a quantum protocol. The ladder only requires experimental resources that have already been demonstrated in Refs. [3, 17]. Using numerical simulations, we show that existing technology suffices to implement our protocol well in superconducting qubits.

The remainder of the paper is organized as follows. In section 2, we describe the measurement that governs our feedback and describe how this measurement is performed on superconducting transmon qubits. Section 3 derives the optimal feedback unitary as a function of the state of the qubits. Before studying the most complete case, we first consider the case of inefficient measurement in section 4 by neglecting coherence terms of the density matrix that are dephased by the measurement, which yields a semiclassical strategy. This strategy is capable of fully entangling the qubits even for non-unit efficiency, but does not work in the continuum limit. In section 5, we use the full density matrix for the system to derive a strategy that is valid in the continuous measurement limit, allowing us to use the Wiseman-Milburn feedback master equation [22, 23] to study the resulting quantum feedback protocol. For the case of unit efficiency, we derive an analytic solution that reaches unit fidelity with respect to the target entangled state. This is not the case for non-unit efficiency, but in section 6 we combine the semiclassical and quantum protocols to develop a superior hybrid protocol. Section 7 concludes with numerical simulations of a realizable scenario.

We note that although the analysis in this paper most immediately applies to superconducting qubits in microwave cavities, it could be adapted to any system in which one can implement a weak measurement of the half-parity observable defined below.

2. ENTANGLEMENT VIA MEASUREMENT

To study feedback in the context of remote entanglement generation, we must first describe the measurement that feedback relies upon. In order to work with a concrete physical system, we consider two solid state qubits dispersively coupled to two cavities, which are themselves connected by a directional coupler such as a microwave circulator as seen in Fig. 1. A sequential probe of the two cavities is made by a cavity-resonant microwave drive, and subsequent measurement of the net phase shift imparted on the probe beam by homodyne measurement can implement a joint measurement of the two qubits [3, 24]. If the phase shifts and qubit dispersive couplings are engineered properly, the phase shifts from the qubit states $|01\rangle$ and $|10\rangle$ can be made equal [24]. As already demonstrated in Ref. [3], this scheme realizes a *half-parity* measurement observable on two qubits

$$\hat{X} = \frac{1}{2}(\sigma_{z1} + \sigma_{z2}), \tag{1}$$

where σ_{zi} is the Pauli operator acting on the ith qubit. As this observable cannot distinguish the states $|01\rangle$ and $|10\rangle$, such a measurement can project the initially unentangled state $\frac{1}{2}(|00\rangle+|01\rangle+|10\rangle+|11\rangle)$ into the entangled state $\frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$ with 50 percent probability. One important advantage of this geometry is that the transmission line used to couple the cavities does not develop standing wave modes, as the circulator enforces directionality of the propagating field. These modes would hybridize with the cavities and introduce qubit dephasing.

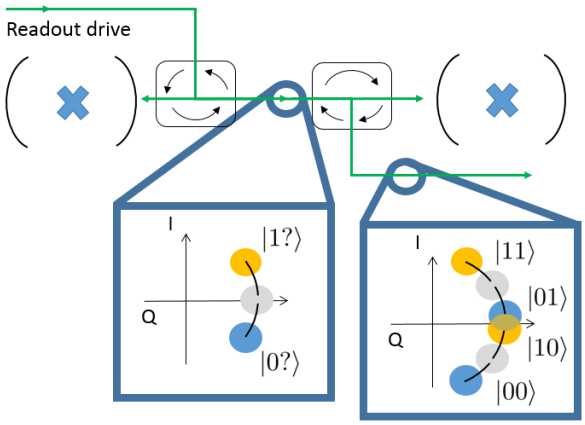

FIG. 1: Joint qubit homodyne measurement with superconducting transmons. Two microwave circulator ensures that the measurement signal reflects off of each cavity only once, so that the imparted phase shifts can cancel in the case that the qubits are in the $|01\rangle$ or $|10\rangle$ states. Inset shows the coherent states after interacting with both qubits. The phase shifts imparted by the qubit on the coherent states manifests as a rotation in the IQ plane. When the coherent states for $|01\rangle$ and $|10\rangle$ overlap, readout maps to a half parity measurement given by Eqn. (1)

In the case that this observable is measured using the sequential probe setup in Fig. 1, the measurement is weak, and a projection onto one of the eigenstates of $\hat{X}$ only occurs after a resolvable time period, during which one can perform feedback on the system based on information already gained. During the measurement, the time-dependent homodyne voltage is described by

$$dV = \langle X \rangle(t)dt + \frac{dW(t)}{\sqrt{8\eta k}}, \tag{2}$$

where $\langle\cdot\rangle$ denotes an expectation value under $\rho(t)$, the state of the two qubit system, k is the strength of the measurement, and $0 \leq \eta \leq 1$ is its efficiency. $dW(t)$ is a Wiener increment satisfying the properties $dW(t)dW(t') = \delta(t-t')dt$, and represents the quantum

noise in the homodyne detection. Note that we have chosen dV to have units of time, so that the average signal $\Delta V = \frac{1}{\Delta t}\int dV$ is unitless. Experimentally, the appropriate scale factor can be extracted simply by preparing states with $\langle X\rangle = \{1, 0, -1\}$ and measuring $\langle \Delta V\rangle$ for any pair.

The state of the system, conditioned on this continuous measurement record is given by the following stochastic master equation (which generates a *quantum trajectory*):

$$\rho_{dV}(t+dt) = \rho(t) + \mathcal{D}[M]\rho(t)dt + \mathcal{H}[M]\rho(t)\sqrt{\eta}dW(t);$$
$$M \equiv \sqrt{2k}X, \tag{3}$$

where $\mathcal{D}[M] \equiv M\rho M^\dagger - 1/2(M^\dagger M\rho + \rho M^\dagger M)$, $\mathcal{H}[M]\rho \equiv M\rho + \rho M^\dagger - \langle M + M^\dagger\rangle\rho$, and the measurement operator is now normalized by the measurement strength k [1]. Note that this equation is in the interaction picture with respect to the free Hamiltonian for the two qubit system, $H_0 = -\frac{\omega_1}{2}\sigma_{z1} - \frac{\omega_2}{2}\sigma_{z2}$.

These equations provide a description of the conditioned dynamics in continuous time. In the following we will also be interested in the conditioned dynamics at discrete time intervals, and in order to obtain such a description of the system, we must compute the finite-time generalized measurement, or POVM [23, 25], associated with the above weak measurement. Such finite-time POVM descriptions are generally difficult to compute for weak measurements, but in the case of such quantum nondemolition (QND) measurements, where $[H_0, X] = 0$, this can be easily done. As shown in Appendix A, the finite time POVM for such weak measurements is composed of a set of *effects* [25], $\Omega_{\Delta V}$, corresponding to the measurement of a voltage increment ΔV defined above during a finite time interval Δt. These effects satisfy the completion property $\int \Omega^\dagger_{\Delta V}\Omega_{\Delta V}d(\Delta V) = \mathbf{1}$, where the integral is performed over the domain of the voltage increment. As shown in Appendix A, the effect corresponding to the above weak measurement, when $\eta = 1$, is given by

$$\Omega_{\Delta V} = \left(\frac{4k\Delta t}{\pi}\right)^{1/4} \exp\left[-2k\Delta t(\Delta V - \hat{X})^2\right]. \tag{4}$$

The state update corresponding to a measurement outcome ΔV is given by the familiar relation $\rho \to \Omega_{\Delta V}\rho\Omega^\dagger_{\Delta V}/\mathrm{Tr}[\Omega_{\Delta V}\rho\Omega^\dagger_{\Delta V}]$. In order to model inefficient measurement, we note that a weak QND measurement with strength k and efficiency $\eta \neq 1$ can be viewed as a sequence of two weak measurements of the same observable, the first with strength $k\eta$ and the second with strength $k(1-\eta)$ which we integrate out to model the loss of this portion of the measurement signal. In view of this, we define a new effect, additionally parameterized by η, as

$$\Omega_{\Delta V,\eta} = \left(\frac{4\eta k\Delta t}{\pi}\right)^{1/4} \exp\left[-2\eta k\Delta t(\Delta V - \hat{X})^2\right]. \tag{5}$$

Then the two-qubit state after evolution by Eqn. (3) for a finite time Δt, conditioned on the finite voltage increment ΔV is

$$\rho_{\Delta V}(t+\Delta t) = \int_{-\infty}^{\infty} \Omega_{\Delta V',1-\eta}\frac{\Omega_{\Delta V,\eta}\ \rho(t)\ \Omega^\dagger_{\Delta V,\eta}}{\mathrm{Tr}[\Omega_{\Delta V,\eta}\ \rho(t)\ \Omega^\dagger_{\Delta V,\eta}]}\Omega^\dagger_{\Delta V',1-\eta}d(\Delta V'). \tag{6}$$

These equations define a general measurement of the observable X, and projective measurement can be recovered by letting $k \to \infty$ or $\Delta t \to \infty$. Measuring X on an intialliy separable state can project into a maximally entangled state, but the probability cannot exceed 50%. In the following sections, we will see that feedback can increse this probability to 1.

3. OPTIMAL ROTATION IN A SINGLE TIME STEP

The only nonlocal resource in our deterministic entanglement scheme is the half-parity measurement detailed in the last section. Due to the lack of direct qubit-qubit coupling, the remaining control resource, the feedback, is restricted to local rotations that act on each qubit individually. Specifically, we define the feedback unitary as

$$U_F[\theta_1, \theta_2, \phi_1, \phi_2] = U(\theta_1, \phi_1) \otimes U(\theta_2, \phi_2) \tag{7}$$

where

$$U(\theta, \phi) \equiv \hat{I}_2 \cos\theta/2 - i\hat{\mathbf{n}}\cdot\hat{\sigma}\sin\theta/2$$

is a general single qubit unitary rotation ($\hat{\sigma}$ is a 3-vector of the Pauli matrices, $\hat{\mathbf{n}}$ is a real 3-vector of unit norm and $\hat{I}_2$ is the identity matrix). It will be sufficient to constrain rotations to a fixed ϕ plane of the Bloch sphere of each qubit, and hence we set $\hat{\mathbf{n}}\cdot\hat{\sigma} = \sigma_x\cos\phi + \sigma_y\sin\phi$. In the presence of dephasing we could also allow σ_z rotations in order to rotate between the $|s\rangle$ and $|t0\rangle$ states, however, our measurement operator is not capable of distinguishing these states, and so we do not have any direct information on how to apply this operation. Hence our construction of optimal strategies will assume negligible decoherence, although we retain ρ_{ss} throughout the analysis in order to study the impact of dephasing in section 7.

Our goal is to find the optimal values of θ_i and ϕ_i as a function of time and measurement outcomes to maximize

[1] Throughout this paper we use units such that $\hbar = 1$.

entanglement. We choose fidelity with respect to the entangled state $|t0\rangle \equiv \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$ as a proxy for entanglement, since this yields a figure of merit that is linear in the state. We simplify the setup further by enforcing identical local feedback unitaries satisfying the properties $\theta_1 = \theta_2$ and $\phi_1 = \phi_2 = \pi/2$. The triplet subspace is closed under local unitary rotations satisfying these constraints. We assume that the initial state is in the triplet subspace, so that we do not require more general rotations that connect the singlet and triplet subspaces.

In order to find the optimal feedback rotation after a measurement, we parameterize the density matrix as follows:

$$\rho = \begin{bmatrix} \rho_{t-t-} & \rho_{t-t0} & \rho_{t-t+} & \rho_{t-s} \\ \rho^*_{t-t0} & \rho_{t0t0} & \rho_{t0t+} & \rho_{t0s} \\ \rho^*_{t-t+} & \rho^*_{t0t+} & \rho_{t+t+} & \rho_{t+s} \\ \rho^*_{t-s} & \rho^*_{t0s} & \rho^*_{t+s} & \rho_{ss} \end{bmatrix}$$

where $|t-\rangle \equiv |00\rangle$, $|t+\rangle = |11\rangle$ and $|s\rangle \equiv \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. As we are restricting to σ_y rotations ($\phi = \pi/2$) and our measurement is represented by a real matrix, we may restrict ρ to be a real matrix. The fidelity with respect to $|t0\rangle$ after applying identical σ_y rotations on both qubits is given by

$$\langle t0|\rho_1|t0\rangle = \rho_{t0t0} + \frac{1}{4}(\sqrt{8}\sin 2\theta(\rho_{t-t0} - \rho_{t0t+}) + (1-\cos 2\theta)(1 - 3\rho_{t0t0} - 2\rho_{t-t+} - \rho_{ss})), \quad (8)$$

where $\rho_1 = U_F[\theta, \theta, \frac{\pi}{2}, \frac{\pi}{2}]\ \rho\ U_F^\dagger[\theta, \theta, \frac{\pi}{2}, \frac{\pi}{2}]$. The optimal rotation angle θ, found by maximizing $\langle t0|\rho_1|t0\rangle$ over θ is,

$$\theta_{\text{opt}}[\rho] = \frac{1}{2}\arctan[\sqrt{8}(\rho_{t-t0} - \rho_{t0t+}), 3\rho_{t0t0} + \rho_{ss} + 2\rho_{t-s} - 1] \quad (9)$$

with $\arctan[y, x]$ behaving like $\arctan[y/x]$, but choosing θ in the correct quadrant, *i.e.*,

$$\arctan[y, x] = \begin{cases} \arctan[y/x] & x > 0 \\ \pi - \arctan[y/x] & x < 0 \end{cases} \quad (10)$$

3.1. *Average-sense* local optimality

Eqn. (9) defines the optimal feedback as a function of the density matrix in a single time step. This relation will define a *locally optimal* protocol (sometimes referred to as a *greedy strategy*), meaning that the controller maximizes the figure of merit at each time step [2]. However this may not be viable for practical applications. Although the controller does in principle have access to the density matrix at every time step, actually calculating $\rho(t)$ (*e.g.*, using Eqn. (6) or Eqn. (3)) amounts to dynamical state estimation, which may be too computationally intensive to implement real-time, and hence experimentally infeasible. Ideally, one would like to provide a protocol that only depends on the measurement outcome from the previous iteration, and also perhaps some deterministic, time-dependent function of the system that can be computed beforehand.

In order to obtain an experimentally feasible protocol we reassess our notion of optimality. Instead of asking that the protocol is locally optimal at each time step for *every* possible measurement trajectory, we instead ask that the protocol is locally optimal on average. For this purpose define

$$\bar{\rho}(t + \Delta t) = \int U_F[\bar{\rho}_{\Delta V}(t)]\bar{\rho}_{\Delta V}(t)U_F^\dagger[\bar{\rho}_{\Delta V}(t)]d(\Delta V) \quad (11)$$

for discrete measurement or

$$\bar{\rho}(t + dt) = \int U_F[\bar{\rho}_{dV}(t)]\bar{\rho}_{dV}(t)U_F^\dagger[\bar{\rho}_{dV}(t)]d(dV) \quad (12)$$

for continuous measurement. $U_F[\rho] = U_F[\theta_{\text{opt}}[\rho], \theta_{\text{opt}}[\rho], \pi/2, \pi/2]$ and $\bar{\rho}$ is the average state after the measurement and applying feedback strategy U_F which is allowed to depend on the previous (and only previous) measurement outcome $\{\Delta V, dV\}$ (to be refered to as V).

In the following sections, we will derive average-sense (a.s.) locally optimal protocols using the following procedure; Instead of using the current state at each time based on the full measurement record, we take the average state as defined above, perform the state update corresponding to the most recent measurement result (using equations 3 or 6), and then insert this state into Eqn. (9) to define $\theta_{\text{opt}}[\bar{\rho}(t), V_i]$ where V_i is the most recent measurement outcome. Explicitly, we define the a.s. local optimal rotation angle at time $t + \{\Delta t, dt\}$ as:

$$\bar{\theta}_{\text{opt}}[V, t] = \frac{1}{2}\arctan[\sqrt{8}(\bar{\rho}_{t-t0,V} - \bar{\rho}_{t0t+,V}), 3\bar{\rho}_{t0t0,V} + \bar{\rho}_{ss,V} + 2\bar{\rho}_{t-s,V} - 1], \quad (13)$$

where $\bar{\rho}_{ij,V}$ are elements of the updated average state, given by Eqn. (3) or Eqn. (6). Because for a given initial state $\bar{\rho}$ evolves deterministically, $\bar{\theta}_{\text{opt}}[V, t]$ can be computed beforehand without relying on dynamical state estimation.

It is useful to observe the symmetries of the framework we have outlined above. Suppose we start in an initial state which is symmetric under the transformation $|0\rangle \leftrightarrow |1\rangle$ on both qubits (henceforth refered to simply as a symmetric state), which implies that $\rho_{t+t+} = \rho_{t-t-}$, $\rho_{t-t0} = \rho_{t0t+}$ and $\langle X\rangle = 0$. Because the target state $|t0\rangle$ respects this constraint, a single round of measurement and feedback will stochastically but symmetrically

[2] The term "local" refers to time-local here, and not spatially local (as in local unitary rotation).

break this symmetry, so that the average state $\bar{\rho}$ will again be symmetric. Thus we assume throughout that $\bar{\rho}$ is symmetric in this way. This constraint is not vital to any derivations that follow, but it simplifies their functional forms. Although the ground state $|11\rangle$ breaks this symmetry, it is locally optimal to apply a $\pi/2$ rotation on both qubits, after which we will be in the state $(|0\rangle + |1\rangle) \otimes (|0\rangle + |1\rangle)$.

Feedback strategies designed in this manner may not perform as well as those in which dynamical state estimation is used at each time step, but the added simplicity makes the resulting protocols substantially simpler to implement experimentally. In addition, in many physical scenarios the real figure of merit is average-case performance of a feedback protocol, and this approach can reproduce the average-case performance of any Markovian, local-in-time (greedy) feedback strategy that performs real-time state estimation, if the cost function is linear in the state (*e.g.*, fidelity with some target state). This is because we have optimized the figure of merit for each possible measurement outcome V_i and the figure of merit is linear, we have ensured that U_F also optimizes the fidelity after integrating over V_i. Note that until Section 6, none of the strategies we formulate address the problem of global optimality.

4. DISCRETE FEEDBACK AND THE SEMICLASSICAL PROTOCOL

Having defined optimality, we can now move on to studying specific feedback protocols. In this section only, we make the additional assumption that the quantum efficiency of the measurement is small, which is a highly relevant scenario in many experimental settings. In this regime, measurement induced dephasing quickly reduces the off-diagonal elements of the density matrix to zero, so that the controller only has access to the classical probabilities associated with the three measurement eigenstates. Without knowledge of the coherences, we will end up deriving a semiclassical protocol, which will provide a useful comparison to the quantum protocol derived in the following section. In section 6, we will see that this feature is important for any non-unit efficiency, and that both the semiclassical and quantum protocols can be useful in most practical settings.

In the small η limit, off-diagonal elements will be approximatley 0. However to study feedback in this case, we explicitly set them to zero when using Eqn. (**??**) to calculate $\theta_{\rm opt}$, so that the first argument of the arctan function in Eqn. (13) is zero. Therefore, $\bar{\theta}_{\rm opt}$ equals 0 or $\pi/2$, with the value decided by whether $3\rho_{t0t0,\Delta V} - 1 \lessgtr 0$ (assuming for simplicity that the singlet subspace is unpopulated). Using equation 6, we see that given an initially symmetric state with fidelity $\bar{\rho}_{t0t0}$, the fidelity after measurement is given by

$$\bar{\rho}_{t0t0,\Delta V} = \frac{\bar{\rho}_{t0t0}}{\bar{\rho}_{t0t0} + \frac{1-\bar{\rho}_{t0t0}}{2}(e^{-4\eta k\Delta t(1+2\Delta V)})(1+e^{16\eta k\Delta t\Delta V})} \tag{14}$$

This yields a threshold behavior for the feedback strategy, in which the desired operation is to do nothing unless $|\Delta V|$ exceeds some critical value $V_{T,\rm opt}$, in which case one should apply a full $\pi/2$ rotation on both qubits (the sign is unimportant). The optimal threshold voltage is given as

$$V_{T,\rm opt} = \frac{1}{8\eta k\Delta t}\text{arccosh}\left[\frac{2\bar{\rho}_{t0t0}}{1-\bar{\rho}_{t0t0}}\exp(4\eta k\Delta t)\right], \tag{15}$$

which defines the *semiclassical protocol*. Note that $V_{T,\rm opt} = 1/2$ in the projective measurement limit $kdt \gg 1$. Eqn. (15) has a simple interpretation. If the state is entangled with high probability, the controller does nothing. If the probability of being in $|t-\rangle$ or $|t+\rangle$ is above a certain threshold, one applies a $\pi/2$ pulse to both qubits, which essentially resets the state to the product state $|\psi_0\rangle \equiv \frac{1}{2}(|0\rangle + |1\rangle) \otimes ((|0\rangle + |1\rangle)$, and gives the joint measurement another chance to collapse to the entangled $|t0\rangle$ state. This strategy is classical in the sense that the optimal feedback could just as easily be calculated by using the classical Bayes rule to combine prior state knowledge with the information gathered from the previous measurement to determine if it is beneficial to apply feedback or do nothing. Note that the larger $\bar{\rho}_{t0t0}$ is, the wider the threshold voltage. Qualitatively, this is because as the fidelity with $|t0\rangle$ increases, it becomes more likely that the qubits are in the entangled state, and so new information needs to suggest with higher probability that we are in $|t-\rangle$ or $|t+\rangle$ before applying a rotation is beneficial, consistent with the Bayesian interpretation.

If one instead uses a fixed threshold value for feedback, one can analytically solve for the steady state fidelity under this protocol:

$$\langle t0|\bar{\rho}(\infty)|t0\rangle = 1 - \frac{4\text{erfc}[\frac{V_T}{\sigma}]}{\text{erfc}[\frac{(V_T+1)}{\sigma}] + 4\text{erfc}[\frac{V_T}{\sigma}] + \text{erfc}[\frac{(V_T-1)}{\sigma}]}, \tag{16}$$

where $\sigma \equiv 1/\sqrt{4k\eta\Delta t}$ is the standard deviation of the measurement operator, and erfc is the complement error function, $1 - \text{erf}(x)$. We can see by taking limits of this expression that one can only reach unit steady-state fidelity in the limit that $V_T \gg \sigma$, or in other words, the threshold voltage is much greater than the variance of the Gaussian weak measurement effect, Eqn. (5).

Fig. 2 shows the fidelity with $|t0\rangle$ as a function of discrete feedback iteration number using two constant threshold strategies and the a.s. locally optimal threshold strategy based on Eqn. (15). Maximally mixed states are an interesting initial state to consider in feedback, but

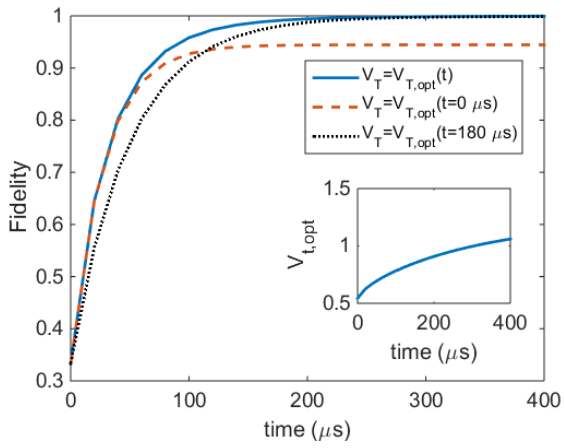

FIG. 2: Discrete-time feedback simulations, showing fidelity with the $|t0\rangle$ state as a function time, starting from maximally mixed state in the triplet subspace. For these simulations, $\eta = 0.1$, $k = 1$MHz and $\Delta t = 20\mu s$. The inset shows the optimal threshold voltage as a function of time. Note that for a smaller threshold, fidelity increases quickly at first but then saturates to a value significantly less than one, while for a larger constant threshold, fidelity increases slowly at first but then surpasses the former and approaches unity, though it does not asymptote to 1 (see Eqn. (16)). The locally optimal strategy, which increases the threshold as a function of time surpasses both fixed-threshold strategies at all times.

as we have no way to detect population in s, we instead set the initial state to the maximally mixed state within the triplet subspace. Even though the initial state is diagonal, for which the semiclassical protocol of Eqn. (15) is a.s. locally optimal, off-diagonal elements not exactly zero after the first iteration due to the application of $\pi/2$ rotations, and thus the strategy is not a.s. locally optimal for later iterations. As shown in Fig. 2, among the thresholding strategies, only this semiclassical strategy approaches unit fidelity, because $V_{T,\mathrm{opt}} \gg \sigma$ at late times.

The above defines a viable discrete feedback protocol. In fact, a similar result using a full parity measurement (as opposed to the half-parity measurement we use here) has already been used to deterministically entangle qubits located in the same cavity [18] (without proof of any optimality properties).

5. THE CONTINUOUS-TIME CASE AND THE QUANTUM PROTOCOL

If one attempts to derive a continuous-time protocol using the above result, the increase in fidelity becomes arbitrarily slow in the small Δt limit. One can see the issue immediately when the $\Delta t \to 0$ limit of the threshold voltage in Eqn. (15) is taken. $V_{\mathrm{T,opt}}$ diverges as $1/dt$, but the standard deviation of the measurement outcomes diverges more slowly, as $1/\sqrt{dt}$ (the fact that these quantities diverge is an artifact of our normalization convention for ΔV. No physical observable diverges in this limit). Thus the probability that this feedback strategy will have us perform any operation on the state vanishes. We will have to include coherences, and take the full density matrix into account, in order to derive a viable continuous-time protocol.

Eqn. (13) itself defines an a.s. optimal protocol, and we can use it to derive our continuous-time feedback strategy directly without setting the off-diagonal terms of ρ to zero. Fig. 3 shows the performance of this discrete time strategy for various choices of Δt, and compares it against the semiclassical protocol derived in the last section. Not surprisingly, Fig. 3 shows that the quantum protocol has strictly better performance over the semiclassical protocol. However, the performance gap between them closes as Δt increases. The reason for this is that the measurement-induced dephasing and the time interval between feedback operations increases as Δt increases, and therefore the density matrix that the feedback is calculated for becomes closer and closer to a diagonal density matrix One can see this behavior explicitly in Fig. 3 (b), in which we plot the applied feedback as a function of the measurement outcome. At late times, it resembles the semiclassical protocol. We will revisit this in section 6.

So far we have only considered discrete feedback, in which the measurement and feedback unitary rotations act sequentially. To derive a continuous protocol, we take the measurement strength to be small and assume that infinite strength rotations are available (although this unphysical assumption be will unnecessary, as shown below).

Assuming $\bar{\rho}$ is symmetric, $dV \propto dW$ because $\langle X\rangle = 0$. Furthermore, Eqn. (3) yields that the quantity $\rho_{t-t0,dV} - \rho_{t0t+,dV}$ appearing in the first argument of the arctan function in Eqn. (13) is infinitesimal and proportional to dW. Given this, we can substitute the continuous-time measurement update in Eqn. (3) into Eqn. (13), and expand in a Taylor Series in $\bar{\rho}_{t-t0,dV} - \bar{\rho}_{t0t+,dV} \propto dW$. Assuming the second argument in the arctan function is greater than zero [3], this yields a proportional feedback strategy in which the feedback rotation is proportional to the measurement outcome dW via

$$\bar{\theta}_{\mathrm{opt}} = \frac{4\sqrt{k\eta}\bar{\rho}_{t-t0}}{3\bar{\rho}_{t0t0} + \bar{\rho}_{ss} + 2\bar{\rho}_{t-t+} - 1} dW, \tag{17}$$

or,

$$\bar{P}_{\mathrm{opt}} \equiv \frac{\bar{\theta}_{\mathrm{opt}}}{dV} = \frac{8\sqrt{2}k\eta\bar{\rho}_{t-t0}}{3\bar{\rho}_{t0t0} + \bar{\rho}_{ss} + 2\bar{\rho}_{t-t+} - 1} \tag{18}$$

[3] If we start in the separable state $(|0\rangle+|1\rangle)\otimes(|0\rangle+|1\rangle)$ as is most practical, this constraint will hold true for all later times. If we start in a state that does not satisfy this property, it is locally optimal to apply a $\pi/2$ rotation on both qubits, after which the second argument will be greater than zero.

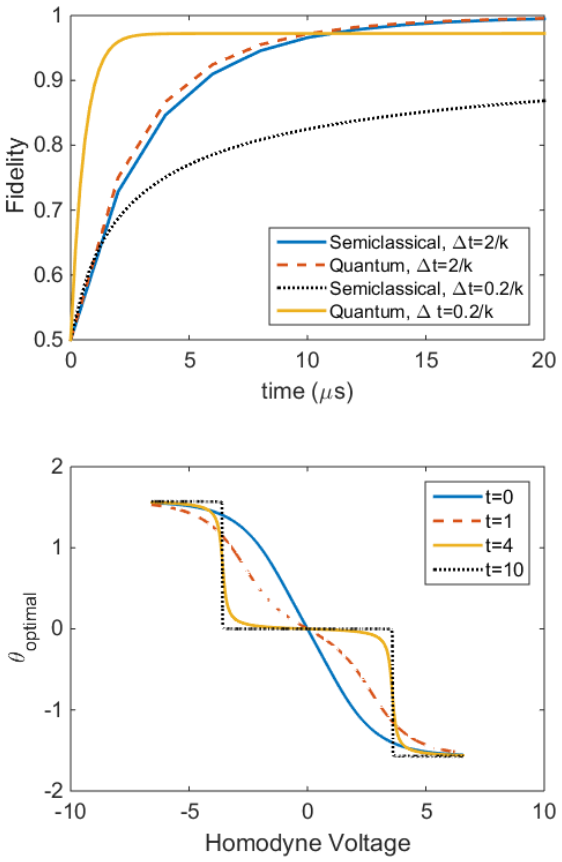

FIG. 3: Discrete feedback simulations, showing fidelity with the $|t0\rangle$ state as a function time under the semiclassical and quantum feedback (Eqn. (13)) strategies, starting in the product state $|\psi_0\rangle$ and assuming a unit efficiency measurement and $k = 1$MHz. The performance also changes according to the discrete time step, Δt, and we show two representative cases.

This shows that the a.s. locally optimal feedback is in fact *direct feedback* in the continuous-time limit, where the feedback rotation angle is directly proportional to the measurement value. This type of feedback is generated by a Wiseman-Milburn feedback master equation [22, 23] that we can use to study this protocol further, given by

$$d\rho = \mathcal{D}[M]\rho dt + \mathcal{H}[M]\rho\sqrt{\eta}dW - i[H_F,\rho]\frac{dW}{\sqrt{8\eta k}} - i[H_F,\{M,\rho\}]\frac{dt}{\sqrt{8k}} + \mathcal{D}[H_F]\rho\frac{dt}{8k\eta}, \quad (19)$$

with

$$H_F = \frac{\bar{P}_{\rm opt}}{2}(\sigma_{y1}+\sigma_{y2}). \quad (20)$$

$d\bar{\rho}$ may be calculated using the averaged version of Eqn. (19) *i.e.*, integrating out dW. Substituting Eqn. (18) into Eqn. (19) yields the equations of motion for the state under the *a.s. locally optimal, quantum, continuous-time protocol*. This equation is difficult to solve in the general case, however, for $\eta = 1$ it admits an analytic solution when the initial state is pure and satisfies the usual symmetry property. If we take the ansatz that the state remains pure throughout its evolution, $\rho = |\psi\rangle\langle\psi|$ takes the form:

$$|\psi\rangle = \left[\sqrt{\tfrac{1-\rho_{t0t0}}{2}},\ \sqrt{\rho_{t0t0}},\ \sqrt{\tfrac{1-\rho_{t0t0}}{2}},\ 0\right]^{\mathsf{T}} \quad (21)$$

Substituting this form into Eqn. (19) yields the following differential equation for ρ_{t0t0}

$$\frac{d\rho_{t0t0}}{dt} = 2k(1-\rho_{t0t0}) \implies \rho_{t0t0}(t) = 1-(1-\rho_{t0t0}(0))e^{-2kt} \quad (22)$$

Crucially, the terms in Eqn. (19) that depend on dW cancel upon substitution. This cancellation ensures that the actual state equals the average state at all times, so that the non-averaged evolution is also deterministic (and hence the omission of the overbar on ρ above). Note that equality between the average sense and actual evolution implies that dynamical state estimation is not necessary to implement the locally optimal strategy. Having a solution for the evolution of the full density matrix also yields an analytic solution for the optimal proportionality coefficient as a function of time:

$$P_{\rm opt} = \frac{4k(1-\rho_{t0t0}(0))}{\sqrt{(1-\rho_{t0t0}(0))(\rho_{t0t0}(0)-1+e^{2kt})}} \quad (23)$$

This equation displays a marked similarity with the optimal feedback for single qubit purification [4], though the functional form differs by two minus signs. Like our protocol outlined above, the 1 qubit state evolution under optimal feedback is deterministic, and has been shown to be globally optimal for $\eta = 1$ [4, 6, 7]. These parallels lead us to speculate that the protocol given by Eqn. (23) may be globally optimal as well.

Fig. 4 shows the fidelity as a function of time using the quantum protocol for $\eta = 1$. The asymptotic behavior of these protocols is interesting to study. Numerical simulations show that with the a.s. locally optimal strategy, the fidelity asymptotes quickly to one, but one might wonder if fixing P could also yield unit fidelity. By fixing $d\rho = 0$ in Eqn. (19), we obtain a system of linear equations which can be solved to yield the steady state fidelity for fixed P of

$$\mathrm{Fid}_{\rm SS} = \frac{P^2+16k^2\eta(1+8\eta)}{3P^2+16k^2\eta(3+8\eta)}, \quad (24)$$

valid for $P \neq 0$ (when $P = 0$, any state that commutes with the measurement is a steady state). For large P, the steady state fidelity tends to 1/3, and in the limit $P \to 0$, the steady state fidelity is $(1+8\eta)/(3+8\eta)$. Moreover, the denominator in Eqn. (24) is greater than the numerator for all values of parameters, and therefore $\mathrm{FID}_{\rm SS} < 1$.

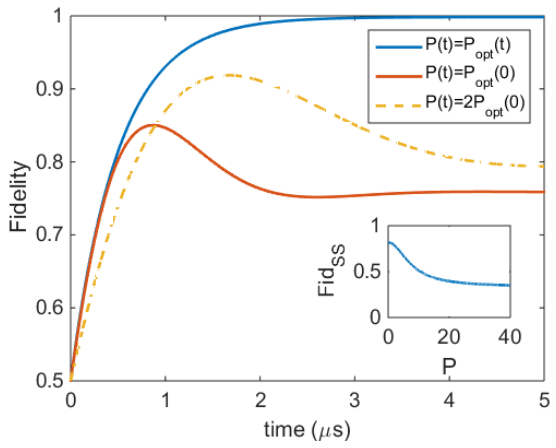

FIG. 4: Continuous feedback simulations, showing fidelity with the $|t0\rangle$ state as a function time under the continuous-time a.s. locally optimal protocol prescribed by Eqn. (18). Simulation parameters are $k = 1$MHz, $\eta = 1$ and a time step $dt \ll 1/k$. Fidelity as a function of time for several values of constant P are also shown for comparison. Inset shows the steady state fidelity as a function of P as given by Eqn. (24). As we have taken the small dt limit, the threshold strategy would not change from its fidelity at $t = 0$. Also note that if we were to plot the locally optimal strategy *i.e.*, one using dynamical state estimation, it would coincide exactly with the curve for $P_{\rm opt}(t)$.

This means that a time-varying direct feedback protocol is necessary to reach fidelity 1 at long times.

This protocol has several appealing features for experimental realization. Firstly, proportional feedback is realizable simply by using a mixer or analog multiplier [17], which have very low latency. Furthermore, unlike with threshold feedback, we are able to increase the fidelity to first order in $\bar{\theta}_{\rm opt}$, which means that only infinitesimal rotations are required.

6. INEFFICIENT MEASUREMENT AND THE CASE FOR DISCRETE FEEDBACK

For unit efficiency, the fidelity quickly reaches 1 when continuous-time feedback is used. But for $\eta < 1$, the fidelity asymptotes to a value less than 1 (see Fig. 5). Qualitatively, this happens because the off-diagonal terms that drive feedback in the continuous-time case decay faster relative to the feedback terms. If we implement discrete-time feedback with the same $\eta < 1$, the fidelity increases more slowly at first, but eventually surpasses the asymptotic fidelity of the continuous-time strategy. Moreover, we know from Sec. 4 that the semiclassical strategy is unaffected by decay of the off-diagonal terms, and this strategy does reach a fidelity of 1. However this strategy is only viable if the measurement time is finite. These facts suggest that we consider a *switching protocol* that transitions between continuous-time and discrete-time feedback; *i.e.*, has a variable measurement duration.

To study how to make this transition, we divide the system evolution into 200 discrete steps each of duration Δt_i subject to the constraint $\sum_i \Delta t_i = T_{\rm final}$. At each time step, discrete-time measurement, according to Eqn. (5), followed by feedback according to Eqn. (13) is applied. Then an optimization over all Δt_i was performed (using the gradient descent algorithm) to minimize the cost function $1 - \langle t0|\bar{\rho}(T_{\rm final})|t0\rangle$.

The results of this optimization are shown in Fig. 5, where we plot fidelity as a function of time, and also and compare to cases where Δt_i is fixed to a finite or an infinitesimal value. The optimization consistently finds a minimum in which the majority of the Δt_i intervals are small and approximately equal, while a few at the end are large. In other words, the optimal solution found by gradient descent is a sudden transition between continuous-time and discrete-time feedback. After the switching time, the off-diagonal elements can be observed to be small, and thus the applied feedback closely resembles that of the semiclassical protocol.

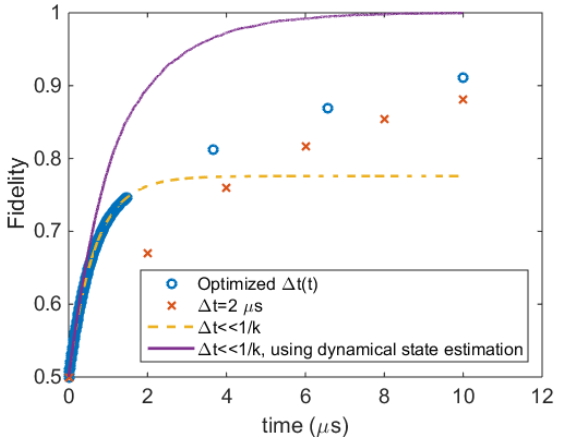

FIG. 5: a) Fidelity with the $|t0\rangle$ state as a function time under the switching protocol . Measurement of a constant rate is on continuously, and feedback is applied instantaneously at each point in the graph. The parameters are $k = 1$MHz and $\eta = 0.4$ to illustrate the benefit of this optimized strategy for non-unit efficiency. The dense cluster of points before $t = 2\mu s$ shows that the gradient search placed almost all of the times to apply feedback at the very beginning. We also plot feedback with two fixed values of Δt for comparison. Also shown is the fidelity using dynamical state estimation and locally optimal feedback.

As we keep the measurement rate (k) constant in the above optimizations, an alternative way to view this change in measurement time is to view finite duration measurements as a series of infinitesimal measurements. With this perspective, we see that discrete feedback abstains from applying feedback at one time only to apply stronger feedback at a later time, and thus is not locally optimal. Although we cannot prove any optimality properties of this switching protocol, its superior performance

shows that it is closer to global optimality than the fixed Δt protocols derived in previous sections.

Fig. 5 also shows that there exists a non-Markovian strategy which outperforms the switching protocol described above. Simulating continuous feedback and averaging over 300 trajectories, we plot the fidelity as a function of time assuming that the controller can perform dynamical state estimation at each time step, and therefore apply Eqn. (9). This strategy surpasses the switching protocol at both early and late times, showing that unlike in the $\eta = 1$ case, dynamical state estimation can yield a better protocol than one based only on the average state. However, as noted in section 3.1, this protocol is more difficult to implement in practice.

7. EXPERIMENTAL REALIZATION

In any experimental realization of the proposed feedback schemes to achieve entanglement between remote qubits, numerous imperfections will complicate the dynamics studied above. Firstly, finite coherence of the qubits will limit the fidelity. Next, the above results apply only to truly Markovian dynamics, in which the controller acts instantaneously solely based on the current measurement outcome. In practice, the homodyne measurement will have finite bandwidth, the feedback will act with some finite delay. In this section we present simulations of the continuous-time quantum feedback protocol specified in Eqn. (18), with these imperfections incorporated. and experimental parameters that have already been achieved in Ref. [3] for the joint homodyne measurement, and in Ref. [17] for continuous proportional feedback in superconducting circuit systems with the 3D transmon qubits. We average over 1500 stochastic trajectories in order to obtain the average performance, and the results are shown in Fig. 6. In order to apply the correct feedback protocol, we have made the approximation that non-Markovian effects are small. In this limit, we can run a series of simulations to calculate the average evolution $\bar{\rho}(t)$ using Eqn. (12) which incorporate decoherence but not feedback delay or finite bandwidth effects. From these simulations, we extract $\bar{P}_{\rm opt}(t)$ and apply this to the stochastic simulations. We only apply continuous feedback, because for these particular parameters, dephasing prevents the fidelity from reaching the switching point observed above.

Figure 6 shows that existing technology suffices to implement our proposal and deterministically entangle two transmon qubits, as both fidelity and concurrence substantially surpass the entanglement threshold to a more than detectable degree. Although our feedback protocols are motivated by deterministic entanglement generation, one can also add in post-selection as was necessary for measuremnt induced entanglement without feedback. In this light, we can also use feedback to enhance the fidelity. Figure 6 shows the fidelity and concurrence by using a 50% post-selection, in which the trajectories for which

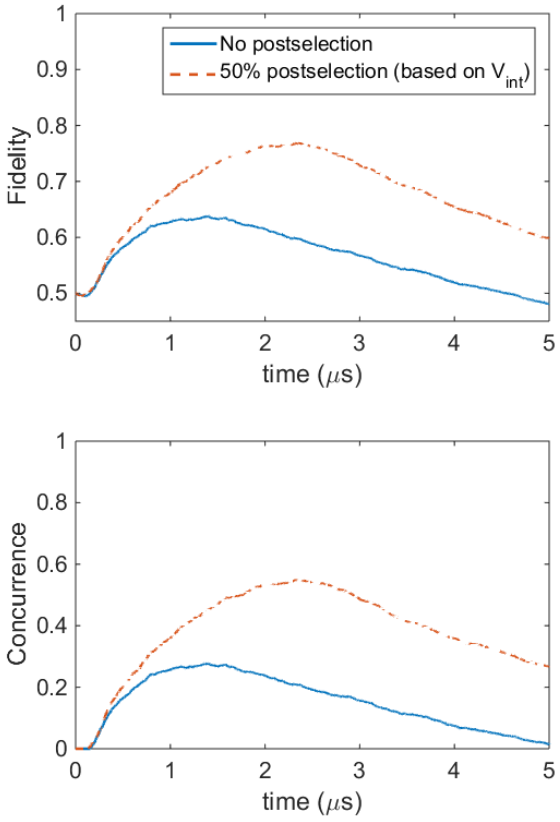

FIG. 6: Continuous feedback simulations, showing fidelity with the $|t0\rangle$ state as a function time under the continuous-time a.s. locally optimal protocol prescribed by Eqn. (18), and incorporating experimentally realistic parameters and imperfections. The parameters used are $\eta = 0.4$, and feedback delay and feedback loop bandwidth of $100ns$ and $1.6MHz$ respectively. Loss of 0.04 in amplitude units in between cavities (mainly due to a circulator) and an intrinsic qubit dephasing rate of $30\mu s$ yields an effective dephasing rate of $6.9\mu s$ for the first qubit and $30\mu s$ on the second. In addition, $T_1 = 20\mu s$, a relaxation time routinely achieved with 3D transmon qubits. The fidelity reaches a much lower peak value than it would in the ideal case before decaying due to finite dephasing. Despite this, the fidelity substantially surpasses 50%, which is the threshold for entanglement. To directly quantify entanglement, we also plot the concurrence [26] as a function of time.

the integrated signal from 0 to $2.7\mu s$ is less than the median. This post-selection method does not require any advanced processing of the signal, and greatly enhances the fidelity and concurrence.

In Ref. [17], propagation through cables introduces the largest contribution to delay, and this delay has a measurable effect on the results of the simulation. It is also conceivable that the delay could be much larger, if for example the qubits are physically separated to a large extent, and one might wonder if our feedback protocol remains

viable in this context. The answer to this question lies in the fact that the measurement strength k effectively sets the timescale for all feedback dynamics, as is evident in Eqn. (19). One can make the substitutions $k \to k/s$, $dt \to dts$ and $dW \to dW/\sqrt{s}$ and recover the same equation. Thus the effect of delay can be negated simply by reducing the measurement strength, which slows down all of the dynamics. Of course, one still encounters a limit set by the other absolute time scales inherent to the problem: T_1 and T_2 processes. However, this argument shows that one can always strike a balance between delay and dephasing to optimize fidelity, and only coherence times will fundamentally limit the achievable entanglement.

In practice, increasing the distance between the qubits will also increase photon loss between cavities. As the measurement probe has not yet interacted with the second qubit at this point, loss between cavities introduces additional dephasing on the first qubit. Indeed, loss between the cavities in Ref. [3], which was dominated by the microwave circulator, presented a significant source of dephasing, and this source of loss is likely to present a fundamental limitation to the maximum achievable entanglement using this scheme.

8. CONCLUSION

We have developed several new and experimentally viable protocols for remote entanglement generation via feedback. By using measurement as a resource for entanglement, we find several protocols are capable of deterministically creating entanglement, in contrast to the probabilitic methods of heralded entanglement. The first strategy considered, the semiclassical protocol, considers only the classical probabilities for being in the entangled or unentangled states. This feedback strategy resembles the strategy that is found to be locally optimal for the single qubit purification case in the small η limit, although in this scenario, feedback is only used to counter the effects of qubit decay [7]. In the case of our quantum feedback scheme, we found that the optimal feedback is simple proportional feedback, which is easily modelled using the Wiseman-Milburn equation and relatively easy to implement in experiment using a mixer. In the case of unit efficiency, we found an analytic solution for the form of the optimal feedback in which the state evolves deterministically. This solution bears a striking resemblance to the optimal feedback in the case of qubit purification for unit efficiency measurement. In the case of inefficient measurement, we identified a transition in which one could benefit by switching between the quantum and semiclassical feedback protocols. Again, an analogous behavior is observed in single qubit purification, though we have taken the additional step of optimizing the transition time.

Having found several protocols for entanglement generation using feedback, we showed that these protocols suffice to deterministically generate entanglement using existing superconducting 3D transmon technology. We also show that if the controller is allowed to use post-selection, the fidelity can be further enhanced using a simple scheme based on the integrated signal. Our work could be applied to atoms as well, provided that a suitable joint homodyne measurement could be performed.

We believe there are several future directions that would be interesting from both a theoretical and an applied point of view. Firstly, one could apply the verification theorems to address the question of whether or not our protocols are globally optimal [27]. In the case of unit efficiency and no decoherence, we speculate that our solution is globally optimal, though we have not attempted to prove this statement. The deterministic evolution of the state under the protocol we derived shows that it is promising, in part because this feature is also observed in the globally optimal protocol for single-qubit purification. Secondly, while fidelity is a useful figure of merit for generating known states, concurrence may be more suitable for a detailed theoretical study and could yield additional insights. For instance, unlike fidelity, concurrence is invariant under the allowed feedback Hamiltonians, which agrees with the behavior of entanglement. Finally, we have ignored dephasing until studying our protocols in a realistic experimental setting, in which we found that it presents the main limitation. Other protocols might focus on correcting for dephasing, and hence creating a means for remote entanglement stabilization.

Acknowledgments

We thank Mollie Schwartz and Irfan Siddiqi for many useful discussions. The effort of LM was supported by grants from the National Science Foundation. Sandia is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL8.

Appendix A: Derivation of finite time POVM

As discussed in the main text the stochastic master equation (SME) associated with a continuous-time QND measurement of the Hermitian observable X with strength k, on a system with free Hamiltonian H is:

$$d\rho = -i[H,\rho]dt + 2k\mathcal{D}[X]\rho dt + \sqrt{2k}\mathcal{H}[X]\rho dW(t). \quad \text{(A1)}$$

The superoperators $\mathcal{D}$ and $\mathcal{H}$ are defined in the main text. We assume that the measurement is unit efficiency, and since it is QND, $[H,X]=0$. This equation describes the evolution of the system conditioned on the measured voltage

$$dV(t) = \langle X \rangle (t)dt + \frac{dW(t)}{\sqrt{8k}} \quad \text{(A2)}$$

The *linear* stochastic master equation [28–30] associated with this equation is:

$$d\bar{\rho} = -i[H,\bar{\rho}]dt + 2k\mathcal{D}[X]\bar{\rho}dt + \sqrt{2k}\bar{\mathcal{H}}[X]\bar{\rho}dW(t), \quad \text{(A3)}$$

with $\bar{\mathcal{H}}[A]\bar{\rho} \equiv A\bar{\rho} + \bar{\rho}A^\dagger$. While Eqn. (A1) is nonlinear in ρ and produces a normalized density matrix (*i.e.*, $\mathrm{tr}(d\rho) = 0$), Eqn. (A3) is linear in $\bar{\rho}$ and produces an unnormalized density matrix. Further, since this equation is linear we can restrict our focus to its action on pure states (since any density matrix is a convex sum of pure state density matrices). This defines the associated linear stochastic Schrodinger equation (SSE) [28]:

$$d\left|\bar{\psi}\right\rangle = \left[\left(-iH - kX^2\right)dt + \sqrt{2k}dW(t)X\right]\left|\bar{\psi}\right\rangle \quad \text{(A4)}$$

This is again a linear equation in the state, $\left|\bar{\psi}\right\rangle$.

Both the linear SME and linear SSE sacrifice the normalization of the resulting state for linearity. To see what this means physically, note that both equations are consistent with $\langle X\rangle = 0$, and hence $dV(t) = \frac{dW(t)}{\sqrt{8k}}$, which means that we are generating conditional dynamics according to statistics associated with some fictitious (time-independent) state with property $\langle X\rangle = 0$. The real state $|\psi\rangle$ may not have this property and this is why the normalization is incorrectly predicted by Eqn. (A4).

Despite this issue with unnormalized states, the utility of Eqn. (A4) is that it is sometimes possible to analytically solve for $\left|\bar{\psi}\right\rangle$, see for example Ref. [30]. In fact, the linear SSE in Eqn. (A4) is in the easiest class of such equations to solve since all the operators in it commute. We are interested in the solution to Eqn. (A4) over a small finite time Δt, which is explicitly [30])

$$\begin{aligned}\left|\bar{\psi}_{\Delta W}(t+\Delta t)\right\rangle &= e^{-iH\Delta t}e^{-2kX^2\Delta t+\sqrt{2k}X\Delta W}\left|\bar{\psi}(t)\right\rangle \\ &= e^{-iH\Delta t}e^{-2kX^2\Delta t+4kX\Delta V}\left|\bar{\psi}(t)\right\rangle \\ &\equiv \bar{\Omega}_{\Delta V}\left|\bar{\psi}(t)\right\rangle, \end{aligned} \quad \text{(A5)}$$

for any initial state $\left|\bar{\psi}(t)\right\rangle$ (including normalized states), and where $\Delta W = \int_t^{t+\Delta t} dW(t')$, is a Gaussian random variable. In the second line we have used the relation between the measured voltage and ΔW to write the solution in terms of the measured quantity, ΔV, which has probability distribution

$$P_{\Delta V} = \sqrt{\frac{4k}{\pi\Delta t}}e^{-\frac{4k\Delta V^2}{\Delta t}}. \quad \text{(A6)}$$

Eqn. (A5) generates states that do not quite yield correct predictions since the normalization of the state is incorrect. Consider a normalized initial state $|\psi(t)\rangle$. Then the probability of evolving according to the stochastic variation ΔV to time $t+\Delta t$ is not actually $\langle\bar{\psi}(t+\Delta t)|\bar{\psi}(t+\Delta t)\rangle$, but rather [28–30]

$$P_{\Delta V}\langle\bar{\psi}(t+\Delta t)|\bar{\psi}(t+\Delta t)\rangle = P_{\Delta V}\,\langle\psi(t)|\,\bar{\Omega}^\dagger_{\Delta W}\bar{\Omega}_{\Delta W}\,|\psi(t)\rangle\,.$$

This expression immediately tells us that the effect associated to the true finite time evolution is

$$\Omega_{\Delta V} = \sqrt{P_{\Delta V}}\bar{\Omega}_{\Delta V}, \quad \text{(A7)}$$

since the true probability of this stochastic evolution should be $\langle\psi(t)|\,E^\dagger E\,|\psi(t)\rangle$ for some effect E.

Now, solving for this effect for the QND measurement, we get

$$\begin{aligned}\Omega_{\Delta V} &= \left(\frac{4k}{\pi\Delta t}\right)^{\frac{1}{4}} e^{-iH\Delta t}e^{-\frac{2k\Delta V^2}{\Delta t}-2kX^2\Delta t+4kX\Delta V} \\ &= \left(\frac{4k}{\pi\Delta t}\right)^{\frac{1}{4}} e^{-iH\Delta t}e^{-2k\Delta t\left(\frac{\Delta V}{\Delta t}-X\right)^2} \end{aligned} \quad \text{(A8)}$$

This is the finite-time effect associated with the $\eta = 1$ weak QND measurement used in the main text (with $H = 0$).

For completeness we can also show that the effects defined above constitute a POVM since,

$$\begin{aligned}&\int_\infty^\infty d(\Delta V)\;\Omega^\dagger_{\Delta V}\Omega_{\Delta V} \\ &= \int_\infty^\infty d(\Delta V)\;P_{\Delta V}e^{-4kX^2\Delta t+8kX\Delta V} \\ &= \sqrt{\frac{4k}{\pi\Delta t}}e^{-4kX^2\Delta t}\int_\infty^\infty d(\Delta V)\;e^{-\frac{4k\Delta V^2}{\Delta t}+8kX\Delta V} \\ &= \sqrt{\frac{4k}{\pi\Delta t}}e^{-4kX^2\Delta t}\sqrt{\frac{\pi\Delta t}{4k}}e^{4kX^2\Delta t} = \mathbf{1}, \end{aligned}$$

where the integral on the third line is a Gaussian integral that can be evaluated by completing the square in the exponent.

[1] H. J. Carmichael, *An open systems approach to quantum optics* (1993).
[2] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, Nature **502**, 211 (2013).
[3] N. Roch, M. E. Schwartz, F. Motzoi, C. Macklin, R. Vijay, A. W. Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi, Phys. Rev. Lett. **112**, 170501 (2014).
[4] K. Jacobs, Phys. Rev. A **67** (2003).
[5] J. Combes and H. M. Wiseman, J. Phys. B: At. Mol. Opt. Phys. **44**, 154008 (2011).
[6] H. M. Wiseman and L. Bouten, Quantum Information

Processing **7**, 71 (2008).
[7] H. Li, A. Shabani, M. Sarovar, and B. K. Whaley, Phys. Rev. A **87**, 032334 (2013).
[8] H. M. Wiseman and J. F. Ralph, New J. Phys. **8**, 90 (2006).
[9] J. Combes and K. Jacobs, Phys. Rev. Lett. **96**, 010504 (2006).
[10] C. Ahn, A. C. Doherty, and A. J. Landahl, Phys. Rev. A **65**, 042301 (2002).
[11] M. Sarovar, C. Ahn, K. Jacobs, and G. J. Milburn, Phys. Rev. A **69**, 052324 (2004).
[12] A. Brańczyk, P. Mendonça, A. Gilchrist, A. Doherty, and S. Bartlett, Phys. Rev. A **75** (2007).
[13] H. M. Wiseman, Phys. Rev. Lett. **75**, 4587 (1995).
[14] M. Sarovar and K. B. Whaley, Phys. Rev. A **76**, 052316 (2007).
[15] J. E. Reiner, W. P. Smith, L. A. Orozco, H. M. Wiseman, and J. Gambetta, Phys. Rev. A **70**, 023819 (2004).
[16] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune, et al., Nature **477**, 73 (2011).
[17] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, and A. N. K. I. Siddiqi, Nature **490**, 77 (2012).
[18] D. Ristè, M. Dukalski, C. A. Watson, G. de Lange, M. J. Tiggelman, Y. M. Blanter, K. W. Lehnert, R. N. Schouten, and L. DiCarlo, Nature **502**, 350 (2013).
[19] M. Sarovar, H.-S. Goan, T. P. Spiller, and G. J. Milburn, Phys. Rev. A **72**, 062327 (2005).
[20] Z. Liu, Z. Liu, L. Kuang, L. Kuang, K. Hu, K. Hu, L. Xu, S. Wei, S. Wei, L. Guo, et al., Phys. Rev. A **82**, 032335 (2010).
[21] N. Nickerson, Y. Li, and S. Benjamin, Nature Commun. **4**, 1756 (2013).
[22] H. M. Wiseman, Phys. Rev. A **49**, 2133 (1994).
[23] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, 2009).
[24] F. Motzoi, K. B. Whaley, and M. Sarovar, arXiv.org (2015), 1502.XXXXX.
[25] K. Kraus, *States, Effects, and Operations: Fundamental Notions of Quantum Theory*, vol. 190 (1983).
[26] S. Hill and W. K. Wootters, Phys. Rev. Lett. **78**, 5022 (1997).
[27] K. Jacobs and A. Shabani, Contemp. Phys. **49**, 435 (2008).
[28] P. Goetsch and R. Graham, Phys. Rev. A **50**, 5242 (1994).
[29] H. M. Wiseman, Quantum Semiclass. Opt. **8**, 205 (1996).
[30] K. Jacobs and P. L. Knight, Phys. Rev. A **57**, 2301 (1998).