Supplemental Material for "Demonstrating a Continuous Set of Two-qubit Gates for Near-term Quantum Algorithms"

B. Foxen*,[1,2] C. Neill*,[2] A. Dunsworth*,[2] P. Roushan,[2] B. Chiaro,[1] A. Megrant,[2] J. Kelly,[2] Zijun Chen,[2] K. Satzinger,[2] R. Barends,[2] F. Arute,[2] K. Arya,[2] R. Babbush,[2] D. Bacon,[2] J.C. Bardin,[2,3] S. Boixo,[2] D. Buell,[2] B. Burkett,[2] Yu Chen,[2] R. Collins,[2] E. Farhi,[2] A. Fowler,[2] C. Gidney,[2] M. Giustina,[2] R. Graff,[2] M. Harrigan,[2] T. Huang,[2] S.V. Isakov,[2] E. Jeffrey,[2] Z. Jiang,[2] D. Kafri,[2] K. Kechedzhi,[2] P. Klimov,[2] A. Korotkov,[2] F. Kostritsa,[2] D. Landhuis,[2] E. Lucero,[2] J. McClean,[2] M. McEwen,[1] X. Mi,[2] M. Mohseni,[2] J.Y. Mutus,[2] O. Naaman,[2] M. Neeley,[2] M. Niu,[2] A. Petukhov,[2] C. Quintana,[2] N. Rubin,[2] D. Sank,[2] V. Smelyanskiy,[2] A. Vainsencher,[2] T.C. White,[2] Z. Yao,[2] P. Yeh,[2], A. Zalcman,[2] H. Neven,[2,b] and John M. Martinis[1,2]

[1] *Department of Physics, University of California, Santa Barbara, CA*
[2] *Google Research, Santa Barbara, CA 93117*
[3] *Department of Electrical and Computer Engineering, University of Massachusetts Amherst, Amherst, MA*
[b] *neven@google.com*
[*] *these authors contributed equally to this work.*

(Google AI Quantum)
(Dated: June 29, 2020)

I. FSIM CONTROL MODEL

A generic representation of a Fermionic Simulation (fSim) gate corresponding to a two-qubit photon conserving unitary requires five parameters. We may separate out the single and two-qubit parameters as follows: a $|01\rangle \leftrightarrow |10\rangle$ swap angle, θ, a $|11\rangle$ state conditional phase, ϕ, and three single qubit phases, Δ_+, Δ_-, and $\Delta_{-,off}$ yielding a generic fSim parameterization,

$$\text{fSim}(\theta, \phi, \Delta_+, \Delta_-, \Delta_{-,off}) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{i(\Delta_+ + \Delta_-)}\cos\theta & -ie^{i(\Delta_+ - \Delta_{-,off})}\sin\theta & 0 \\ 0 & -ie^{i(\Delta_+ + \Delta_{-,off})}\sin\theta & e^{i(\Delta_+ - \Delta_-)}\cos\theta & 0 \\ 0 & 0 & 0 & e^{i(2\Delta_+ + \phi)} \end{pmatrix} \quad (1)$$

We are interested in performing a two-qubit gate, which is independent of the single-qubit rotations. Therefore, we can focus on the matrix where Δ_+, Δ_-, and $\Delta_{-,off}$ are all zero, leading to the notation,

$$\text{fSim}(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -i\sin\theta & 0 \\ 0 & -i\sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix} \quad (2)$$

used to designate an arbitrary gate within the excitation preserving subspace.

II. FSIM GATE NUMERICS

The qubit dynamics presented in the main paper (Figure 2) are well described by numerics simulating two interacting qutrits (e.g. a pair of coupled three-level anharmonic oscillators) evolving with a time dependent detuning, $\Delta(t)$, and coupling, $g(t)$. We truncate the full two-qutrit Hamiltonian limiting our simulation to states with 1 or 2 excitations. Operating with the basis $|01\rangle, |10\rangle, |11\rangle, |20\rangle, |02\rangle$, the Hamiltonian describing the system is given by:

$$\text{H(g}, \Delta, \eta) = \begin{pmatrix} 0 & g & 0 & 0 & 0 \\ g & \Delta & 0 & 0 & 0 \\ 0 & 0 & \Delta & \sqrt{2}g & \sqrt{2}g \\ 0 & 0 & \sqrt{2}g & 2\Delta + \eta & 0 \\ 0 & 0 & \sqrt{2}g & 0 & \eta \end{pmatrix} \quad (3)$$

where η is the nonlinearity of each qubit, which we assume is the same for both qubits (240 MHz). Using this model, we may estimate the unitary operation enacted by arbitrary time-domain control of the coupling strength and the qubit detuning by discretizing these time domain control waveforms and performing a time ordered integral of H(t).

In Figure S1 we qualitatively reproduce the experimental results in Figure 2 of the main text by simulating 15 ns rectangular control pulses defining both g and Δ. In Figure S2 we illustrate the broadening effect that using shorter pulse lengths has on the Rabi interactions of both the $|01\rangle \leftrightarrow |10\rangle$ and $|11\rangle \leftrightarrow |02\rangle$ interactions by simulating rectangular pulses that are 10 ns, 15 ns, and 20 ns long. In Figures S2 and S3, we have omitted points where the leakage exceeds a 1% threshold which identifies the parameter space where we can perform fSim gates with low error. Experimentally we have chosen to implement our CPHASE gates with 13 ns long rectangular pulses with a 1 ns pad on either side—when we made the gate length shorter, leakage increased (data not shown). Here, in Figure S2a, we qualitatively see that the width of the 1% leakage band where we perform the CPHASE gate begins to pinch off and the $|2\rangle$ state Rabi interaction reaches all the way to the on-resonance iSWAP-like parameter space (dotted white line) when the gate

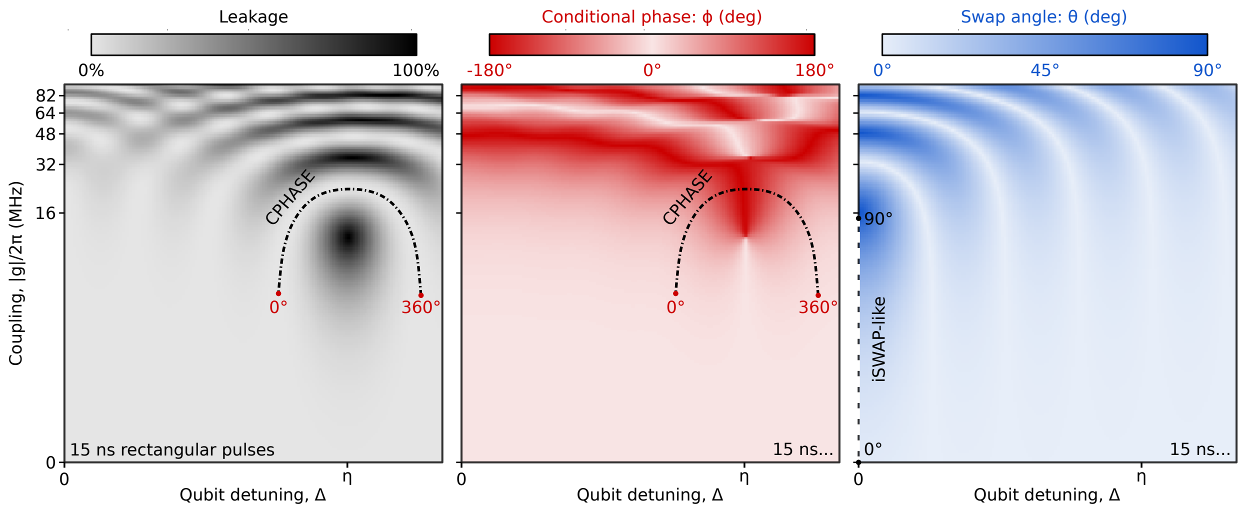

FIG. S1. Numeric simulation of two interacting qutrits reproducing the data from our experiments in Figure 2 of the main text. We simulate qubits with a fixed nonlinearity (240 MHz) with 15 ns long rectangular control pulses defining the qubit detuning, Δ, and their coupling, g.

length is 10 ns. Both these results qualitatively reproduce what we observed experimentally when attempting iSWAP-like gates shorter than 11 ns or the CPHASE gate shorter than 13 ns. Finally, in Figure S3 we simulate the effect of smoothing the control pulses by simulating 20 ns long coupler pulses that are rectangular, rectangular with 3 ns Gaussian smoothing, and cosine shaped (all detuning pulses are rectangular and have the same length). Here we see that smoothing reduces the extent of leakage from the second and third $|11\rangle \leftrightarrow |02\rangle$ swap lobes expanding the available low-error fSim control space. This indicates that pulse smoothing may be an important consideration of any future fSim implementation that aims to perform an arbitrary fSim using a single coupler pulse instead of the two discrete rectangular pulses we have used in this work.

III. GATE CHARACTERIZATION

We use a variety of techniques to characterize the performance of our single and two-qubit gates. In lieu of full process tomography, we use depth one population based measurements to perform unitary tomography to quickly assess the unitary operation performed by a given set of control pulses. We then turn to benchmarking techniques that amplify gate errors and allow for the characterization of small error rates. We use Clifford based benchmarking to characterize our single-qubit microwave gates and cross-entropy benchmarking (XEB) to characterize our two-qubit entangling gates.

A. Computing and reporting Pauli error rates

Before jumping in to gate characterization, a quick aside on Pauli error rates. We report Pauli error rates which are independent of the Hilbert space dimension and thus add linearly as the circuit's Hilbert space grows. In the past, many have reported average single and two-qubit depolarizing errors, e_d, as exponential decay constants of a sequence fidelity, $F = Ae^{me_d} + B$ where A and B are fit parameters to compensate for state preparation and measurement (SPAM) errors, m is the number of gate repetitions in the sequence, and e_d is the depolarizing error per cycle. The Pauli error, e_p, is related to e_d by the dimension of the Hilbert space:

$$e_p = e_d \times \left(1 + \frac{1}{D}\right) \tag{4}$$

where $\mathrm{D} = 2^{\mathrm{n}}$ is the dimension of the Hilbert space for an n-qubit gate. We note that this results in an increase in the reported error by a factor of 1.5 for single-qubit

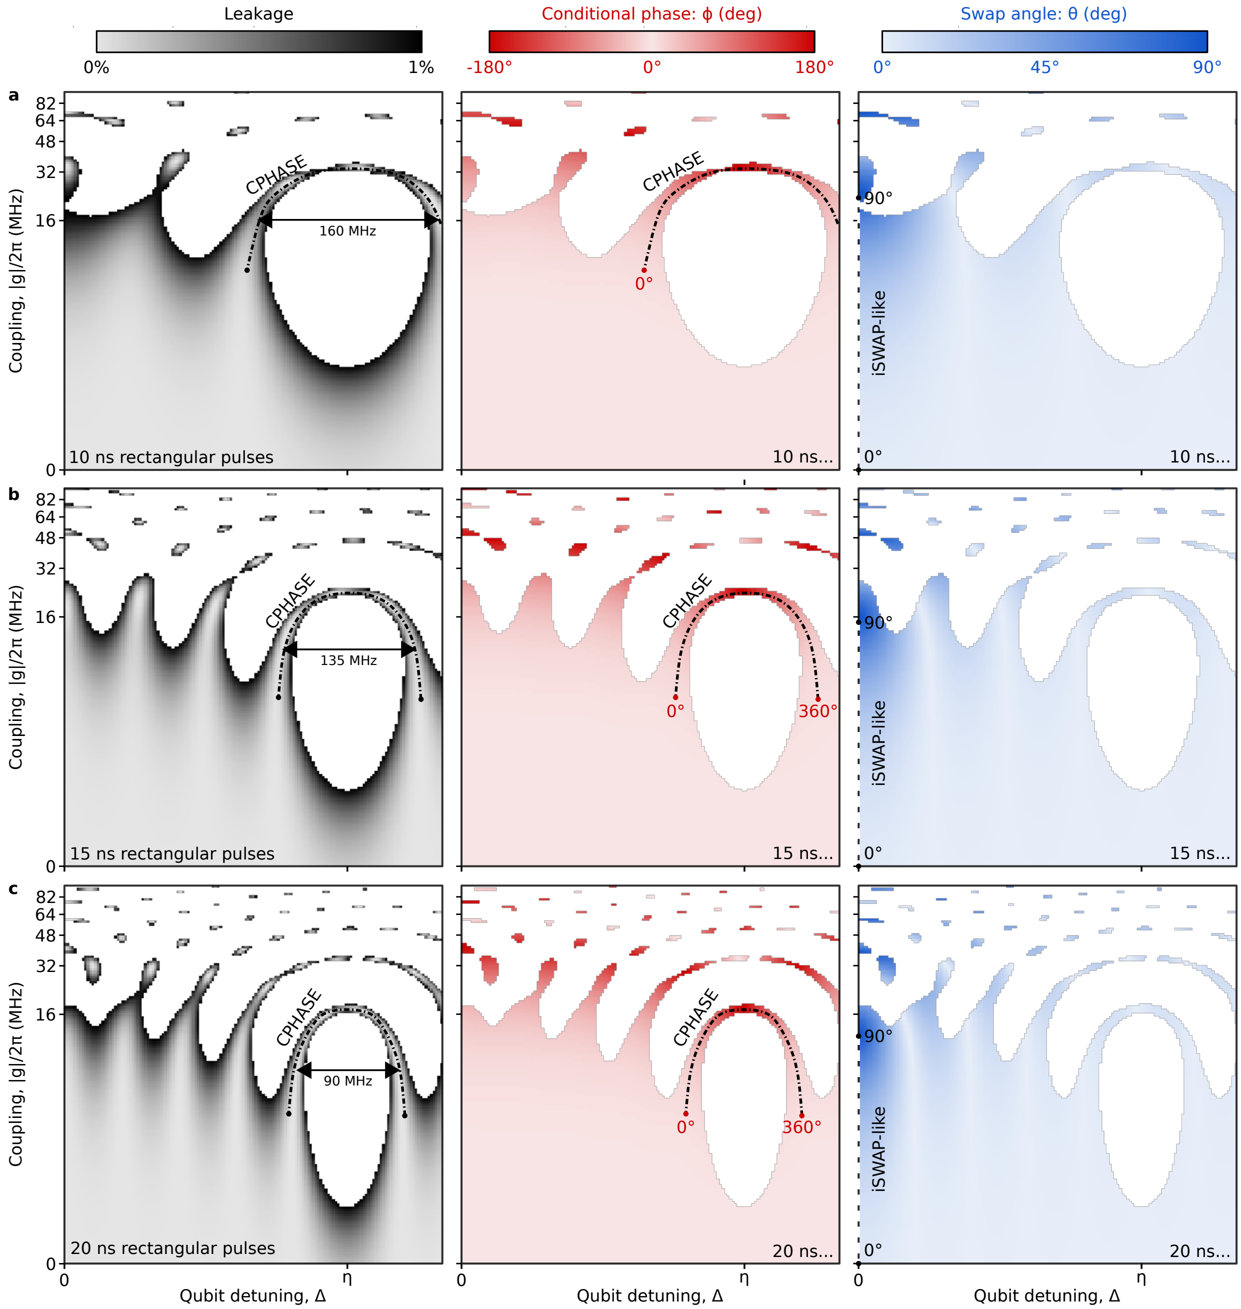

FIG. S2. Numeric simulation of **a**, 10 ns, **b**, 15 ns, and **c**, 20 ns rectangular control pulses showing the fSim parameter space where leakage is less than 1% (white regions are where leakage exceeded this threshold). Experimentally we chose to perform our CPHASE gate with 13 ns long pulses and the iSWAP-like gate with 11 ns control pulses (both of which had 1 ns pads on either side)—as we found that shorter implementations of either gate increased leakage and the overall gate error. Here, these numerics demonstrate that for 10 ns long gates, the low-leakage lobe where we perform the CPHASE gate narrows considerably and the $|2\rangle$ state Rabi interaction reaches the on-resonance iSWAP-like line cut near $\theta = 90°$, both of which agree with our experimental results.

gates ($n = 1$) and a by a factor of 1.25 for two-qubit gate errors ($n = 2$) [1].

When performing two-qubit XEB, we measure the exponential decay constant per cycle, $e_{r,\mathrm{cycle}}$ where each cycle consists of the application of one single-qubit gate per qubit and one fSim entangling gate involving both qubits. In order to extract the error per fSim gate, we can convert this to a Pauli error per cycle, $e_{p,\mathrm{cycle}}$, and subtract off the two single-qubit Pauli gate errors, e_{p,q_1} and e_{p,q_2}, which we estimate using single-qubit Clifford based randomized benchmarking performed on each qubit in isolation.

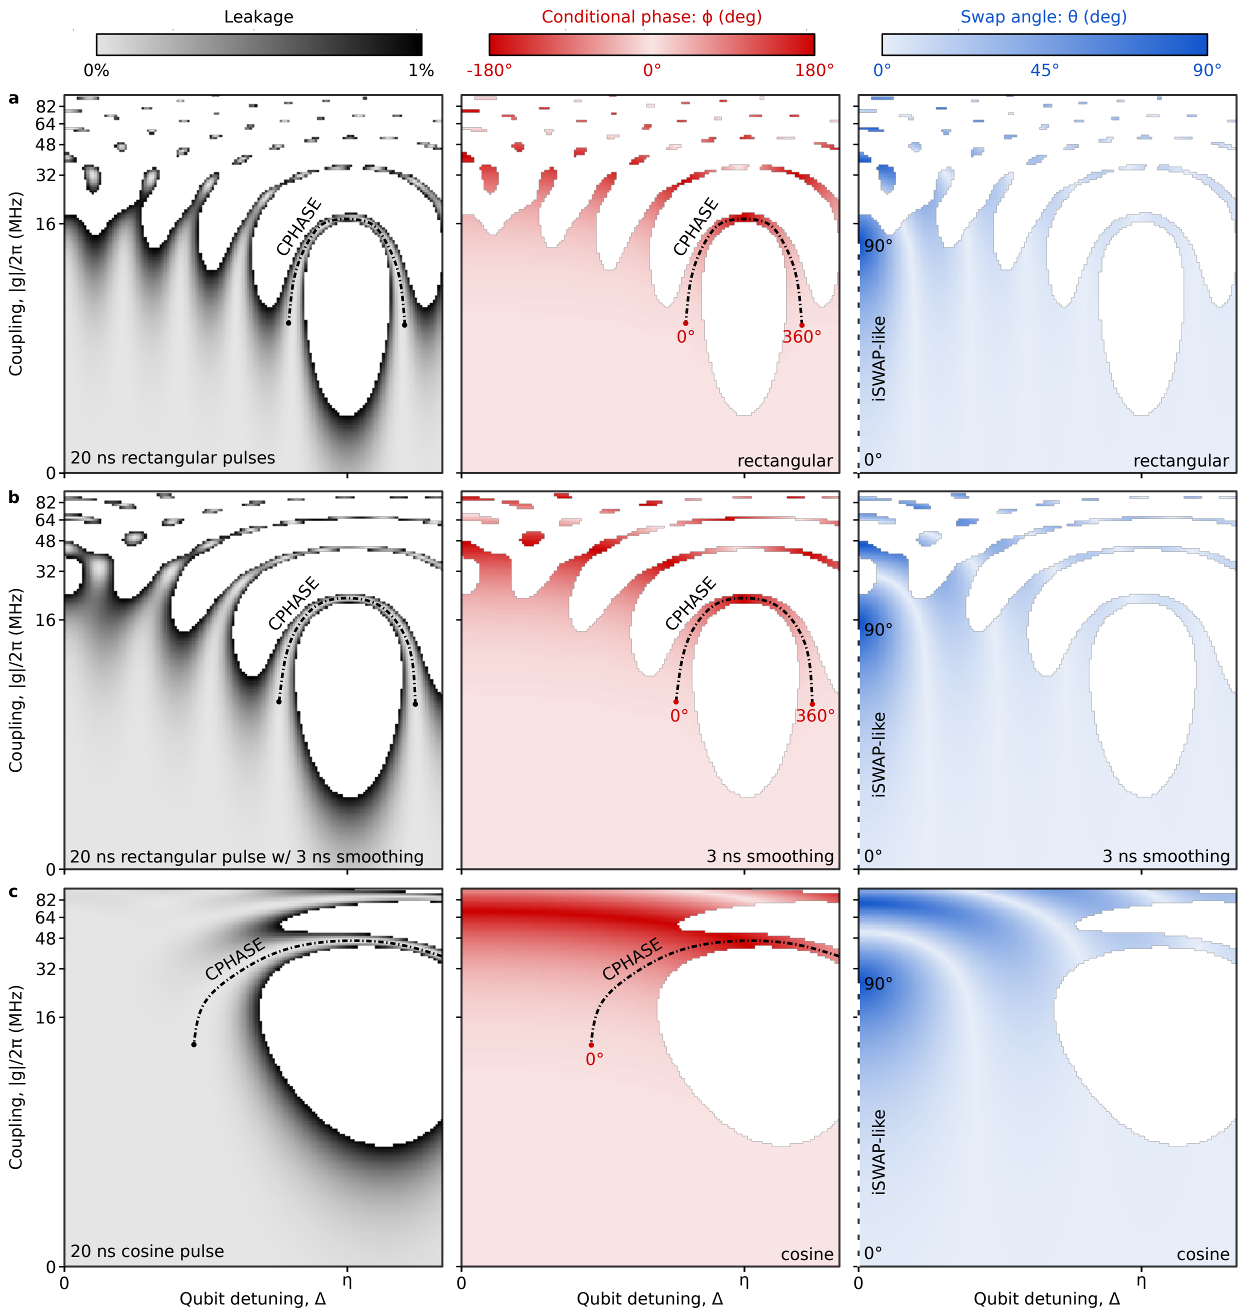

FIG. S3. Numeric simulation of **a**, a 20 ns rectangular coupler pulse, **b**, a 3 ns rise time rectangular pulse, and **c**, cosine coupler pulse showing the fSim parameter space where leakage is less than 1%. We observe that as the coupler pulses become more smooth, the fSim parameters space where leakage is less than 1% expands considerably. This indicated that pulse shaping and or smoothing may play an important role in any future implementation of the fSim gate set that aims to implement the gate set with a single pulse.

$$e_{p,2\mathrm{q}} = e_{p,\mathrm{cycle}} - (e_{p,\mathrm{q}_1} + e_{p,\mathrm{q}_2}) \tag{5}$$

For simplicity, all two-qubit Pauli errors have been computed assuming single-qubit Pauli errors of 7.5×10^{-4} per gate per qubit consistent with our typical single-qubit error rates immediately following a successful run of our standard single-qubit gate calibration procedure (see supplement III B).

$$e_{p,\mathrm{two\_qubit}} = e_{p,\mathrm{cycle}} - 2\cdot(7.5\times10^{-4}) = e_{p,\mathrm{cycle}} - 1.5\times10^{-3} \tag{6}$$

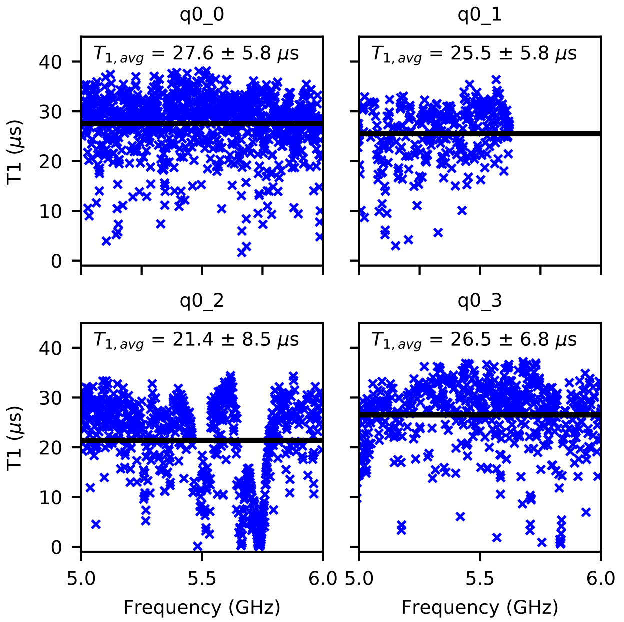

FIG. S4. Swap spectroscopy for four qubits from 5 to 6 GHz characterizing the qubit T_1 as a function of qubit frequency. For all four qubits on this chip over the available frequencies in the range of 5-6 GHz we find an average $T_1 = 25.3 \pm 7.3 \mu s$.

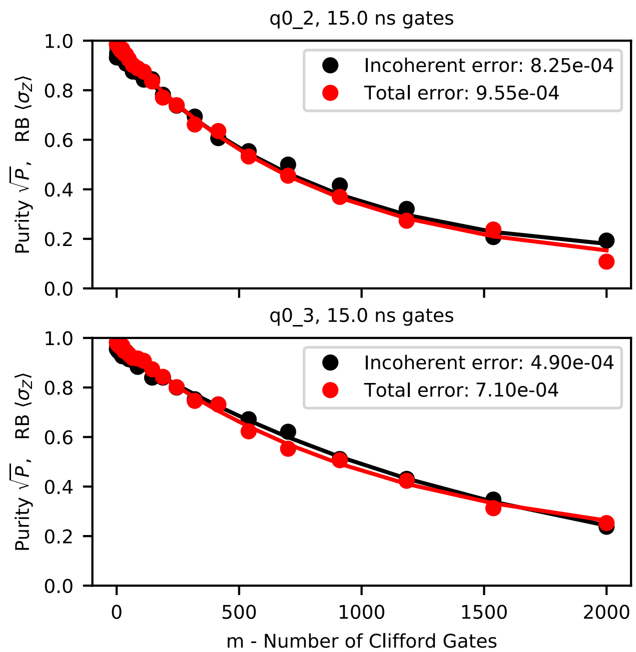

FIG. S5. Representative single-qubit Clifford-based randomized benchmarking results used to characterize the average error of our single-qubit gates. With a typical calibration, the single-qubit Pauli errors for both qubits are usually in the range of $5 - 10 \times 10^{-4}$. When computing the two-qubit gate error from the XEB per cycle error throughout this paper, we assume a moderately conservative error of 7.5×10^{-4} per single-qubit gate.

B. Single-qubit coherence and gates

Qubit coherence, in conjunction with gate duration, places a lower bound on both our single and two-qubit gate error rates. In Figure S4 we characterize T_1 for four qubits over a frequency range of 5 to 6 GHz. To perform this measurement we calibrate single-qubit gates, readout, and flux bias frequency control for a given qubit idle frequency. We then excite the qubit to the $|1\rangle$ state and detune the qubit to another frequency for a variable amount of time before detuning back to the idle frequency for readout. For each detuned frequency, T_1 is extracted as an exponential decay of the population over time, $P\,|1\rangle \propto Ae^{-t/T_1} + B$, where A and B are fit parameters to compensate for state preparation and measurement errors. In Figure **??** we plot representative data for four qubits on one processor over a frequency range from $5-6$ GHz. We find $T_1 = 25.3 \pm 7.3\,\mu s$ (one standard deviation) by averaging the available data from all four qubits ($f_{\max}$ for the second qubit was anomalously low so we include data for this qubit from $5 - 5.61$ GHz).

We use single-qubit purity [2] and Clifford-based randomized benchmarking [3, 4] to characterize the average error of our single-qubit gates. In Figure S5 we present representative results for a pair of qubits demonstrating purity-limited (incoherent error-limited) performance. These gate errors drift over time, but immediately following a successful run of our standard calibration procedure we typically observe single-qubit error rates at or slightly higher than the 7.5×10^{-4} level [5]. As such, we use this estimate in computing two-qubit error rates throughout this paper. These error rates are consistent with the coherence limit, for $T_{\text{gate}} = 15$ ns and $T_1 = 30\,\mu s$, giving $e_{p,\text{inc}} \approx 1.5 \times T_{\text{gate}}/3T_1) = 2.5 \times 10^{-4}$, with the remainder of the error coming from leakage and T_2 [6].

C. Unitary tomography

Section II of the main text describes shallow circuits used to characterize leakage and the two-qubit control parameters, θ and ϕ. Here, we detail the procedure used to directly measure all the non-zero matrix elements composing an arbitrary photon conserving unitary operation and the algebra used to convert these matrix elements into the five fSim control parameters (in Eq. 1). We use the resulting fSim model to compute the XEB sequence fidelity which we may then use as a cost function to optimize some, or all, of the fSim model parameters.

In order to efficiently characterize the unitary operation performed by a given set of control pulses, we initialize and measure a set of circuits as summarized in Table S1. If we consider a general photon conserving unitary the non-zero matrix elements will take the form:

TABLE S1. Summary of the two-qubit unitary tomography measurement sequences. Here, $\{u_{11}, u_{12}, u_{22}, u_{21}\}$ are the complex matrix elements of the two-qubit unitary in $[|01\rangle, |10\rangle]$ subspace. The two additional measurements ($u_{12,\text{excited}}$ and $u_{22,\text{excited}}$) are repeated measurements of u_{12} and u_{22} but with the other qubit placed into the excited state. This additional information is used to construct the conditional phase, ϕ.

Matrix element	Initial state	Measure qubit
u_{11}	(x, 0)	0
u_{12}	(0, x)	0
u_{22}	(0, x)	1
u_{21}	(x, 0)	1
$u_{21,\text{excited}}$	(1, x)	0
$u_{22,\text{excited}}$	(1, x)	1

$$\begin{array}{cc} & \begin{array}{cccc} |00\rangle & |01\rangle & |10\rangle & |11\rangle \end{array} \\ \mathrm{U} = & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & u_{11} & u_{12} & 0 \\ 0 & u_{21} & u_{22} & 0 \\ 0 & 0 & 0 & u_{33} \end{pmatrix} \begin{array}{c} |00\rangle \\ |01\rangle \\ |10\rangle \\ |11\rangle \end{array} \end{array} \tag{7}$$

Where u_{nm} denotes a non-zero element. We measured u_{nm} by initializing excited qubit in the basis ket of column m with an $X/2$ gate, and measuring the expectation value of $\sigma_x + i\sigma_y$ of the excited qubit in the basis ket denoted by row n. e.g. for u_{21} we initialize the left qubit, apply the fSim gate, and then measure $\sigma_x + i\sigma_y$ of the right qubit—this is the complex value of u_{21}. This procedure works for the single excitation subspace (e.g. n, m in $[1, 2]$), but u_{33} is computed from repeated measurements of $u_{12,\text{excited}}$ and $u_{22,\text{excited}}$ where the previously uninitialized qubit is instead placed into the $|1\rangle$ state as summarized in Table S1. This procedure is similar to process tomography, but requires considerably fewer measurements to characterize the fSim matrix. We note that an optimal measurement sequence would require only $2n$-1 circuits (for a $n \times n$ matrix) [7]. Even with several thousand repetitions of each circuit, characterizing the matrix with this method takes only a few seconds. Our series of six circuits is intentionally over-complete to avoid singular behavior when some matrix elements are small. In table S2 we list the conversion between the matrix elements and the five parameters of our fSim control model. These are useful measurements for building an fSim model, but we cannot characterize small gate errors ($\approx 10^{-3}$) using this method due to the limitations of state preparation and measurement (SPAM) errors which are a few percent.

D. Cross-entropy error benchmarking

Cross-entropy benchmarking (XEB) is a powerful technique for characterizing the error of an arbitrary gate [8]. It is particularly useful when implementing non-Clifford gates like the continuous fSim gate set we use here. XEB

TABLE S2. Computing fSim model parameters from the results of our unitary tomography protocol. The "condition" column is present because we compute $u_{33} = u_{22,\text{excited}}/u_{11}^*$ or $u_{33} = u_{12,\text{excited}}/u_{21}^*$ depending on if u_{11} or u_{21} is larger to ensure the result is non-singular. ψ_{10} is the phase difference accumulated between the two qubits over the gate duration.

fSim parameter	Value	condition
θ	$\arctan(\|u_{12}\|/\|u_{11}\|)$	none
ϕ	$\Delta_+ - \angle(u_{12,\text{excited}} \times u_{21})$	$\|u_{21}\| > \|u_{11}\|$
ϕ	$\angle(u_{22}) - \angle(u_{22,\text{excited}})$	$\|u_{21}\| < \|u_{11}\|$
Δ_+	$\angle(-u_{11} \times u_{21})$	$\|u_{21}\| > \|u_{11}\|$
Δ_+	$\angle(u_{11} \times u_{22})$	$\|u_{21}\| < \|u_{11}\|$
Δ_-	$2 \times \angle(u_{11}) - \Delta_+$	none
$\Delta_{-,\text{off}}$	$-2\left(\angle(-u_{12}/i) + \psi_{10}\right) + \Delta_+$	$\psi_{10} = (\omega_{q_1} - \omega_{q_0}) * t_{gate}$

uses a repetitive gate sequence to amplify small errors where each cycle consists of a random single-qubit gate from the set {X/2, Y/2, ±X/2±Y/2} applied to each qubit followed by the fSim gate we are benchmarking. We extract the error per cycle as an exponential decay in the XEB sequence fidelity, $\mathcal{F}_{\text{XEB}}$. The sequence fidelity is computed using the cross-entropy between two probability distributions P and Q, $S(P, Q) = -\sum_i p_i ln(q_i)$, by comparing the expected, measured, and incoherent probability distributions for a given gate sequence,

$$\mathcal{F}_{\text{XEB}} = \frac{S(P_{\text{incoherent}}, P_{\text{expected}}) - S(P_{\text{measured}}, P_{\text{expected}})}{S(P_{\text{incoherent}}, P_{\text{expected}}) - S(P_{\text{expected}})} \tag{8}$$

The numerator is the difference between the measured and expected cross-entropy and the denominator serves as a normalization so that $\mathcal{F}_{\text{XEB}}$ takes a value from [0, 1]. We then use $1 - \mathcal{F}_{\text{XEB}}$ as a cost function to optimize the five parameters of our fSim control model. For a given random sequence, we compute the expected probability distribution using perfect single-qubit gate models and the fSim model obtained from our unitary tomography experiment (supplement III C). Since, the sequence fidelity is dependent on the single and two-qubit gate models used in the cross-entropy calculation, we can use $1 - \mathcal{F}_{\text{XEB}}$ as a cost function to optimize some or all of our fSim gate model parameters, a process termed *ex situ* optimization.

E. RB vs XEB

As a sanity check, one may ask that we compare the result of Clifford based randomized benchmarking (RB) and cross-entropy benchmarking (XEB). Clifford based RB requires an inversion gate, inverting a random gate sequence to map the total ideal gate sequence starting in the $|0\rangle$ state back to $|0\rangle$. For most of the fSim gates, the inversion gate is non-trivial, but, for the special case of a $\text{CZ}_\phi = \text{fSim}(0°, 180°)$, which is part of the Clifford gate set, this comparison is possible.

In Figure S6a we perform single-qubit Clifford based randomized benchmarking (gate sequence inset), extracting average single-qubit Pauli errors $e_{p,q1} = 0.7 \times 10^{-3}$

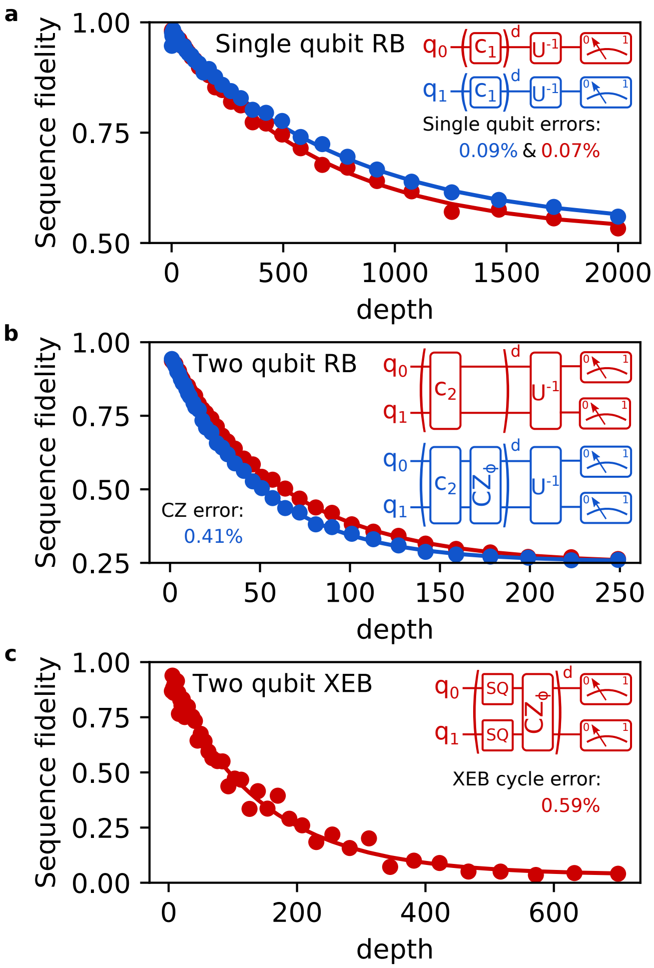

FIG. S6. Comparison of Clifford-based randomized benchmarking (RB) and cross-entropy benchmarking (XEB). **a**, Single-qubit Clifford based randomized benchmarking measuring average Pauli errors of 0.09% and 0.07% for each qubit. **b**, Two-qubit Clifford based randomized benchmarking with (blue) and without (red) an interleaved CZ_ϕ gate, allowing us to extract the Pauli error per $CZ + \phi$ of 0.41%. **c**, Two-qubit cross-entropy benchmarking where each cycle includes two single-qubit gates and a CZ_ϕ gate yielding a Pauli error per cycle of 0.59%. Here we find that the sum of the single and two-qubit errors measured with Clifford based RB (0.09% + 0.07% + 0.41% = 0.57%) corresponds well to the XEB error per cycle (0.59%).

and $e_{p,q2} = 0.9 \times 10^{-3}$. In Figure S6b we perform two-qubit Clifford based randomized benchmarking with and without an interleaved CZ_ϕ gate (sequences inset), extracting a Pauli error per CZ_ϕ of 4.1×10^{-3}. Then, in Figure S6c we use XEB to measure the per cycle error of the CZ_ϕ + two single-qubit gates obtaining $e_{p,\mathrm{cycle}} = 5.7 \times 10^{-3}$. If we then sum the Clifford based errors for each SQ gate and the CZ_ϕ $(0.7 + 0.9 + 4.1) \times 10^{-3} = 5.7 \times 10^{-3}$ we find good agreement with the XEB error per cycle $e_{p,\mathrm{cycle}} = 5.9 \times 10^{-3}$.

F. Error budgeting

In this section, we use various techniques to provide a more thorough budget of our XEB per cycle errors. As we have discussed, XEB measures the total error per cycle, $e_{p,\mathrm{cycle}}$. This includes coherent and incoherent errors for one single-qubit gate per qubit and one fSim gate. We use single-qubit Clifford-based randomized benchmarking to characterize the average total error for single-qubit gates, we use purity benchmarking to characterize incoherent error of both the single-qubit and fSim gates, and we use $|2\rangle$ state readout in conjunction with XEB to characterize per cycle leakage (which is included in the incoherent error). Here we focus on the two-qubit gate errors by assuming purity-limited single-qubit Pauli gate errors of 7.5×10^{-4} as described in supplement III B—this effectively means we subtract 1.5×10^{-3} from $e_{p,\mathrm{cycle}}$ to obtain $e_{p,\mathrm{2q}}$ for both error per gate and purity loss per gate measurements.

In Figure S7a we perform Purity benchmarking for each XEB gate sequence and obtain an average Purity loss of 3.76×10^{-3} per fSim gate. In Figure S7b we plot $e_{p,\mathrm{2q,unitary\_tomography}}$, the Pauli error per fSim gate using the fSim gate model obtained from unitary tomography. The average $e_{p,\mathrm{2q,unitary\_tomography}}$ is 5.07×10^{-3} indicating a coherent error of 1.31×10^{-3} per fSim. In Figure S7c we perform *ex situ* optimization of our fSim gate model to reduce the coherent error by changing the three single-qubit detuning model parameters. We hold the values of θ and ϕ fixed to the sampling grid, but allow the single-qubit phases in the fSim model to be optimized. With this improved gate model coherent error is nearly eliminated. The average error $e_{p,\mathrm{2q,ex\,situ}}$ is 3.83×10^{-3} reducing the average coherent error to 7×10^{-5} per gate.

We characterize leakage by directly measuring the $|2\rangle$ state population as a function of the XEB sequence depth. In Figure S8 we perform this measurement for a line cut of fSim control pulses that sweep the coupler bias on either side of the low-leakage bias used to perform a CPHASE gate. We find leakage to be minimized to a value of $5 - 6 \times 10^{-4}$ for a range of coupler biases spanning nearly 10 “clicks” of our 13-bit bipolar DAC $(2/2^{13} \approx 0.0002)$.

In total, these metrics indicate that we have achieved incoherent-error-limited gates with fairly low leakage (if necessary, leakage may be reduced further by optimizing the gate length at the potential expense of additional incoherent error). Additionally, we find that we are able to perform the desired fSim(θ, ϕ) gate we want without incurring additional coherent error. A critical component in achieving these results was eliminating the non-gate-like behaviors induced by long settling tails on our flux bias pulses. As such, we will now detail the procedure used to calibrate our flux control pulses.

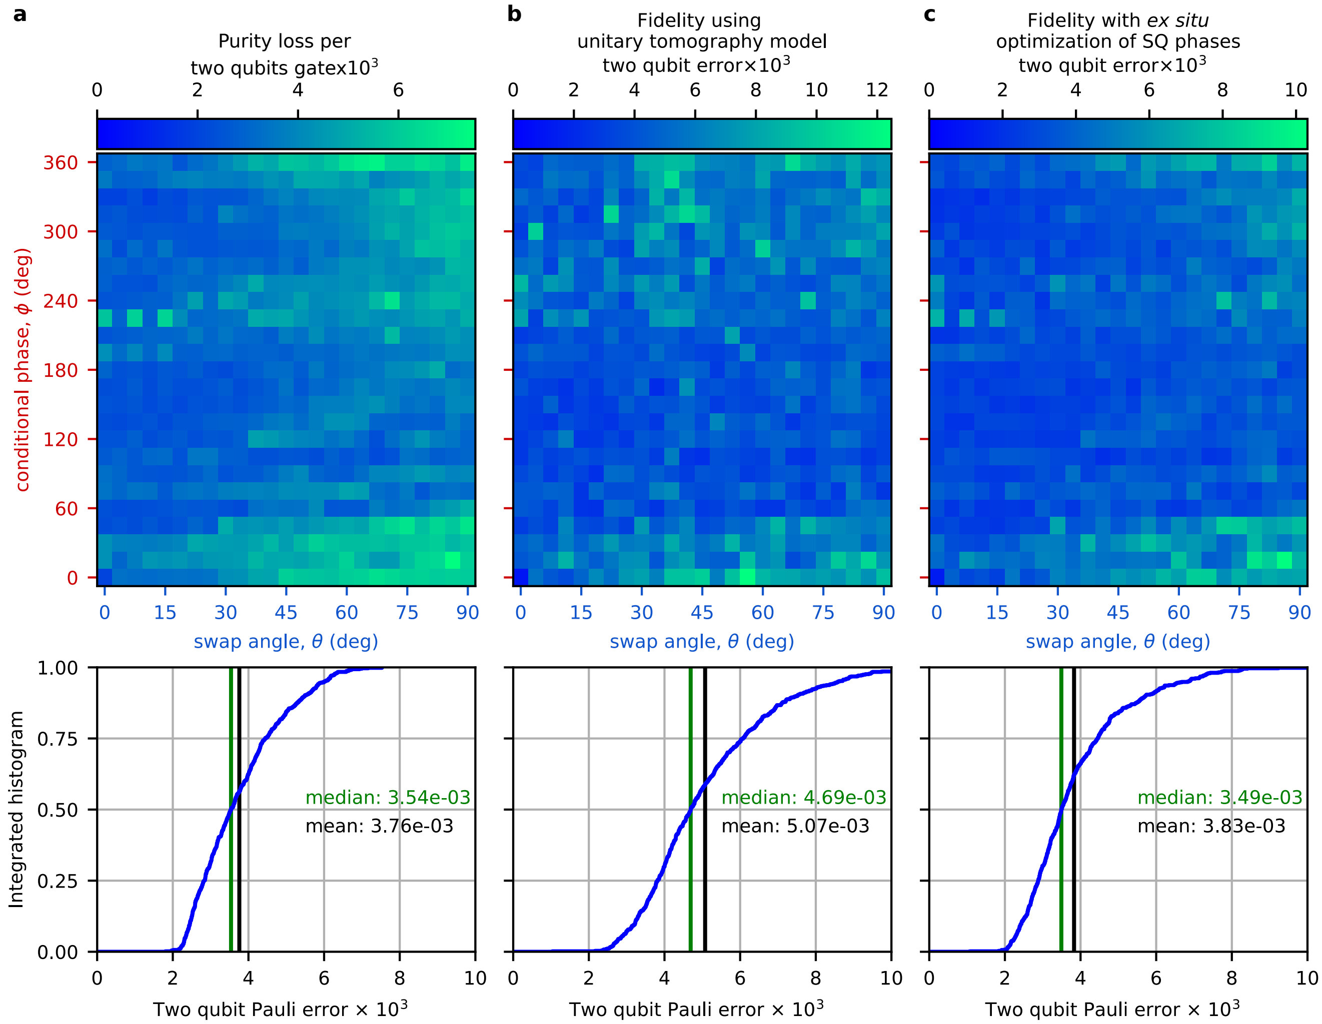

FIG. S7. Comparison of purity benchmarking and cross-entropy benchmarking with and without a constrained *ex situ* optimization of the fSim control angles. **a**, Purity loss per two qubit gate. **b**, XEB error per gate using the fSim gate model obtained from unitary tomography (supplement III C). **c**. XEB error after a constrained *ex situ* optimization of the fSim gate parameters where θ and ϕ were held fixed to the grid and the single-qubit phases were optimized.

G. Unitary overlap

The unitary overlap of two unitary matrices, e.g. some target fSim, U_{target}, and the actual fSim, U_{actual}, is defined as $Tr(U_{\text{target}} \cdot U_{\text{actual}})/D$, where D is the dimension of the Hilbert space. The unitary overlap is related to the Pauli error, $e_p = 1-(Tr(U_{\text{target}} \cdot U_{\text{actual}})/D)^2$. The Pauli error in an fSim gate for small deviations in either θ or ϕ is proportional to the square of the deviation angle. In Figure S9 we plot the additional coherent error incurred if you assume some actual $\text{fSim}_{\text{actual}} = \text{fSim}(\theta+\delta\theta, \phi+\delta\phi)$ is instead some target $\text{fSim}_{\text{target}} = \text{fSim}(\theta, \phi)$. This plot indicated that a deviation of either 2.5° in θ or 4° in ϕ with result in an additional coherent error of 1×10^{-3}. In our case (Figure S7), after a constrained optimization where θ and ϕ were fixed to a grid, our average error was approximately 1×10^{-4} higher than the purity limit which corresponds to a deviation of about 1° in either θ or ϕ.

IV. CONTROL PULSE CALIBRATION

In a world without flux settling tails, we would be able to implement an arbitrary fSim gate with a fidelity that is the sum of the requisite CPHASE and iSWAP-like gates by just merging the control pulses into a composite fSim gate pulse sequence. Unfortunately, due to flux settling tails, further calibration, described in IV C, was required. The keystones of this calibration were two-fold: 1) When performing two flux control based gates back to back (e.g. 2 ns separation), adjust the amplitude of the second pulse based on the first. 2)When implementing a composite gate, perform a CPHASE gate followed by the iSWAP-

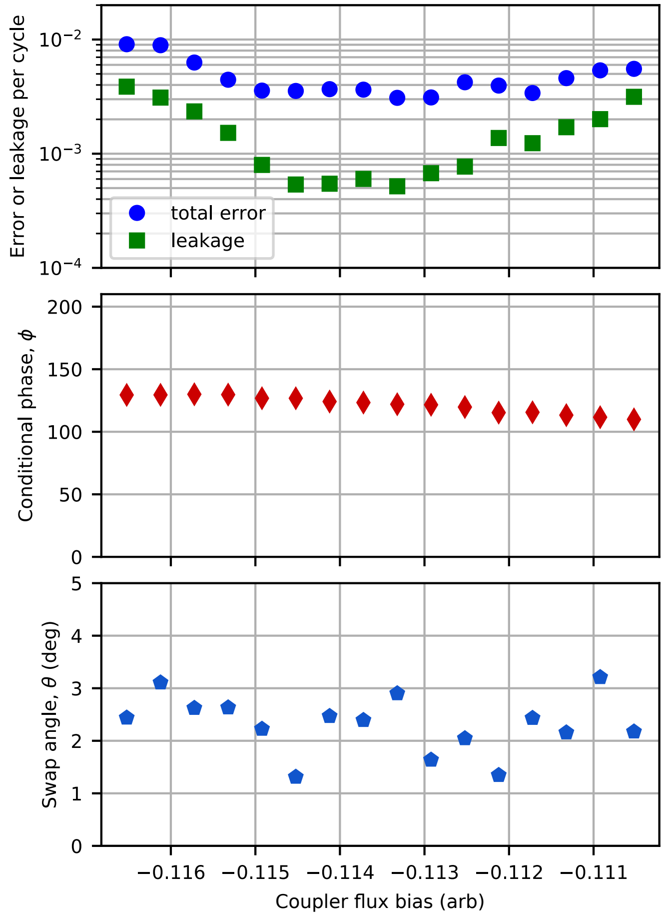

FIG. S8. Plot of leakage and total XEB error per cycle for a 13 ns CPHASE gate as a function of the coupler bias. The increment on the x-axis is four times the minimum increment of our DAC ($2/2^{14}$, e.g. a 14-bit bipolar DAC ≈ 0.0001). In this case we find that leakage reaches a minimum of $5-6\times 10^{-4}$ for a range of coupler amplitudes approximately $10x$ our minimum DAC adjustment.

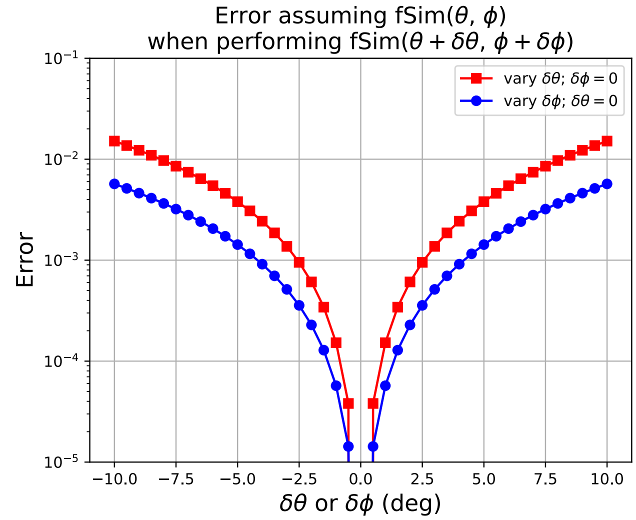

FIG. S9. We may choose to interpret some $\text{fSim}_{\text{actual}} = \text{fSim}(\theta+\delta\theta, \phi+\delta\phi)$ as some $\text{fSim}_{\text{target}} = \text{fSim}(\theta, \phi)$, by accepting additional coherent error. For small deviations in either θ or ϕ the error is proportional to the square of the deviation.

like gate so that bleed through is well behaved; in the reverse order, bleed through of the iSWAP-like coupler pulse into the CPHASE gate pulses will result in leakage to the $|2\rangle$ state which is an error in the fSim model. Using these two principles, we were able to implement a robust calibration of the complete fSim gate set.

As we have demonstrated numerically in supplement II, our desired implementation of the fSim gate set is possible with less than 1% error using simple rectangular control pulses. Unfortunately, the system transfer function (electronics and wiring) is imperfect and cannot produce these ideal waveforms exactly. Fortunately, as explored numerically in Figure S3, our fSim implementation is mostly sensitive to the integral of our control pulses rather than the shape. This likely remains true unless the spectral content of our flux control pulses approaches the qubit frequency. However, we must be very careful to ensure our control pulses do not bleed into each other which requires careful calibration of our flux bias settling tails.

We can consider settling non-idealities at two time scales: 1) pulse distortion during the duration of a gate (roughly 15 ns), and 2) pulse settling that occurs after the intended gate duration. Distortion at short times may, for instance, make it difficult to place the qubits exactly on resonance during a gate—this may make it difficult to achieve a swap angle, θ, of 90° swap amplitude (Rabi oscillation amplitude $= g^2/(g^2+\pi\Delta^2/2) = 1$ if and only if the qubits are on resonance), but fortunately these distortions do not have a huge impact on the rest of the fSim parameter space. Due to the periodic nature of Rabi oscillations the resulting fSim is mostly dependent on the integral of the control pulses. Pulse settling that occurs outside the intended gate interval means that adjacent gates will bleed in to each other. If the tails are relatively short (a few ns), it is possible to mitigate this error just by placing a short idle time between gates. Pulse settling at longer times is particularly nefarious because it becomes no longer feasible to pad gates with idle times and setting times of 5-1000+ ns have been observed in superconducting qubit systems [9, 10]. If left uncompensated, the performance of the m^{th} 15 ns long gate would be dependent on the preceding 1-60+ gates. This runs contrary to the entire notion of gate-based local operations and certainly would not fit within our static fSim control model used with XEB. As such, it is this long-time settling in particular that requires a careful calibration to enable the sensible control strategy employed throughout this letter.

The full fSim gate calibration happens in three stages. In the first stage, we calibrate the electronics to eliminate

TABLE S3. Summary of the settling parameters for two qubits. The average of the settling compensation for these two qubits was applied to the coupler.

	α_1 (%)	τ_1 (ns)	α_2 (%)	τ_2 (ns)	α_3 (%)	τ_3 (ns)
q_2	-0.46	858	-1.00	104	-4.94	10
q_3	-0.61	996	-0.82	94	-5.97	9
coupler (avg q_2 & q_3)	-0.53	927	-0.91	99	-5.45	10

the long-time settling flux settling. In the second stage, we describe the calibration procedure for the CPHASE and iSWAP-like gate sets. Then, for the fSim gate family, we perform further calibrations of the composite fSim gates to achieve the best possible gate performance by adjusting the control amplitude of the second pulse dependent on the first rather than adding longer buffer times between flux pulses.

A. Electronics calibration

On this device there are a total of seven flux bias lines, four for the qubits and three for the couplers. Each channel is driven by a dedicated 1 GS/s, 14-bit DAC controlled by an FPGA to form an arbitrary waveform generator. Each flux bias line uses nominally identical cabling, attenuation, and filtering from room temperature down to the sample's chip mount as shown in Figure S10. To compensate for non-idealities in each line, we first measure the qubit's response to a flux pulse, fit the response using three exponential decay time constants, and then use this model to pre-distort our control pulses as in previous work [10, 11]. This allows us to directly measure and compensate for the transfer function of each qubit's flux bias wiring. Implementing a similar *in situ* calibration of the coupler bias lines is the subject of on-going work. For now, we have found it sufficient to simply apply the average of the two adjacent qubit settling models to the coupler. The pulse calibration parameters for the pair of qubits and the coupler used to benchmark the fSim gate set are summarized in Table S3.

After performing the electronics calibration we find the unitary gate interactions of our fSim gates to be well characterized by either unitary tomography, performed with a depth-1 circuit, or cross-entropy bench-marking using a depth N circuit where N varies from 5 to 700. This fact is illustrated by Figure S7 panels b and c where

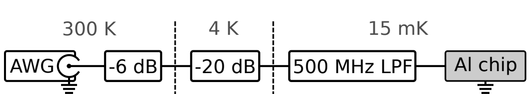

FIG. S10. Flux bias wiring: Each flux bias line was routed via a dedicated 50 Ohm coaxial cable from room temperature to the mK stage of the cryostat with 200 MHz Gaussian filters integrated with the arbitrary waveform generator (AWG), 6 dB of attenuation at 300K, 20 dB of attenuation at 4K, and a 500 MHz reflective low pass filter at 10 mk.

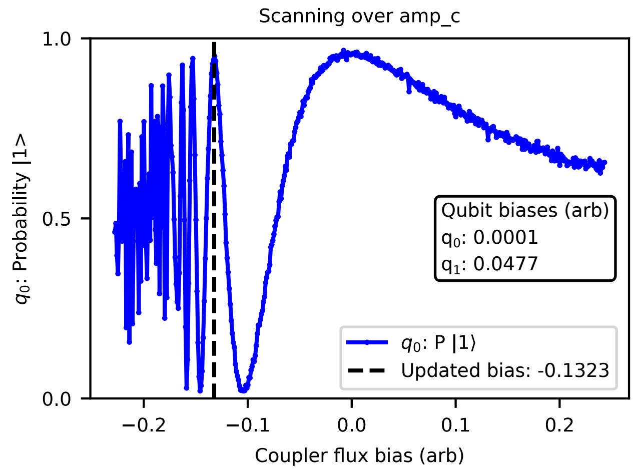

FIG. S11. CPHASE gate calibration scan. We use the pre-calibrated bias-to-qubit frequency function to choose a desired qubit-qubit detuning and then sweep the amplitude of the coupler flux bias to identify the amplitude that completes a diabatic $|11\rangle \leftrightarrow |02\rangle$ swap indicated by the dotted vertical line.

the difference in the average error of all 525 fSim gates differs by only 1.2×10^{-3} with and without optimizing the single-qubit unitary parameters—this provides an upper bound on the effects of pulse bleed through on gate fidelity. If we consider the gate timing of the cross-entropy benchmarking sequence in Figure S7, which used 28 ns fSim gates interleaved with 15 ns single-qubit gates, this result indicates that our settling is well compensated at times longer than 15 ns. This result also indicates that the qubit biases are settled enough to have a minimal impact on the single-qubit gate errors. If this were not the case then we would require a circuit-depth-dependent gate model to reach the purity limit. However, the settling of the coupler bias flux signal at times less than 15 ns becomes non-negligible and merits special consideration when calibrating fSim gates composed of a CPHASE gate in close proximity to an iSWAP-like gate. So, we will now detail the calibration procedure for each of the component gate families in the next section and finish our calibration discussion with a description of composite fSim calibration procedure.

B. CPHASE and iSWAP-like calibrations

We calibrate the CPHASE interaction by repeating the leakage experiment described in main Figure 2a to fine tune the coupler bias amplitudes and to identify combinations of qubit detunings, Δ, and corresponding coupler bias amplitudes that yield low-leakage gates. We use the qubit frequency bias transfer function to choose qubit biases that set the desired qubit detuning, Δ in the vicinity of η. The frequency range around η is set by the width of the Rabi interaction which, for a fixed pulse length, is inversely proportional to gate length since shorter gates require stronger coupling, g (see Figure S2 in supplement II). We use 15 ns pulses (13 ns rectangular pulses with a 1 ns padding on either side) which makes the Rabi interaction span about 75 MHz on either side of the qubit nonlinearity, η. For each detuning in this range, we repeat the experiment in main Figure 2a varying the coupling strength to minimize leakage. An example of the raw data from this experiment is provided in Figure S11 where the dotted line indicates the low-leakage coupler bias amplitude that achieves one full swap from $|11\rangle$ to $|02\rangle$ and back. We initialize the $|11\rangle$ state, apply the CPHASE control pulses, and measure the $|1\rangle$ state population of the lower frequency qubit to identify when the population has completed a full swap. Then, for each combination of Δ and the corresponding low-leakage coupler bias we repeat the experiment from main Figure 2b to measure the conditional phase. This procedure works well for $\phi : [-130°, 130°]$ until the Rabi swap amplitude becomes small and the peak becomes broad along the coupling strength line-cut. At that point, we extrapolate towards the zero coupling bias while measuring the conditional phase to fill out the rest of the conditional phase control space.

In Figure S12 we calibrate a 13 ns iSWAP-like gate (11 ns rectangular pulses, with 1 ns padding) by repeating the experiment from main Figure 2c three times with the qubits on resonance ($\Delta = 0$ MHz) to fine-tune the pulse amplitudes needed to reach $\theta = 90°$. Then, for $0° < \theta < 90°$ we simply interpolate the coupler bias between the "OFF" bias and the $\theta = 90°$ bias. For each iSWAP-like tune up experiment we initialize one qubit to the $|1\rangle$ state, apply the iSWAP-like pulses to the qubits and coupler, and then measure the $|1\rangle$ state population of the other qubit. In Figure S12a, we first use our pre-calibrated qubit frequency bias DC transfer functions to choose qubit flux bias amplitudes that place both qubits at the same frequency, and we sweep the coupler bias from the "OFF" bias to the maximum coupling bias to identify the amplitude that achieves exactly one a swap from the first to the second qubit corresponding to $\theta = 90°$ (dotted line). In Figure S12b, we repeat the experiment using the coupler bias from Figure 12a while sweeping the bias of one qubit to maximize the amplitude of the swapped population, thus placing the qubits on resonance. Finally, in S12c, we repeat the experiment using the new qubit flux biases to fine-tune the coupler flux bias.

C. fSim calibration

Once we have calibrated the iSWAP-like and CPHASE gates, we should nominally be able to use one of each to implement any fSim gate. Unfortunately our pulse response is imperfect at short times, as described in supplement IV A. This was less of an issue for the iSWAP-like and CPHASE gates in Section IV of the main text

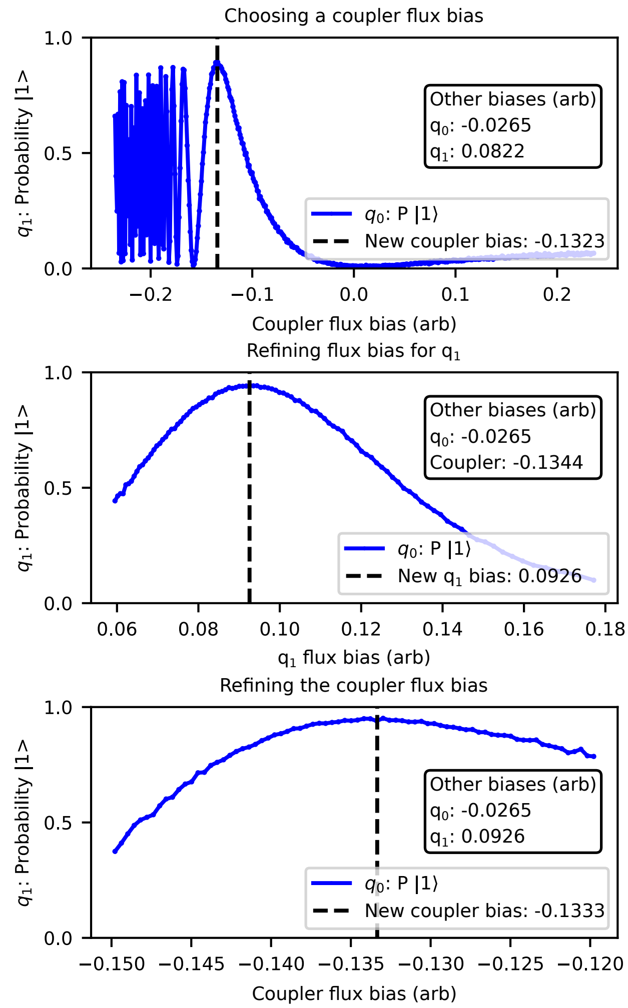

FIG. S12. iSWAP-like gate calibration for $\theta = 90°$ performed in three steps. In each experiment we initialize the $|01\rangle$ state and measure the population of the first qubit to identify the control biases to complete a full swap to $|10\rangle$ **a**, We use prior calibrations to bias the qubits on-resonance and scan the coupler bias amplitude to find the bias that completes one swap. **b**, Using the new coupler flux bias we scan the bias of one qubit to place them on resonance. **c**, Using the updated qubit biases, we again scan the coupler bias to find tune the coupler bias.

because the XEB gate sequence alternated (10 ns) single and (13 ns or 15 ns) two-qubit gates; in those cases, the uncompensated flux bias settling tails resulted in a small detuning at the qubit idle frequencies. However, when we perform an fSim gate as a composition of a CPHASE followed immediately by an iSWAP-like gate, the tail of the first coupler pulse bleeds into the second coupler pulse. Even a small settling tail adding to the amplitude of the coupler pulse can drastically change the coupling during the second gate due to the large coupler flux sensitivity at strong couplings (main Figure 1c). In the future, this problem may be mitigated by identifying and removing

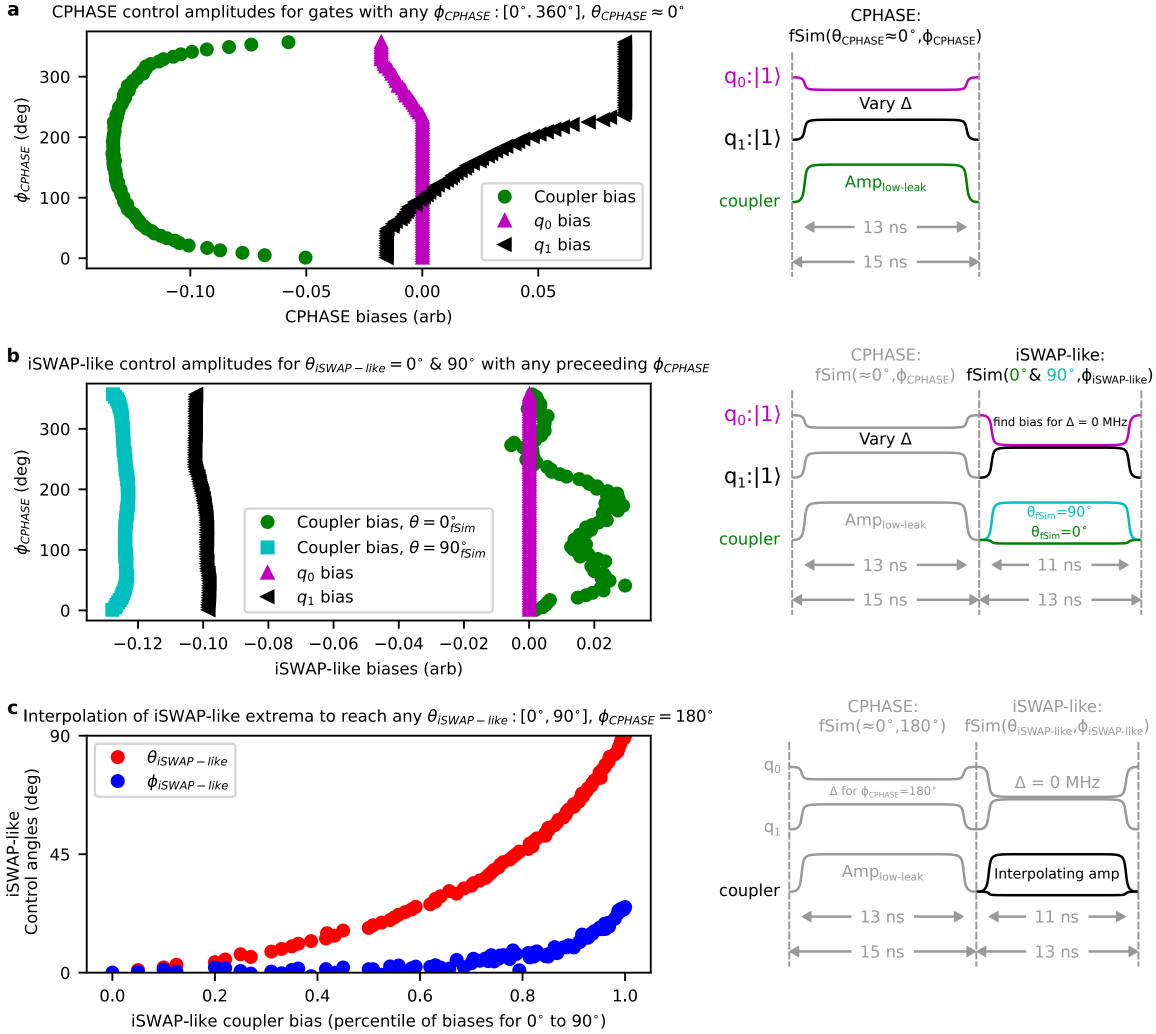

FIG. S13. Three steps to calibrating the fSim gate to account for pulse distortion due to the first (CPHASE) pulses bleeding into the second (iSWAP-like) pulses. **a**, Follow the usual CPHASE calibration procedure to bring-up a full CPHASE gate family corresponding to fSim($\theta \approx 0°, \phi : [0°, 360°]$). **b**, Follow the iSWAP-like tune up procedure, but play the CPHASE control pulses before the iSWAP-like pulses. Use the sequence to identify the flux bias amplitudes that achieve fSim gates with $\theta = 0°$ and $\theta = 90°$ for each proceeding CPHASE gate. **c**, for a preceding CPHASE gate with $\phi = 180°$, bring up gates corresponding to fSim($\theta : [0°, 90°], \phi = 180°$).

the physical origin of these settling tails, with a more thorough *in situ* calibration procedure for the couplers, or by placing longer idle times between gates.

In this work, we deal with pulse bleed through by calibrating composite fSim gates where the amplitudes of the second set of pulses in the composite fSim sequence is dependent on the first, thus eliminating the need for excessive idle times between gates. Conveniently, the tune up procedure for each gate in the composition is the same as in the isolated iSWAP-like or CPHASE case, just with the two sets of pulses played back-to-back—this works because each experiment in our usual bring-up procedure operates within an isolated manifold (e.g. one excitation for θ or two excitations for ϕ) when performing fSim gates. The ordering of the gates within the fSim gate is chosen to place the CPHASE gate before the iSWAP-like gate. Since both coupler pulses have the same sign, if the CPHASE coupler amplitude bleeds into the iSWAP-like coupler amplitude, this results in slightly more swapping which is easily measured and adjusted for by reducing the iSWAP-like coupler amplitude to compensate. If we ordered the gates in the reverse order, pulse bleed through would generate leakage during the CPHASE gate which is much more difficult to characterize and remove.

For the purpose of building a robust registry of gates, we erred on the side of over-calibration for this demonstration. However, we find these control parameters are well behaved and it should be possible to sample more sparsely in the future to simplify calibration of the full fSim gate set. Figure S13 outlines the three steps used to calibrate our composite fSim gates. In Figure S13a, we first calibrate many CPHASE gates spaced every 1°

using control pulses for just the CPHASE gate as shown on the right following the procedure outlined in supplement IV B. Then, in Figure S13b, for each preceding CPHASE gate we follow the iSWAP-like calibration procedure (also supplement IV B) to identify qubit and coupler bias amplitudes to achieve both a $\theta = 0°$ and $90°$ gate. Finally, in Figure S13c, for a CPHASE conditional phase, $\phi_{\mathrm{CPHASE}} = 180°$, we tune up iSWAP-like gates for θ from $0°$ to $90°$ in $1°$ increments by interpolating between the mininimum and maximum amplitudes determined in S13b. We use this calibration to produce a spline for $\theta_{\mathrm{iSWAP-like}} \rightarrow \%(bias_{90°} - bias_{0°})$ and another for $\theta_{\mathrm{iSWAP-like}} \rightarrow \phi_{\mathrm{iSWAP-like}}$.

With the fSim gate registry in hand we set out to benchmark specific fSim gates. For a given target fSim, we first look up the iSWAP-like pulse amplitudes that achieve the correct swap angle $\theta_{\mathrm{iSWAP-like}}$, and subtract the conditional phase due to the iSWAP-like gate from the total target to choose pulse amplitudes for the desired CPHASE gate (e.g. $\phi_{\mathrm{CPHASE}} = \phi_{\mathrm{target}} - \phi_{\mathrm{iSWAP-like}}$). We then performed unitary tomography (supplement III C) using the pulse amplitudes we looked up in the registry to quickly assess the resulting fSim control angles of the composite gate. If either control angle is off by more than $1°$, we used the registry to adjust the corresponding iSWAP-like (θ) or CPHASE (ϕ) control amplitudes by $\pm 1°$ accordingly. This process converged to an fSim gate within $1°$ of both θ_{target} and ϕ_{target} for the target fSim with fewer than 9 adjustments for each of the 525 fSim gates we benchmarked. Once the unitary tomography experiment indicated the composite fSim gate produced a unitary operation near the target fSim gate, we performed purity and cross-entropy benchmarking.

V. SYSTEM STABILITY

As the size of quantum processors grows (number of qubits), so too does the time it takes to calibrate a device (at least until fully parallel calibrations are possible). As the system drifts from these calibrations over time, the performance of a processor will fall and calibrations must be revisited. If the required calibration time is long compared to the scale of drift, then the device becomes unusable in practice. While electronics drift with both time and temperature must be considered when designing a system, one particularly worrisome issue is the time dependence of two level system (TLS) defects entering and leaving the qubit spectrum [12].

Here we present a promising snapshot of the stability of our system. In the process of calibrating the fSim gate set, we started by calibrating the single-qubit gates and readout. We then operated with the same single-qubit calibrations for several days while we were working on the fSim gate ultimately obtaining our primary fSim benchmarking dataset about a week after the initial single-qubit calibration. Shortly after this, a TLS showed up near one of the qubit's idle frequencies significantly limiting its coherence. Then, after about another week, we returned to the original calibration parameters to benchmark a subset of the same fSim(θ, ϕ) gates presented in the main text. We were pleasantly surprised to find that the original calibration was still good enough to produce high-fidelity gates.

In Figure S14 we used the two-week-old iSWAP-like and CPHASE calibrations to benchmark a less dense grid of fSim gates. While the average performance has degraded by a factor of two from the initial calibration, the average error is still less than 1%, but that is not the whole story. These fSim gates were benchmarked in a random order—if we look at a plot of the gate error as a function of time for these 91 (figure S14b), we see a strong time dependence where the first 50 gates (benchmarked over the course of an hour) have an average error much lower than gates #50 to #80. This would seem to indicated that the two-week-old electronics calibration is stable enough to maintain high-fidelity gates for weeks, and that the decreased fidelity is likely due to the residual and/or intermittent presence of a TLS interacting with one of the qubits. In an ideal world, we would be able to prevent or remove TLS defects, but, at least presently, we do not know how to do this. Instead, relying on the stability of our electronics, an optimal strategy for maintaining up-time on a large-scale quantum processor will likely involve calibrating a number of idle frequency configurations and being able to quickly vet and switch to an old configuration if and when a TLS shows up.

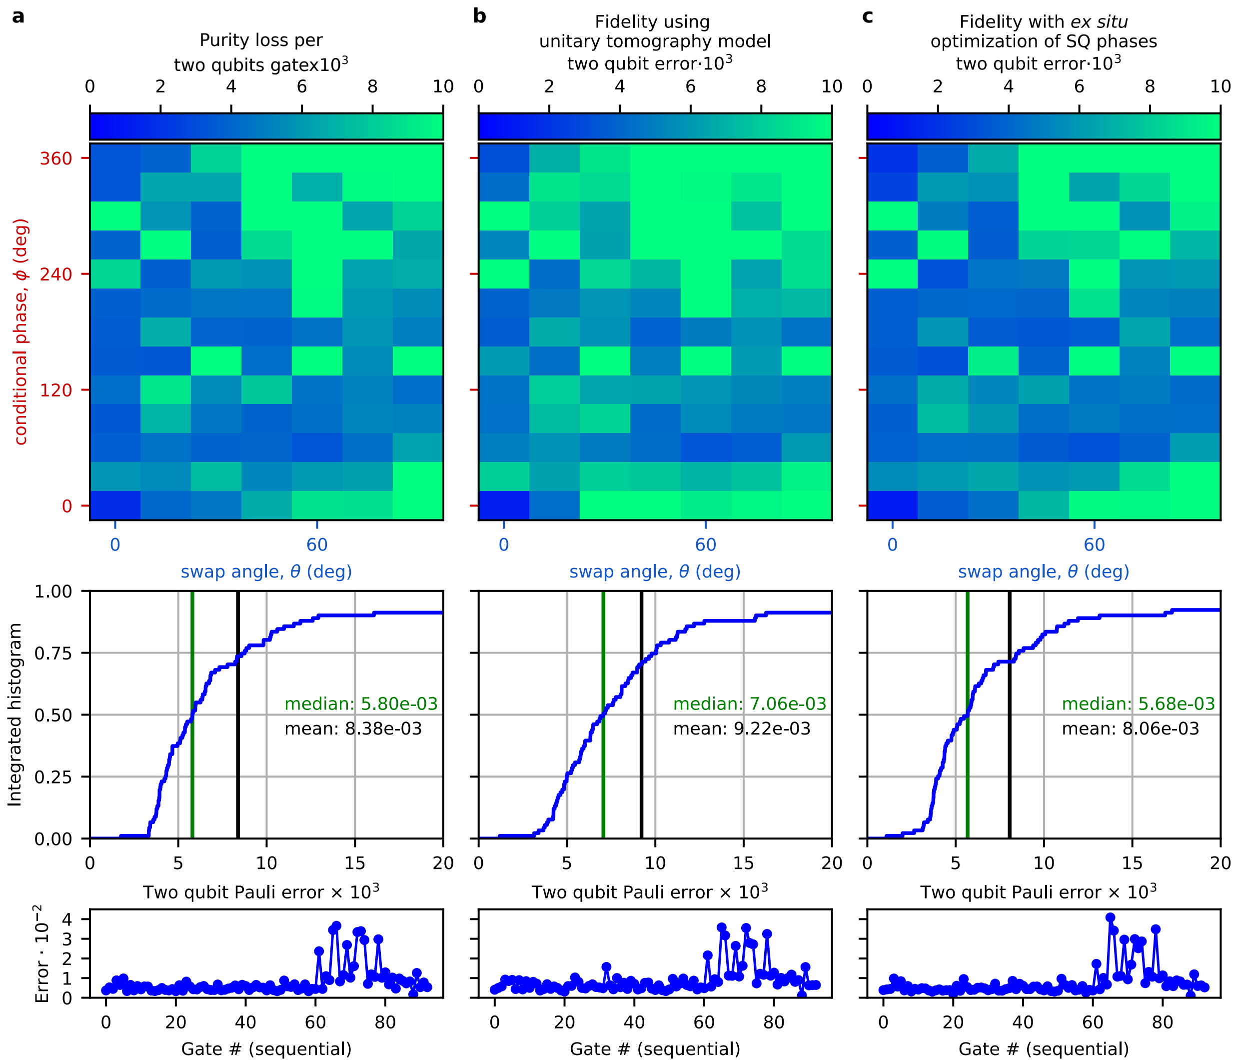

FIG. S14. Snapshot of system stability. A few days after taking the data in main Figure 4, a TLS showed up at one of the qubit idle frequencies effectively breaking the calibration. After about another week, we returned to the original calibration and repeated the fidelity measurement on a subset of 91 fSim gates which we present here (**top** and **middle** rows). We find that the average gate fidelity had decreased somewhat, but is still above 99%. Furthermore, if we look at the gate error rates sorted in the order they were measured (**bottom** row), a strong time-dependence becomes apparent. Many of the gates presenting low errors ($\approx 5 \times 10^{-3}$) as they did after the initial calibration. It is not until gates numbered 60 to 80 or so where large errors show up. This indicates that our control electronics are stable enough to maintain a high-fidelity calibration on the timescale of weeks, and that TLSs are likely the biggest threat to maintaining long term calibrations.

[1] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019.

[2] Joel Wallman, Chris Granade, Robin Harper, and Steven T Flammia. Estimating the coherence of noise. *New Journal of Physics*, 17(11):113020, nov 2015.

[3] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.*, 106:180504, May 2011.

[4] A. D. Córcoles, Jay M. Gambetta, Jerry M. Chow, John A. Smolin, Matthew Ware, Joel Strand, B. L. T. Plourde, and M. Steffen. Process verification of two-qubit quantum gates by randomized benchmarking. *Phys. Rev. A*, 87:030301, Mar 2013.

[5] Zijun Chen, Julian Kelly, Chris Quintana, R. Barends, B. Campbell, Yu Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Lucero, E. Jeffrey, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, and John M. Martinis. Measuring and suppressing quantum state leakage in a superconducting qubit. *Phys. Rev. Lett.*, 116:020501, Jan 2016.

[6] P. J. J. O'Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Qubit metrology of ultralow phase noise using randomized benchmarking. *Phys. Rev. Applied*, 3:044009, Apr 2015.

[7] Saleh Rahimi-Keshari, Matthew A. Broome, Robert Fickler, Alessandro Fedrizzi, Timothy C. Ralph, and Andrew G. White. Direct characterization of linear-optical networks. *Opt. Express*, 21(11):13450–13458, Jun 2013.

[8] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018.

[9] M. A. Rol, F. Battistel, F. K. Malinowski, C. C. Bultink, B. M. Tarasinski, R. Vollmer, N. Haider, N. Muthusubramanian, A. Bruno, B. M. Terhal, and L. DiCarlo. Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits. *Phys. Rev. Lett.*, 123:120502, Sep 2019.

[10] B Foxen, J Y Mutus, E Lucero, E Jeffrey, D Sank, R Barends, K Arya, B Burkett, Yu Chen, Zijun Chen, B Chiaro, A Dunsworth, A Fowler, C Gidney, M Giustina, R Graff, T Huang, J Kelly, P Klimov, A Megrant, O Naaman, M Neeley, C Neill, C Quintana, P Roushan, A Vainsencher, J Wenner, T C White, and John M Martinis. High speed flux sampling for tunable superconducting qubits with an embedded cryogenic transducer. *Superconductor Science and Technology*, 32(1):015012, dec 2018.

[11] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508:500 EP –, Apr 2014.

[12] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Yu Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, Erik Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, Daniel Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and John M. Martinis. Fluctuations of energy-relaxation times in superconducting qubits. *Phys. Rev. Lett.*, 121:090502, Aug 2018.