Editors' Suggestion

Fast microwave-driven two-qubit gates between fluxonium qubits with a transmon coupler

Siddharth Singh,[1,*,†] Eugene Y. Huang,[1,*,‡] Jinlun Hu,[1] Figen Yilmaz,[1] Martijn F. S. Zwanenburg,[1] Piranavan Kumaravadivel,[2] Siyu Wang,[1] Taryn V. Stefanski,[1,3,§] and Christian Kraglund Andersen[1,‖]

[1] *QuTech and Kavli Institute of Nanoscience, Delft University of Technology, 2628 CJ Delft, The Netherlands*
[2] *Netherlands Organization for Applied Scientific Research (TNO), Delft, The Netherlands*
[3] *Quantum Engineering Centre for Doctoral Training, H. H. Wills Physics Laboratory and Department of Electrical and Electronic Engineering, University of Bristol, BS8 1FD, Bristol, United Kingdom*

(Received 8 May 2025; revised 25 September 2025; accepted 28 October 2025; published 5 February 2026)

Two-qubit gates constitute fundamental building blocks in the realization of large-scale quantum devices. Using superconducting circuits, two-qubit gates have been implemented in various ways, with each method aiming to maximize gate fidelity. Another important goal of a new gate scheme is to minimize the complexity of gate calibration. In this work, we demonstrate a high-fidelity two-qubit gate between two fluxonium qubits, enabled by an intermediate capacitively coupled transmon. The coupling strengths between the qubits and the coupler are designed to minimize residual crosstalk while still allowing for fast gate operations. The gate is based on frequency selectively exciting the coupler using a microwave drive to complete a 2π rotation, conditional on the state of the fluxonium qubits. When successful, this drive scheme implements a conditional phase gate. Using analytically derived pulse shapes, we minimize unwanted excitations of the coupler and obtain gate errors of 10^{-2} for gate times below 60 ns. At longer durations, the gate performance is limited by relaxation of the coupler. Our results show how carefully designed control pulses can speed up frequency-selective entangling gates.

DOI: 10.1103/yxf3-jtx5

I. INTRODUCTION

Superconducting qubits have emerged as a leading candidate for building large quantum processors in recent years, demonstrating significant progress toward quantum simulations [1–8] and quantum error correction [9–12]. Within these experiments, single- and two-qubit gates are fundamental building blocks that dictate overall performance. The most widespread type of qubit used in these applications is the transmon qubit [13]. Two-qubit gates for transmon qubits were originally implemented with high fidelities using dynamic flux pulses to tune the qubits in and out of resonance [14–16] or by applying a strong microwave drive to activate an effective two-qubit interaction [17–21]. However, residual crosstalk often limits such architectures [22,23]. More recently, tunable couplers have consistently reached high gate fidelities, alleviating the crosstalk issue to a great extent [24–30].

Despite transmon-based gates reaching higher fidelities, transmons have a small anharmonicity that eventually limits the speed of gates and can lead to leakage out of the computational subspace [31–33]. The small anharmonicity further imposes strong requirements on qubit frequency targeting to avoid frequency crowding [34,35]. Additionally, dielectric losses are known to limit the coherence times of transmons [36,37].

The fluxonium qubit is one appealing alternative that aims to mitigate these shortcomings, as it is less prone to dielectric loss because of its lower qubit frequency, while also featuring a much larger anharmonicity [38, 39]. Coherence times in the millisecond range have been demonstrated in recent years [40–42]. Fast single-qubit and two-qubit gates have also been demonstrated with high fidelities [43–45]. One challenge in implementing fast two-qubit gates between fluxonium qubits arises from their small charge dipole moment, which suppresses the interaction between the computational states when they are directly capacitively coupled. Thus, direct capacitive coupling schemes for fluxonium qubits may require very large capacitors or very strong driving [45–47]; however, stronger coupling schemes are often accompanied by an increased residual interaction rate. A possible remedy is to

*These authors contributed equally.
†Contact author: Siddharth.Singh@tudelft.nl
‡Contact author: E.Y.Huang@tudelft.nl
§Present address: QphoX, 2628 XG, Delft, The Netherlands.
‖Contact author: C.K.Andersen@tudelft.nl

use inductive coupling [48,49]. However, inductive coupling may be accompanied by additional flux loops and require closer placement of qubits. Similarly to architectures for transmon qubits, an apparent solution is to introduce a tunable coupler to mediate strong interaction between the fluxonium qubits while simultaneously suppressing any residual coupling between the computational states [44,48,50]. In particular, two-qubit gates with high fidelities have been implemented by applying a microwave drive to the coupler [44,51]. While high fidelities have been achieved with this gate scheme, the maximum reported fidelity utilized a reinforcement learning algorithm for gate calibration [44], a potentially time costly procedure. Moreover, while reinforcement learning algorithms can adapt to unknown situations, these black-box algorithms often lead to calibrated pulse shapes that are hard to interpret. Furthermore, the uncertainty in the pulse shape makes it hard to understand the underlying physical mechanisms for drifts in the performance of the gate with time.

In this work, we implement a two-qubit gate between two fluxonium qubits coupled via a transmon coupler circuit; see Fig. 1(a). The gate is implemented by driving a coupler transition that is conditional on the state of the two fluxonium qubits; see Fig. 1(b). The transmon coupler interacts strongly with the higher levels of the fluxonium qubits, which leads to distinct coupler transition frequencies for each computational state. We expect that a large frequency selectivity among these transitions leads to faster two-qubit gates. On the other hand, for faster gates, the spectral overlap between the pulse and unwanted coupler transitions leads to coherent errors in the gate. Here, we employ simple analytical pulse-shaping techniques [52,53] to achieve fast two-qubit gates. Given the analytical form of these pulse shapes, we can interpret how they minimize coherent errors. Specifically, the pulses are designed to minimize the Fourier amplitude at the unwanted transitions. We experimentally obtain two-qubit gate fidelities of 99% for gate times as low as 60 ns.

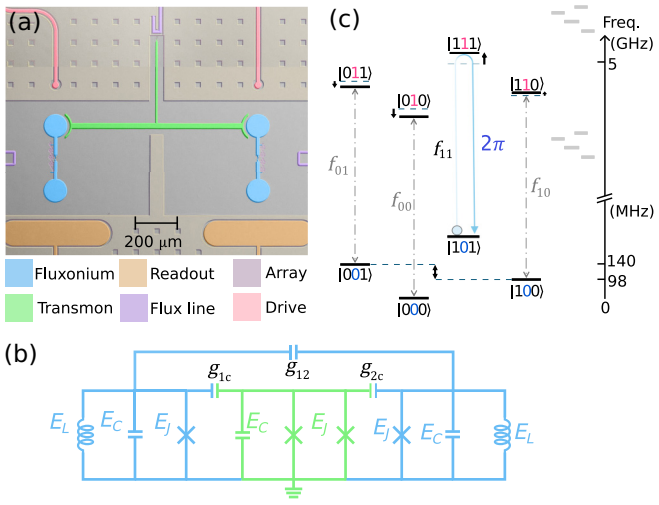

FIG. 1. (a) False-colored microscope image of the device, showing a magnified view of a single set of the combined qubit system. The fluxonium qubits are indicated in blue and are coupled by a transmon displayed in green. All the qubits are flux biased by currents applied to the flux lines (shown in purple), and the fluxonium qubits are driven capacitively by the drive lines (shown in pink). The readout resonators (shown in orange) have large capacitive elements near the qubits, and beyond this image, the resonators extend into $\lambda/2$ resonators. (b) Circuit representation of the two fluxonium qubits in blue and the transmon in green. Energies corresponding to each component of the qubit are indicated, along with the coupling between them. (c) Energy diagram of the combined system showing different levels. We use the notation $|F_1 C F_2\rangle$ to indicate the state of fluxonium 1, the coupler, and fluxonium 2, respectively. For clarity, we mark the system states corresponding to the coupler ground state in blue, constituting the computational subspace. The states with the coupler in the excited state are represented in red. We indicate the transition frequencies of the coupler with $f_{ij} = \omega_{ij}/(2\pi)$, conditioned on the fluxonium qubits' state $|ij\rangle$. Our gate scheme is operated by driving a 2π rotation on the coupler when the fluxonium qubits are in the $|11\rangle$ state, indicated with the blue array. The gray energy levels on the right-hand side indicate the higher excited fluxonium states. These states shift the dressed coupler levels (bold black) away from the bare coupler levels (dash-dot), allowing each coupler transition to be selectively addressed.

II. DEVICE DESCRIPTION

The device used in this work is composed of two fluxonium qubits; see Fig. 1(a). These two fluxonium qubits are both capacitively coupled to a transmon coupler circuit in a geometry that accommodates up to four couplers connected to each fluxonium qubit, as required, e.g., for implementing the surface code [54–57]. Throughout this work, we refer to the two fluxonium qubits as F_1 and F_2 and to the transmon coupler as C. The fluxonium qubits and the transmon couplers feature Josephson junctions fabricated using a Manhattan-style process; see further details about the fabrication in Appendix B and Ref. [58]. The fluxonium qubits include a flux loop, each composed of the central Josephson junction with Josephson energy E_J and a Josephson array consisting of 100 junctions in series with a combined inductive energy of E_L. The flux loops of the fluxonium qubits, as well as the symmetric superconducting quantum interference device (SQUID) loop of the transmon coupler, can be biased using their individual flux lines. Note that the coupler circuit is also coupled weakly to each drive line, such that the coupler can be readily driven by either. The fluxonium qubits are read out using dispersive readout enabled by their individual readout resonators. Similarly, the transmon coupler is also coupled to each readout resonator. The circuit layout of the three qubits is shown in Fig. 1(b). The parameters of the device are described further in Appendix A.

We write the total system Hamiltonian as

$$\hat{H} = \hat{H}_{F_1} + \hat{H}_{F_2} + \hat{H}_C + \hat{H}_I, \tag{1}$$

where each fluxonium qubit is described by their individual Hamiltonian

$$\hat{H}_{F_i} = 4E_{Ci}\hat{n}_{F_i}^2 + \tfrac{1}{2}E_{Li}\hat{\phi}_{F_i}^2 - E_{Ji}\cos(\hat{\phi}_{F_i} - \phi_{i,\text{ext}}), \tag{2}$$

where E_{Ci} is the charging energy, E_{Li} is the inductive energy, E_{Ji} is the Josephson energy of the ith fluxonium, and $\phi_{i,\text{ext}}$ is the external flux bias of each qubit. Similarly, the Hamiltonian for the transmon coupler is

$$\hat{H}_C = 4E_{C,C}\hat{n}_C^2 + E_{J,C}|\cos(\phi_{C,\text{ext}})|\,\cos(\hat{\phi}_C), \tag{3}$$

where $E_{C,C}$ is the charging energy of the coupler and $E_{J,C}$ is the total Josephson energy of the SQUID. The transmon is tuned by the external flux $\phi_{C,\text{ext}} = \pi\Phi_{C,ext}/\Phi_0$, where Φ_0 is the magnetic flux quantum. Since we have capacitive coupling between each fluxonium and the coupler, as well as directly between the two fluxonium qubits, the interaction Hamiltonian is

$$\hat{H}_I = \hbar g_{12}\hat{n}_{F_1}\hat{n}_{F_2} + \sum_{n=1,2}\hbar g_{ic}\hat{n}_{F_i}\hat{n}_C. \tag{4}$$

We describe the quantum state of this coupled system using the basis state $|F_1CF_2\rangle$, referring to the eigenstates of the coupled system unless otherwise noted. As indicated in Fig. 1(b), the higher excited states of each fluxonium qubit couple strongly with the transmon coupler, leading to level repulsions between the transmon levels and the fluxonium levels. In particular, the transition frequency of the coupler will depend on the state of the fluxonium qubits. We refer to the transition frequencies of the coupler as $\omega_{ij} = (E_{|i1j\rangle} - E_{|i0j\rangle})/\hbar$, where $E_{|x\rangle}$ is the eigenenergy of state $|x\rangle$. Since $\omega_{11} \neq \omega_{00}, \omega_{01}, \omega_{10}$, we can drive a selective 2π rotation on the ω_{11} transition such that state $|101\rangle$ picks up a phase of π, thereby implementing a conditional phase gate.

The specific fluxonium qubits that we study in this work are operated with qubit frequencies at 98 and 140 MHz for F_1 and F_2, respectively. The transmon coupler has a transition frequency around 4.7 GHz when operated at zero flux bias and with the fluxonium qubits in their ground state. With these qubit frequencies, the device is in a regime where the residual ZZ coupling is predicted to be below 0.5 kHz; see Fig. 2(a), where we extract the theoretical ZZ coupling based on a numerical diagonalization of the full Hamiltonian. The extracted values of the coupling constant are $g_{12}/2\pi = 40$ MHz and $g_{ic}/2\pi = 236$ MHz, compared to the designed values of $g_{12}/2\pi = 47$ MHz and $g_{ic}/2\pi = 236$ MHz. Here, we define the residual ZZ coupling as

$$\chi_{ZZ} = \omega_{101} - \omega_{100} - \omega_{001} + \omega_{000}. \tag{5}$$

To verify this residual coupling experimentally, we perform a modified Ramsey sequence to precisely measure the accumulated phase of F_2 with F_1 in either state $|0\rangle$ or $|1\rangle$; see also Appendix D. As seen in Fig. 2(b), we extract a residual ZZ coupling of around -20 Hz, well below the predicted value of 0.5 kHz in magnitude. Given the small value of measured ZZ, the calculation is highly susceptible to uncertainties in the model parameters. For example, the extracted transition frequencies from the fit have errors on the order of megahertz, corresponding to a relative error of around 0.1%, but still orders of magnitude larger than the residual ZZ coupling. Similarly, there may also be additional coupling terms that are not included in our model that may affect the coupling strength at the subkilohertz regime. Finally, it is well known (see, e.g., Refs. [58,59]) that a simple lumped capacitive model of the system is insufficient to predict higher-order dispersive features in superconducting qubit systems. However, the Hilbert space here is too large to perform highly accurate high-frequency simulations of the full device. Overall, we expect the residual ZZ coupling to cause little correlated crosstalk errors.

To further characterize the crosstalk between the two qubits, we employ simultaneous randomized benchmarking [60,61]; see Figs. 2(c) and 2(d). For these results, we decomposed the Clifford group into combinations of $\pi/2$ pulses and virtual-Z gates [62]. Note that the readout signal is slightly shifted for the simultaneous protocol due to a shift in the readout signal arising from the microwave reset scheme of the fluxonium qubits; see also Appendix F. We observe that the single-qubit gate fidelity is only slightly affected by simultaneous driving. Because of the small residual ZZ coupling, we attribute the small degradation in the single-qubit fidelities to microwave crosstalk between the fluxonium qubits.

III. TWO-QUBIT GATE IMPLEMENTATION

Our gate scheme involves driving the microwave transitions of the transmon coupler. To verify that we can selectively drive the coupler based on the state of the fluxonium qubits, we perform conditional spectroscopy of the coupler; see also Appendix E. As shown in Fig. 3, we see four distinct frequency transitions between 4.5 and 4.9 GHz for two different flux bias points. Because of fabrication uncertainty, the coupler frequency in the device lies close to the resonance frequency of the readout resonator of F_1 at 4.993 GHz. As a consequence, the coupler is partly hybridized with the readout resonator and we observe an overall degradation of qubit coherence.

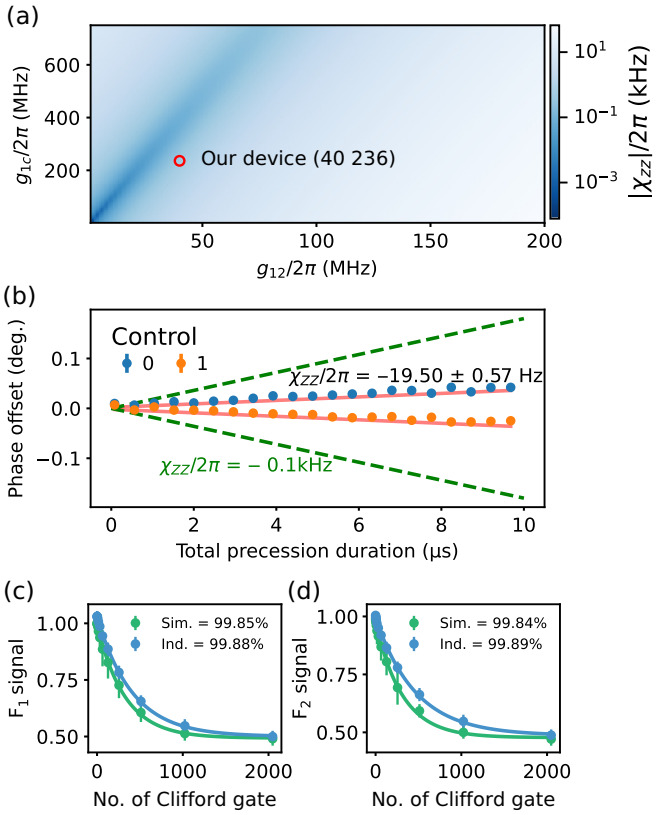

FIG. 2. (a) The residual ZZ rate for a range of coupling strengths between the two fluxonium qubits, g_{12}, and between each fluxonium and the coupler, g_{ic}. Parameter values corresponding to our device are indicated by a red circle. (b) Measured accumulated phase on F_2 dependent on the state of F_1. We extract the residual ZZ rate, χ_{ZZ}, from the slope of the signal. For reference, the theoretical value of $\chi_{ZZ} = 0.1$ kHz is plotted in green. (c),(d) Randomized benchmarking for single-qubit gates on fluxonium qubits F_1 and F_2, respectively, performed on each individual qubit (blue) and simultaneously on both qubits (green). The inset text displays the fitted average fidelity per gate.

To mitigate this detrimental effect, we use the frequency-flexible aspect of our design and flux bias the coupler to $0.13\Phi_0$ for the remainder of this work. At this operating point, the measured coupler lifetimes are of the order of 1.2 μs. Furthermore, at the chosen flux-bias point, four distinct frequency transitions of the coupler are observed; see Fig. 3. A detailed description regarding the lifetimes and frequency selectivity is given in Appendix A. We verify that we can coherently drive conditional Rabi oscillations on the coupler transition by applying a microwave pulse to the driveline of F_2, which also capacitively couples to the transmon coupler. As shown in Fig. 4, depending on which state the fluxonium qubits are prepared in, we observe a Rabi oscillation at a different frequency. Note that, while the transmon coupler does not have an individual readout resonator, we can combine the readout signals of F_1 and F_2 to monitor the coupler dynamics; see also Appendix E.

Importantly, we note that, compared to the other states, the fastest Rabi rate is observed when preparing the $|11\rangle$ state; thus, we expect the fastest possible gate to be achievable by driving this transition. Additionally, the coupler frequency corresponding to the prepared $|11\rangle$ state of the fluxonium qubits is most detuned from the $|1\rangle$–$|2\rangle$ transition of F_1, and the coupler has the longest lifetime in this state; see Appendix A. Specifically, we apply a drive, leading to the Hamiltonian

$$\hat{H}_d(t) = [\hbar\Omega_I(t)\cos(\omega_d t) - \hbar\Omega_Q(t)\sin(\omega_d t)]\hat{n}_{F_2}, \quad (6)$$

where $\Omega_I(t)$ and $\Omega_Q(t)$ are the pulse envelopes of the in-phase and quadrature components of the drive, respectively, and ω_d is the drive frequency. To achieve a fast high-fidelity gate, we must appropriately design the pulse shape of the drive as well as optimize the driving frequency. A black-box approach to finding the optimal driving parameters can be employed using the optimized randomized benchmarking for immediate tune-up protocol [63], or using reinforcement-learning algorithms [44]. However, the iterative feedback loop between the experiment and the classical optimizer may lead to a time-costly calibration procedure. Here, we instead use a straightforward calibration approach with no iteration steps, which allows for a fast and deterministic calibration procedure; see Appendix G. In our gate scheme we choose to drive the population from the $|101\rangle$ state to the $|111\rangle$ state and back again to the $|101\rangle$ state to obtain the conditional phase (CZ) gate. To ensure that we recover all the population back to the $|101\rangle$ state, we first fix the pulse shape, i.e., $\Omega_I(t)$ and $\Omega_Q(t)$, and then we measure the coupler response as a function of the drive frequency and the absolute pulse amplitude of the drive. Thus, for each pulse frequency, we deterministically obtain the pulse amplitude that maximally recovers the population back into the $|101\rangle$ state; see also Appendix G. Next, we measure the conditional phase from the application of 1, 3, or 5 gates. For a well-calibrated conditional phase gate, each of these should yield a conditional phase of exactly π. We measure the conditional phase for a range of frequencies and we pick the drive frequency where the conditional phase is equal to π. Finally, we measure the single-qubit phases of both F_1 and F_2 subject to repeated application of the gate. We correct for these single-qubit phases using virtual-Z gates [62].

Using this calibration routine, we first explore a simple pulse shape where the in-phase drive is a raised cosine

$$\Omega_I(t) = \frac{\Omega}{2}[1 - \cos(2\pi t/t_g)], \quad (7)$$

where t_g is the total duration of the gate. Additionally, we set $\Omega_Q(t) = 0$. Using interleaved randomized benchmarking [60,61], we characterize the gate fidelity of the

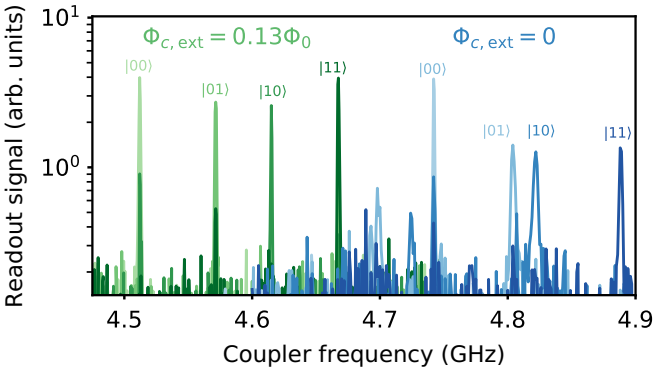

FIG. 3. Conditional spectroscopy of the coupler state when fluxoniums F_1, F_2 are prepared in different states $|F_1, F_2\rangle$, labeled above each peak. Two sets of peaks are observed for different flux operating points of the coupler, corresponding to $\Phi_{c,ext} = 0$ shown in shades of blue and $\Phi_{c,ext} = 0.13\Phi_0$ shown in shades of green.

CZ gate; see Fig. 5(a). For a cosine pulse, we find a maximal gate fidelity of $98.9\% \pm 0.1\%$ for a gate duration of 68 ns. For longer gate times, the fidelity becomes limited by the coupler coherence times; see the extracted CZ errors in Fig. 5(c). For gate durations below 60 ns, the gate error increases significantly, which we interpret as resulting from driving other spurious transitions, such as the transmon coupler frequencies associated with different fluxonium qubit states. A well-known technique to eliminate residual driving of weakly off-resonant transitions is the derivative removal by adiabatic gate (DRAG) technique [52]. Using DRAG, we can eliminate the spurious driving at a detuning of Δ by defining the quadrature drive as

$$\Omega_Q(t) = -\frac{1}{\Delta}\frac{d\Omega_I(t)}{dt}. \tag{8}$$

In our case the nearest transition that we are interested in suppressing is ω_{10} (see also Fig. 3). Thus, we can set $\Delta = \omega_{10} - \omega_d$ to obtain the pulse shape of the drive. Using this pulse shape, we see in Fig. 5(c) that we can obtain a gate fidelity around 99% for a larger range of gate times compared to the simpler cosine pulse shape. However, the DRAG scheme can only be used to suppress one of the three residual transitions. To further suppress the driving of residual transitions, we adopt the Fourier *ansatz* spectrum tuning derivative removal by adiabatic gate (FAST-DRAG) method, as introduced in Ref. [53]. Specifically, we use a pulse envelope given by

$$\Omega_I(t) = A\sum_n c_n\left(1 - \cos\frac{2\pi nt}{t_g}\right), \tag{9}$$

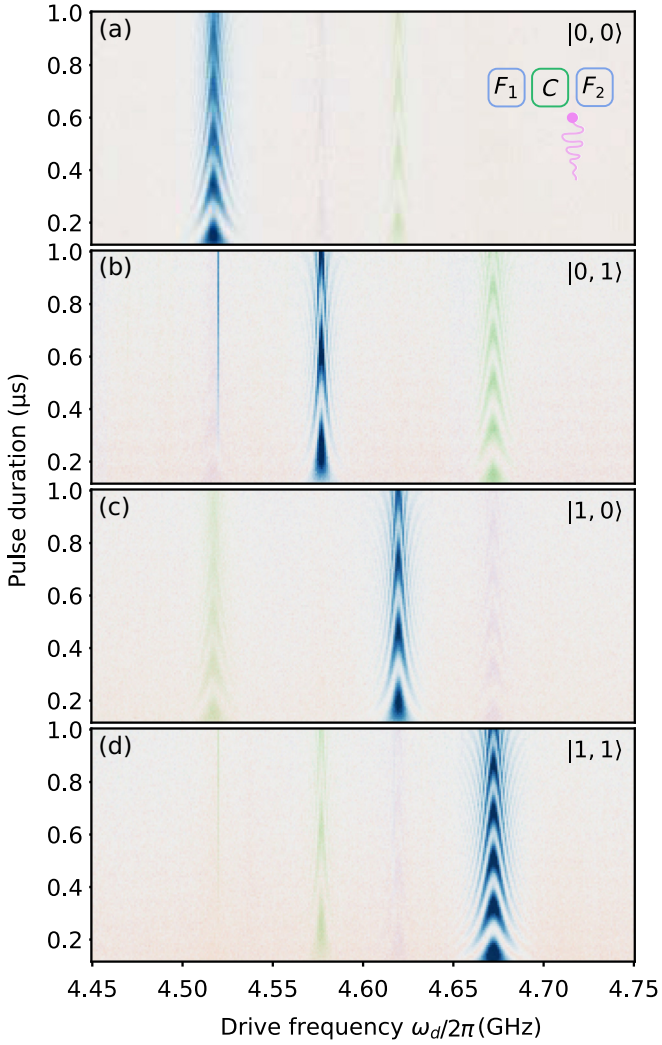

FIG. 4. Rabi oscillations when driving near the coupler frequencies are shown with different initialization of the fluxonium state, $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$ in panels (a),(b),(c), and (d), respectively. The readout signals of both fluxonium qubits are analyzed using principal component analysis, and we visualize the main components in blue-to-green and red-to-purple; see Appendix E. The coupler is driven through the drive line of F_2, as indicated in the inset of (a).

where t_g is the gate duration, A is the overall amplitude, and the c_n are the coefficients for the cosine series. Additionally, we add the quadrature drive following Eq. (8) with the $\Delta = \omega_{10} - \omega_d$, as before. As detailed in Appendix H, we engineer the Fourier spectrum of the pulse by solving a convex minimization problem with three cosine components in the series. In particular, we aim to minimize the Fourier components of the pulse in narrow frequency windows around f_{01}, f_{10}, and f_{00}. Examples of the resulting pulse shapes are displayed in Fig. 5(b), and the resulting average gate fidelities are shown in Fig. 5(c). We find a minimal gate error of $99.0\% \pm 0.1\%$ for a gate time of 64 ns with both the DRAG scheme and the FAST-DRAG

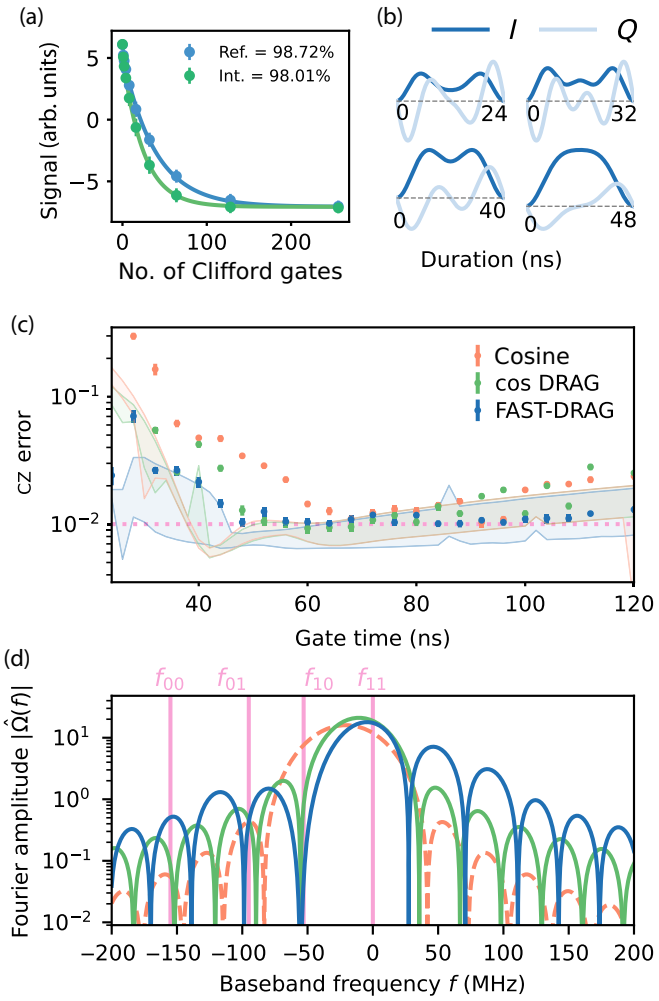

FIG. 5. (a) Clifford interleaved randomized benchmarking for a two-qubit CZ gate. The readout signal for the reference sequence is shown in blue, and in green we show the when interleaving the CZ gate. The extracted fidelities per Clifford gate are shown in the legend. (b) Pulse shapes used in the FAST-DRAG scheme. Normalized I and Q components are shown in dark blue and light blue, respectively. (c) Average CZ gate infidelity for the simple cosine pulse shape (orange), for the DRAG scheme (green), and for the FAST-DRAG pulse shape (blue). Error bars correspond to 1 standard deviation in the least-squares parameter fit. The horizontal dashed line indicates a fidelity of 99%. The shaded areas indicate the numerical simulations; see details in the main text. (d) Fourier spectrum for different pulse shapes is shown for a gate duration of 32 ns. Four vertical lines at the top correspond to the transmon coupler frequencies f_{00}, f_{01}, f_{10}, and f_{11}.

scheme. It is interesting to note that, while the overall lowest error is not improved compared to the DRAG scheme, the FAST-DRAG scheme provides good performance over a broader range of gate durations, showing that we more consistently accumulate fewer coherent errors from residual driving of spurious transitions. This is particularly evident for gates with durations shorter than 50 ns. We visualize the impact of Fourier engineering by plotting the resulting pulse spectrum; see Fig. 5(d). Generally, we find that the simple cosine pulse has large frequency components around all spurious transmon transitions. In contrast, the DRAG scheme minimizes the component at f_{10}, as expected, while the FAST-DRAG scheme broadens the window around f_{10} while also suppressing the frequency components around f_{01}. Given the over-constrained optimization problem, we find that the frequency component at f_{00} remains less suppressed.

To gain further insights into the errors, we simulate the gate numerically. We model the fluxonium qubits and the transmon coupler with the parameters discussed in Appendix A and implement a time-dependent drive using the qiskit dynamics package [64]. In the simulations, we include four energy levels for each fluxonium qubit and three levels for the transmon coupler. From the resulting unitary evolution, we truncate the evolution to the computational subspace of the two fluxonium qubits to obtain the matrix U from which we calculate the process fidelity as $F_p = |\mathrm{Tr}(U^\dagger_{\mathrm{ideal}} U)|^2/d^2$, where $d = 4$ is the Hilbert space dimension. For each type of pulse shape and duration, we numerically optimize the drive amplitude and the drive frequency in order to maximize the gate fidelity. In the simulations, we include a drive on fluxonium qubit F_2 and residual relative driving on F_1 and the transmon coupler. Specifically, we use a drive Hamiltonian $\hat{H}_D \propto \hat{n}_{F_2} + \eta_C \hat{n}_C + \eta_{F_1} \hat{n}_{F_1}$ to describe the relative capacitive crosstalk, η_i, from the driveline of F_2 to the transmon coupler C and fluxonium qubit F_1. We fix $\eta_{F_1} = 0.028$ based on capacitance matrix simulations of the device design. We find the simulated gate fidelities to be very sensitive to crosstalk, so to capture uncertainties in the precise crosstalk, we vary η_C between 0.2 and 1.0. The size of the full Hilbert space of two fluxonium qubits and the transmon coupler, including the higher excited states, prohibits efficient simulations as an open system. Thus, we include transmon T_1 in an ad hoc manner. Specifically, we calculate the process fidelity of each transmon transition as

$$F_{p,ij} = \frac{1}{2}\left(1 + \exp\left\{\int_0^{t_g} \gamma_{1,ij}\, p_{i1j}(t)\, dt\right\}\right), \tag{10}$$

where $p_{i1j}(t)$ is the probability to be in state $|i1j\rangle$ and $T_{1,ij} = 1/\gamma_{1,ij}$ is the lifetime of the coupler when the fluxonium qubits are prepared in state $|ij\rangle$. We calculate the total process fidelity as $F_{P,\mathrm{total}} = F_P \prod_{ij} F_{ij}$ and, finally, extract the average gate fidelity as $F_g = (dF_{P,\mathrm{total}} + 1)/(d + 1)$. As we increase η_C, the residual population in the transmon increases and, thus, the CZ errors increase as well. In Fig. 5(c), the shaded area indicates the span from the minimum to the maximal gate errors as we vary η_C; specifically, we observe that the cross-driving of the coupler can lead to an additional error of up to 1%. For pulse durations longer than 60 ns, we see that the overall behavior

of simulations matches well with the experimental data; thus, the coupler lifetime remains the main limiting factor for longer gate times. We also note that the FAST-DRAG scheme can potentially lead to smaller errors for longer gates since this drive scheme will cause less residual driving of the transmon and, thus, less potential errors due to the coupler lifetime. This effect is also partly visible in the experimental data. For gate times smaller than 50 ns, we expect a lower error from the simulations; thus, the experiment is clearly limited by residual driving of other transitions. From the simulations, we can identify residual driving of the 1–2 transition of F_1 through the transmon coupler as the main cause of coherent errors; see Fig. 8(a) below. Additionally, for gate times smaller then 50 ns, the FAST-DRAG scheme gives experimentally lower error compared to the conventional DRAG scheme. In fact, when there is large crosstalk to the transmon coupler, the simple cosine pulse is expected to outperform the FAST-DRAG scheme since the cosine pulse is overall more localized in frequency. Additionally, the simulations do not include crosstalk to the readout resonator detuned only 320 MHz from the drive frequency. For short gate times, residual driving of the readout resonator may lead to additional dephasing of the qubits.

Finally, the simulations indicate that the simple cosine and the DRAG scheme should perform similarly. However, in the experiment we clearly see an improvement from using DRAG. The reason is likely that the numerical optimization scheme simultaneously optimizes the amplitude and detuning of the gate, finding a good compromise where a larger detuning is fine-tuned to achieve the precise conditional phase. In the experiments, we optimize the population recovery independent of the conditional phase, as described in Appendix G. Thus, our simulations point to a clear direction for future devices, namely that we must minimize microwave crosstalk between the fluxonium qubits and improve the frequency targeting of the resonators and the coupler in a fluxonium-transmon-fluxonium architecture.

IV. CONCLUSION

In this work we have explored a system consisting of two fluxonium qubits coupled directly and through a transmon coupler circuit. The device parameters were designed such that the residual *ZZ* crosstalk between the two fluxonium qubits is suppressed. In this architecture, we can selectively drive the transmon coupler based on the fluxonium qubits' state. Thus, by driving a full 2π rotation conditioned on the fluxonium qubits being in the $|11\rangle$ state, we were able to implement a CZ gate. Using a simple cosine pulse shape for the drive, we obtained a fidelity of $98.9\% \pm 0.1\%$ with a gate duration of 68 ns. By employing pulse-shaping techniques, we optimized the performance further for shorter pulse durations. Specifically, we used the DRAG and FAST-DRAG techniques to obtain fidelities around $99.0\% \pm 0.1\%$ over a larger range of gate durations. These pulse-shaping techniques use simple analytical formulas to suppress the driving of residual coupler transitions, thereby minimizing the coherent errors from the drive.

ACKNOWLEDGMENTS

S.S., F.Y., and S.W. carried out the numerical simulations and calculations for the device energy parameters. S.S., F.Y., and P.K. fabricated the device. S.S., E.Y.H., J.H., M.F.S.Z., and T.V.S. characterized the device and performed the measurements for the experiments. S.S, E.Y.H., J.H., F.Y., M.F.S.Z., and T.V.S. contributed to the measurement setup. S.S. wrote the manuscript with feedback from all coauthors. C.K.A. supervised the execution of the project and the writing of the manuscript. The authors thankfully acknowledge support from the Kavli Nanolab Delft cleanroom staff members, including B. van Asten, E.J.M. Straver, E.F.D. Pot, B.P. van den Bulk, A. van Run, C.R. de Boer, M. Fischer, A.K. van Langen-Suurling, and M. Zuiddam. We also acknowledge fruitful discussions with A. Kamlapure, M. Finkel, H.M. Veen, and D.J. Thoen. The authors also acknowledge support from J. Bauer and L.J. Splitthoff on the junction development process. All authors acknowledge receiving support from the Dutch Research Council (NWO). E.Y.H. acknowledges funding support from Holland High Tech (TKI), J.H. acknowledges funding support from NWO Open Competition Science M, and T.V.S. acknowledges support from the Engineering and Physical Sciences Research Council (EPSRC) under EP/SO23607/1.

DATA AVAILABILITY

The data for the experiments are openly available [65]. Scripts to analyze the data are available from GitLab [66].

APPENDIX A: DEVICE MEASUREMENTS

In this appendix, we describe the device studied in this work and our approach to extract the device parameters. The device incorporates four pairs of fluxonium qubits, each with a transmon coupler, as shown in Fig. 6(a). In this work, we focus on a single set of fluxonium qubits. Qubit spectroscopy for one of the fluxonium qubits is shown in Fig. 7, where pulsed spectroscopy peaks (blue markers) were fitted to a single fluxonium spectrum model. This fitting gives us the individual fluxonium energy parameters listed in Table I. The transmon capacitive charging energy is estimated by its measured anharmonicity, and its Josephson energy is inferred from the measured coupler frequency. At the operational point, we measure the lifetimes and coherence times of the fluxonium qubits and

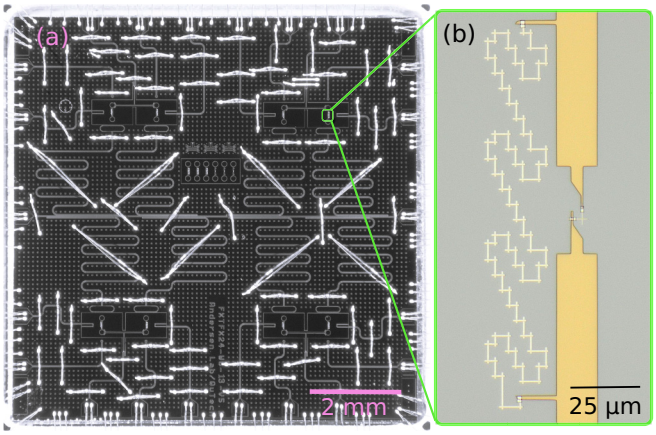

FIG. 6. (a) The device is shown with all four sets of the three-qubit system. The device features aluminum wirebonds on top, connecting ground planes across control lines and coplanar waveguide resonators. (b) A single fluxonium qubit is enlarged to show the array design and the single junction. The array consists of a total of 100 junctions, each with dimensions of 450 × 450 nm^2, and the single junction measures 150 × 150 nm^2.

the coupler using standard time-domain techniques; see Table I.

We use the scqubits package [67] to analyze the energy spectrum of the system in more detail. The energy spectrum is further used to calculate the coupling values (g_{ic}, g_{ij}). The fitting is done with the shift in coupler frequencies when the coupler is in either $|0\rangle$ or $|1\rangle$, corresponding to different initialized fluxonium states. In Fig. 8(a), we see how the transmon transition frequencies are positioned between the $|1\rangle$ to $|2\rangle$ transitions and the $|0\rangle$ to $|3\rangle$ transitions of each fluxonium. In particular, we note a strong avoided crossing between the $|0\rangle$ to $|1\rangle$ transition of C and the $|1\rangle$ to $|2\rangle$ transition of F_1 as we increase the external flux bias of the transmon coupler. Moreover, it is essential to ensure that the coupler is sufficiently detuned from the fluxonium resonators to avoid Purcell decay of the coupler transmon into the readout resonator. At an external flux of $\Phi/\Phi_0 = 0$, the transmon coupler transitions approach the readout resonator frequencies. Biasing the transmon coupler away from zero external flux reduces the Purcell decay rate, although the coupler becomes more sensitive to flux noise. At the external flux of $\Phi/\Phi_0 = 0.13$, the coupler is detuned by $\Delta/2\pi = 321$ MHz, along with a measured $\kappa/2\pi = 3.276$ MHz for the resonator of qubit F_1 and $g/2\pi = 17$ MHz for coupling between the coupler and resonator of qubit F_1, giving a Purcell limited lifetime of 17.3 μs

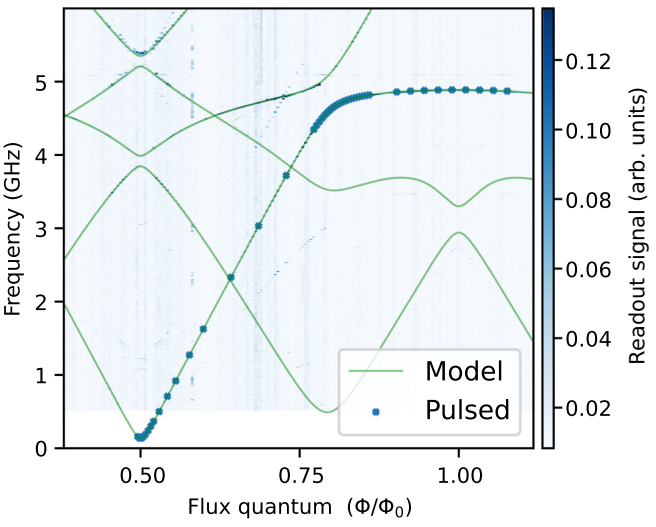

FIG. 7. Two-tone continuous-wave spectroscopy for fluxonium F_2 as a function of external flux and drive frequency is shown in the background, with spectroscopy peaks as blue crosses. A single fluxonium spectrum model is overlaid as solid green lines, with energy parameters fitted to the pulsed spectroscopy data.

Furthermore, as seen in Fig. 8(b), when we bias the transmon coupler away from zero external flux, we also increase the frequency selectivity of the desired transition frequency ω_{11} relative to ω_{00} and ω_{10}. The dispersive shift (χ) corresponding to the resonators of fluxonium qubits F_1 and F_2 is measured. For the readout of both fluxoniums, a readout pulse length of 2.24 μs was used.

APPENDIX B: FABRICATION

The device used in this work has been fabricated using a recipe similar to those described in Refs. [58,68]. We will describe further details here, as several improvements have been implemented in the process. The process begins with cleaning a high-resistivity silicon wafer (> 20 kΩ cm) with a (100) orientation, sourced from Topsil. The wafer is first immersed in a nitric acid bath for 7 min, accompanied by sonication, to remove organic residues from the surface. Next, it is sequentially rinsed in two deionized water baths for 60 s each before being dried with a nitrogen gun. The wafer undergoes an oxide removal step using a 40% HF solution for 6 min, followed by another deionized water rinse as described above. To further passivate the surface and delay oxide formation, the wafer is coated with HMDS vapor [69]. Since this passivation has a limited effective duration, the wafer is immediately loaded into a sputter system. For deposition of NbTiN, we utilized a sputtering system with a confocal target geometry. The deposition is performed using an NbTi target with a chemical composition of Nb:Ti (70:30). We deposit 200 nm of NbTiN using a deposition rate of 28.3 nm/s. To further process the wafer, we dice it into smaller pieces of 18 × 18 mm^2. To remove the protective resist from the dicing step, the device is put in NMP for 2 h at 80 °C and additionally left in NMP at room temperature over night. To remove potential resist residues, we follow the overnight cleaning step with IPA

TABLE I. Measured device parameters, with F_1 and F_2 representing the fluxoniums and C the coupler. The transition frequency and lifetimes of the coupler are measured with the fluxonium states prepared in $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$, written in this order. Here T_2^E of the coupler for the $|00\rangle$ fluxonium state could not be measured, marked as "–".

Measured	F_1	C	F_2
T_1	72.3 μs	0.54, 1.55 μs 1.05, 1.19 μs	89.6 μs
T_2^R	17.86 μs	0.77, 1.23 μs 1.19, 1.19 μs	18.11 μs
T_2^E	21.34 μs	–, 2.38 μs 2.45, 1.93 μs	25.50 μs
$\omega_{01}/2\pi$	98.95 MHz	4.517, 4.577 GHz 4.619, 4.672 GHz	144 MHz
$\omega_{\text{res}}/2\pi$	4.993 GHz	–	5.082 GHz
$g_{ic}/2\pi$	236 MHz	–	236 MHz
$g_{12}/2\pi$	40 MHz	–	40 MHz
E_J/h (GHz)	4.9928	16.87	4.3350
E_L/h (GHz)	0.5008	–	0.4921
E_C/h (GHz)	0.8805	0.1861	0.8829
$\chi/2\pi$	1.029 MHz	–	1.319 MHz

for 1 min, a 65 °C PRS for 20 min, and finally IPA for 1 min.

The next step is to pattern the NbTiN using electron-beam lithography. To prepare the sample for the lithography step, an anisotropic oxygen plasma descum, with 20 sccm O_2 at 0.1 mbar and 20 W for 2 min, is used to remove organic residues. To improve resist adhesion, we bake the sample for 10 min at 150 °C and spin a HMDS layer. We use AR.P6200.18 as the resist, which we spin at 2500 rpm and bake at 160 °C for 3 min. We finally pattern the coplanar waveguide and capacitive structures of the device using a Raith 5200 EBPG system using a beam current of 250 nA and a dose of 350 μC/cm^2. After exposure, the resist is developed in a multistep process: (i) amyl acetate development for 30 s with sonication, (ii) amyl acetate development for 30 s without sonication in a separate beaker, (iii) MIBK for 10 s, (iv) IPA for 60 s, (v) MIBK for 20 s with sonication, (vi) IPA for 15 s in separate beakers, repeated twice. To remove resist residues, an additional 2-min oxygen descum is performed, as in the previous descum step. The pattern is etched using reactive ion etching system in a two-step recipe. The first step uses SF_6:O_2 13.5:4 at 70 W and a pressure of 0.1 mbar until a signal drop from the endpoint detection system. The second step uses SF_6:O_2 4:16 at 50 W

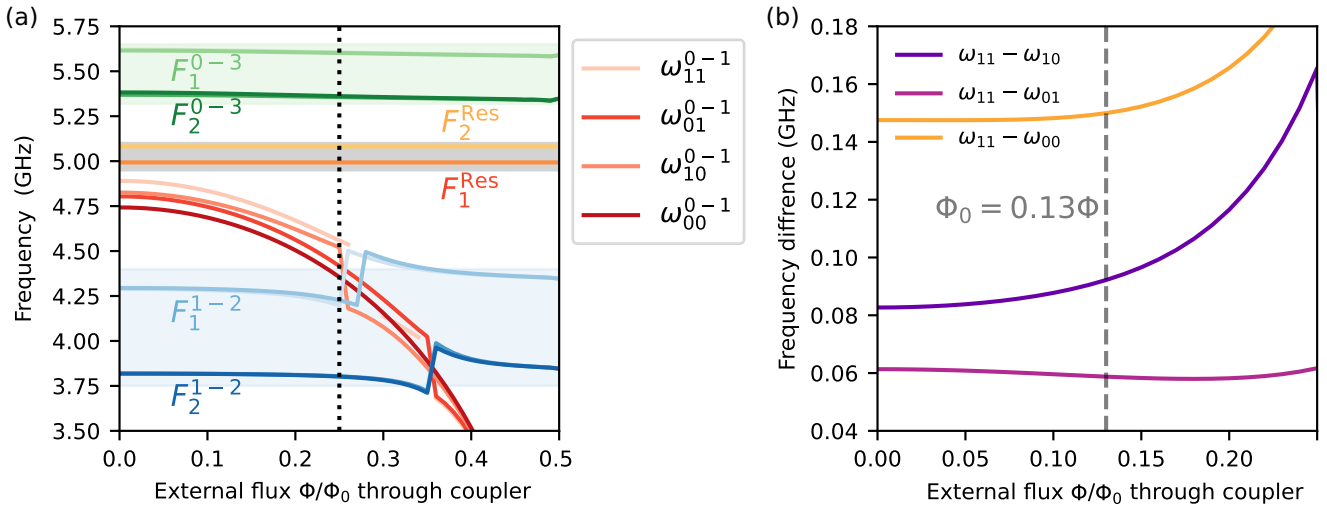

FIG. 8. (a) Energy levels of the fluxonium-transmon-fluxonium system as a function of the external flux of the transmon. The transition frequency from state $|n\rangle$ to $|m\rangle$ for fluxonium i is indicated with F_i^{n-m}. In the blue and green regions we see the 12 and 03 transitions of the fluxoniums, respectively. The gray region contains the two readout resonators. The coupler transitions are labeled ω_{ij}^{0-1}, corresponding to the fluxonium qubits in state $|ij\rangle$. (b) Differences between the undesired transition of the transmon (with the fluxoniums in $|00\rangle$, $|01\rangle$, and $|10\rangle$) and the desired transition $|11\rangle$ as a function of the external flux of the transmon. Here ω_{ij} represents F_1 prepared in the ith state and F_2 in j th state, respectively. A black vertical dashed line is marked to show the value of $\Phi = 0.13\Phi_0$, corresponding to the operational point used in the main text.

and a pressure of 0.08 mbar. We perform a final descum step with a 220 sccm flow of O_2 at 200 W for 30 s.

To remove the resist, we use the resist remover PRS-3000 first for 1 min with sonication followed by 3 h without sonication in a separate beaker. To further remove any organic residues from the surface, we clean for 30 s using nitric acid, followed by 25-min cleaning using 7:1 BOE to remove any native oxides [70].

The Josephson junctions are patterned with electron-beam lithography. This process uses a trilayer stack of MMA El8, PMMA 495 A8, and PMMA 950 A3 which we spin at 3000, 1000, and 3000 rpm, respectively. Each spinning step involves baking at 185 °C for 10 min and cooling down for 2 min. The junctions are defined by exposing the resist using a 5-nm resolution beam and a current of 1 nA. The exposure is designed to define a tapered undercut profile [71] using a selective exposure of the MMA resist. The undercuts are defined using a dose of 170 $\mu C/cm^2$, and we use a ratio of the junction dose to the undercut dose of 6:1. The development of the resist is performed in a H_2O:IPA (1:3) mixture at room temperature with sonication for 3 min 10 s. Subsequently, the device is cleaned in IPA twice for 10 s and 1 min, respectively. To clean resist residues, we do an oxygen descum, followed by cleaning with a 7:1 BOE mixture [72]. Immediately afterwards, we load the device into a Plassys MEB550 system for the electron-beam evaporation process. The junction deposition starts with 30 nm Al deposition followed by a 5-mbar, 20-min oxidation step. Then we deposit 110 nm of Al after rotating the sample 90° followed by another oxidation at 1.3 mbar for 11 mins.

The chip is cleaned using NMP at 80 °C for 2 h and then left overnight at room temperature. Next, the device is cleaned again using NMP, while the residual aluminum film is removed with a pipette. We additionally clean the device two more times with NMP and three additional times with IPA, for 5 min each.

To create a better galvanic connection between the junctions and the base layer, we deposit additional aluminum bandages [73]. We spin a resist stack of MMA EL8 and PMMA 950 A4 at 4000 and 2000 rpm, respectively. Each resist layer is baked at 180 °C for 5 min. The resist is exposed using electron-beam lithography with a beam current of 250 nA and a dose of 1200 $\mu C/cm^2$, and developed with a H_2O:IPA (1:3) mixture at 6 °C for 1 min 30 s followed by 30 s of IPA.

Prior to the Al deposition, an ion-milling step is performed to remove native oxides and resist residues. We then deposit 150 nm of Al film, followed by 1.3-mbar, 11-min oxidation. The lift-off is done similar to the junction lift-off step. The final result is illustrated in Fig. 6(b), where we see both the smaller Josephson junction as well as the larger junctions in the array that were fabricated simultaneously in this process. The whole device is finally diced again and we wirebond the device to a PCB, as shown in Fig. 6(a).

APPENDIX C: EXPERIMENTAL SETUP

The PCB holding the device is attached to a gold-plated copper mount and covered with two rectangular aluminum shields. The assembly is installed in a Blue Fors LD-400 system. The sample holder is further covered with three coaxial shields, consisting of two mu-metal cans and one gold-plated copper can. The PCB is connected to room-temperature electronics with coaxial microwave cables, as shown in Fig. 9.

We use a Zurich instruments (ZI) UHFQA to generate the input and output signals for readout. The signals are up- and down-converted using a ZI HDIQ IQ modulator and a

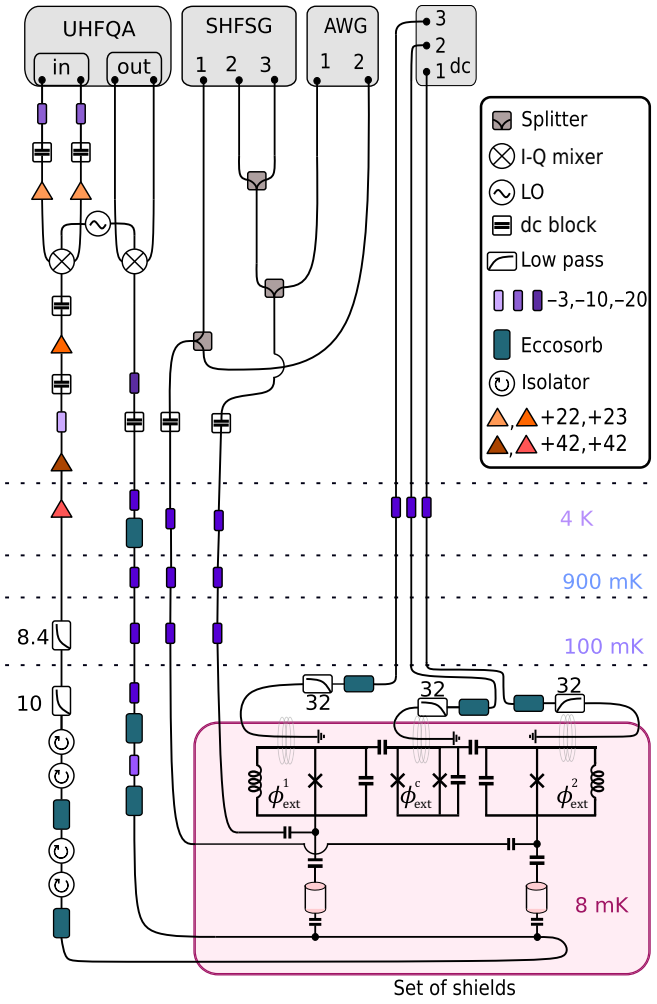

FIG. 9. Schematic of the wiring in the dilution fridge and electronic setup used during the experiment. The inset shows different components, with the numbers representing the attenuation and amplification in decibels. Filters used on the output line of the transmission chain are labeled in gigahertz, while the filters connected to the flux lines are labeled in megahertz.

custom built IQ modular, respectively. The LO is provided by a signal generator from Anapico. The output signal is amplified at the 4K stage using a HEMT amplifier from Low Noise Factory and at room temperature using a series of amplifiers, as shown in Fig. 9. For driving the individual fluxonium qubits, we use a ZI HDAWG that generates signals directly at the qubit frequencies. Additionally, we use a ZI SHFSQ to drive the transmon coupler as well as for qubit reset; see also Appendix F. These signals are combined with power combiners at room temperature. For flux bias of the qubits, we use in-house built current source modules (S4G). The flux lines on each qubit have 32-MHz low-pass filters to filter out noise at the fluxonium qubit frequency.

APPENDIX D: *ZZ* MEASUREMENT

The residual *ZZ* rate, as defined in Eq. (5), was measured using the circuit in Fig. 10. One fluxonium is designated as the control qubit and is initialized in either the $|0\rangle$ or $|1\rangle$ state. The other fluxonium is initialized in an equal superposition state. Both qubits freely evolve for a duration of $\tau/2$, after which we apply refocusing π pulses on both qubits simultaneously. The qubits then freely evolve again for a duration of $\tau/2$. Finally, a $\pi/2$ recovery rotation is applied for a set of different rotation angles θ. During the free evolution time, the target qubit rotates around the Z axis by $+\chi_{ZZ}\tau/2$ ($-\chi_{ZZ}\tau/2$) when the control qubit is prepared in the $|0\rangle$ ($|1\rangle$) state. This phase offset can be extracted from a cosine fit to the target qubit population as a function of θ.

APPENDIX E: COUPLER READOUT AND RABI

Since the transmon coupler does not have a dedicated readout resonator, we measure its state using a Rabi-assisted readout method [74]. In this scheme, both fluxoniums are initialized and kept in a known $|0\rangle$ or $|1\rangle$ state until the transmon is read out, at which point a π pulse is applied to both fluxonium qubits, which will only be resonant if the coupler is in its ground state. If the coupler is in its excited state, the fluxonium qubit frequencies are shifted by more than 60 MHz (see Fig. 4); hence, the π pulse will not flip the state of the fluxoniums. In other words, if the fluxonium qubits have flipped states then the coupler is in its excited state.

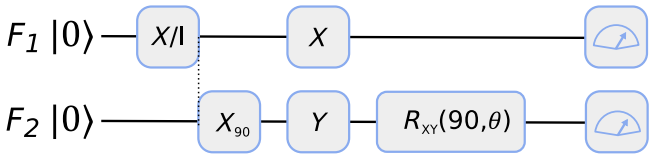

FIG. 10. Circuit for measuring the residual *ZZ* interaction rate.

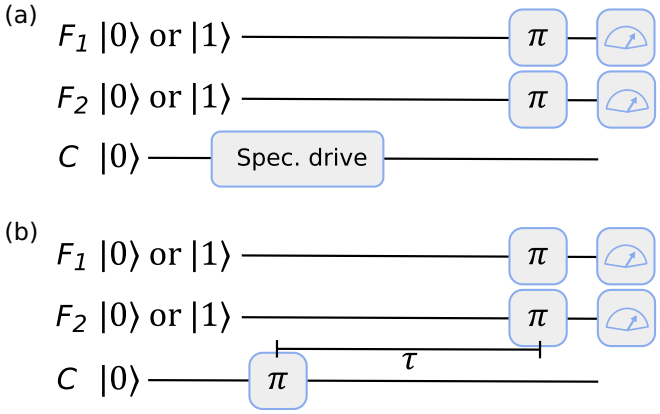

FIG. 11. (a) Circuit to measure the conditional spectroscopy for the coupler and fluxoniums. (b) Circuit to measure the coherence times of the coupler.

To avoid calibrating the state discrimination boundary in Rabi-assisted readout, we work directly with the complex integrated readout traces. Multiplexed readout produces a complex value for each fluxonium, resulting in a total of four real values. To observe Rabi oscillations in the coupler, we perform principal component analysis on the four-dimensional data points. The first principal component should capture the largest variations, corresponding to oscillations in the final coupler state. The second and subsequent principal components should have only small variations if the system is truly confined to a two-state manifold.

To characterize the lifetime of the coupler, we use the circuit shown in Fig. 11(b), where the coupler is excited with a calibrated π pulse. In this measurement, we initialize each fluxonium qubit in either $|0\rangle$ or $|1\rangle$ to extract any state-dependent lifetime of the coupler.

APPENDIX F: RESET

We reset the qubits using a two-tone driving scheme similar to Ref. [75]. The qubit reset protocol uses two drive pulses at the angular frequencies ω_1 and ω_2. The drive frequencies correspond to the transition frequency between $|1,0\rangle$ and $|2,0\rangle$, and to the transition frequency between $|2,0\rangle$ and $|0,1\rangle$, respectively. When the system is in state $|0,1\rangle$, it will decay back into $|0,0\rangle$ by the dissipation of one photon through the resonator; see Fig. 12(a). The reset uses a square bichromatic drive with a duration of 32 μs followed by an idle time of 8 μs to ensure that the resonator is fully dissipated. When the qubits (F_1 and F_2) are not reset, they thermalize into a mixed state with almost equal populations in $|0\rangle$ and $|1\rangle$. As a result, the single-shot readout histogram does not clearly distinguish between the qubit states; see Figs. 12(b) and 12(d). After applying the reset, the qubit is initialized in $|0\rangle$ states, and we now clearly see two well-separated histogram peaks in Figs. 12(c) and

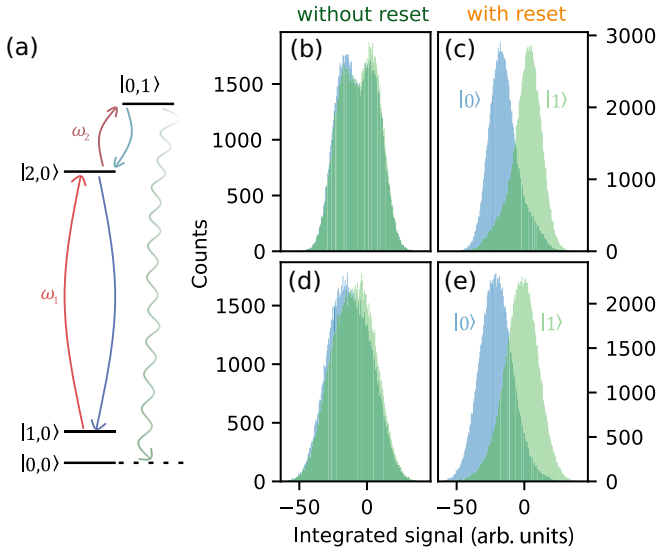

FIG. 12. (a) Reset scheme used for initializing the qubit in the ground state. The energy levels are labeled with $|Q,R\rangle$, where Q is the qubit state and R is the resonator state, shown with simultaneous driving at the angular frequencies ω_1 and ω_2, and with relaxation back into the qubit ground state. (b),(c) Single-shot readout histograms for qubit F_1. Panel (b) shows the outcome without performing the reset, while (c) shows the result after applying the reset. Similarly panel (d) and (e) represent data for qubit F_2.

12(e). The readout assignment fidelity for fluxonium qubits 1 and 2 is 80% and 77%, respectively.

APPENDIX G: CONDITIONAL PHASE GATE TUNE-UP

In this work, to calibrate the gate, we first measure Rabi oscillations on the coupler when the fluxoniums are prepared in the $|11\rangle$ state [see Fig. 13(a)], where the pulse duration and pulse shape are kept fixed. From this measurement, we extract the pulse amplitude for each drive frequency that maximally returns the transmon to its initial state [see the orange markers in Fig. 13(a)], which we fit with a spline function to establish a map from drive frequency to drive amplitude.

Using the extracted drive amplitudes, we determine the conditional phase from the CZ gate as a function of the drive frequency; see Fig. 13(b). Specifically, we perform a measurement of the conditional phase for 1, 3, and 5 gate repetitions. To efficiently sample the frequency axis, we use a noise-aware root-finding algorithm to locate the point of π phase. Once the phase error is within $\pm 5^\circ$, we increase the number of gate repetitions, repeat the procedure, and terminate when five gate repetitions yield a phase error of less than 5°.

Once the pulse amplitude and frequency is calibrated to generate the conditional phase gate, we calibrate the residual single-qubit phases for both fluxonium qubits. The single-qubit phases are measured by amplifying the single-qubit phase errors with repeated applications of the CZ gate; see Fig. 13(c). We linearly fit the accumulated phase data and use the slope as a measure of the acquired phase per gate. The unwanted single-qubit phases are compensated by using virtual-Z gates [62].

APPENDIX H: FAST-DRAG CALCULATIONS

To eliminate the undesired frequency components in the in-phase pulse envelope (Ω_I) of the drive used for the gate, we implemented the FAST-DRAG method [53] to derive analytical solutions for the shape of the pulse. This approach works by attenuating the frequency amplitude across n frequency spans (f_n), each spanning a range from $f_{n,l}$ to $f_{n,h}$. Additionally, the amplitude of each frequency span is suppressed by an amount proportional to a specified weight (w_n) assigned to that span. The method works by defining the pulse envelope, Ω_I, as a cosine series as described in Eq. (9):

$$\Omega_I = A \sum_{n=1}^{N} c_n g_n(t). \tag{H1}$$

Here, $g_n = [1 - \cos(2\pi n t/t_g)]$ is a cosine pulse with n periods for the duration of t_g. The pulse amplitude is given by A and each cosine component is weighted with a coefficient c_n. To solve for the coefficients c_n such that the combined pulse shape minimizes the frequency components at the undesired frequencies, we choose $N = 3$ basis functions and solve the quadratic optimization problem

$$\begin{aligned} \text{minimize} \quad & \sum_{n=1}^{3} w_n \int_{f_{n,l}}^{f_{n,h}} |\hat{\Omega}_I(f)|^2 \, df \\ \text{such that} \quad & \sum_{n=1}^{3} c_n t_g = \pi, \end{aligned} \tag{H2}$$

where $\hat{\Omega}_I(f)$ is the Fourier transform of the pulse envelope. The condition of π comes from the desired angle of rotation for the gate. To provide an experimental basis to the mathematical relations, we define the center frequencies of each frequency span to be the conditional transitions of the coupler $\{f_{00}, f_{01}, f_{10}\}$ not corresponding to the $|11\rangle$ state of the fluxoniums. Additionally, we assign the weights of each spurious frequency as the square of the Rabi rate of the coupler, corresponding to the transition (see Fig. 4). Furthermore, we use a frequency span of 5 MHz centered at each transition.

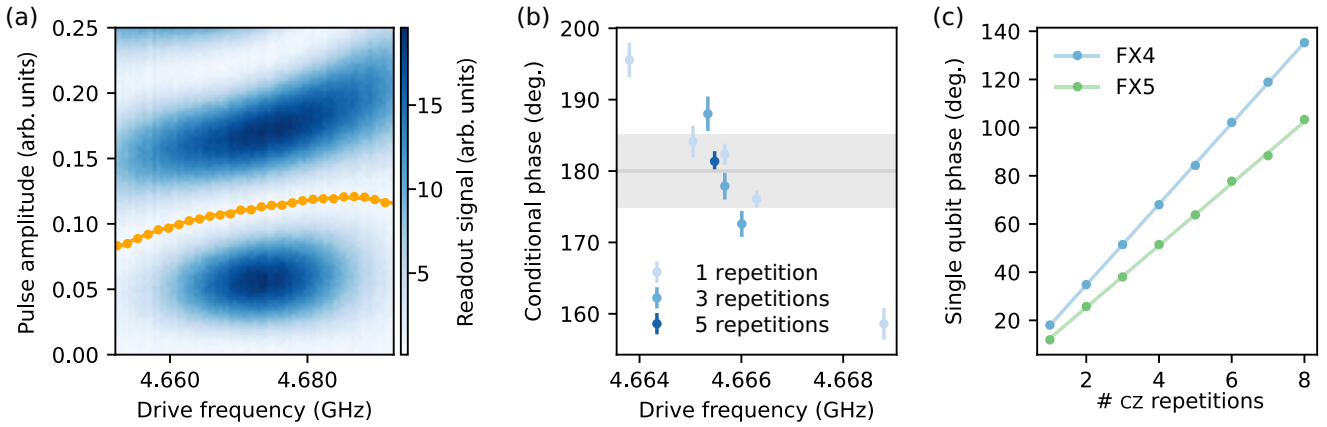

FIG. 13. (a) Coupler response to varying drive frequency and pulse amplitude for a cosine-shaped pulse with a duration of 64 ns. The heatmap corresponds to the magnitude of the difference in the Rabi-assisted readout signal when a drive is played versus when it is not, as described in the main text. The yellow markers indicate the 2π rotation amplitude for each drive frequency, and the solid yellow line is a spline fit through the markers. (b) Conditional phase versus drive frequency with varying gate repetitions. The gray band represents a conditional phase error of $\pm 5^\circ$. (c) Single-qubit phase errors amplified with an increasing number of gate repetitions.

To find the appropriate minimum as given in Eq. (H2), we first consider the Fourier transform of the basis function, as given by

$$\begin{aligned}\hat{g}_n = t_g(&e^{-i\pi t_g f}\,\mathrm{sinc}(\pi t_g f) \\ &- \tfrac{1}{2}e^{i\pi(n/t_g - f)t_g}\,\mathrm{sinc}[\pi(n/t_g - f)t_g] \\ &+ \tfrac{1}{2}e^{i\pi(n/t_g + f)t_g}\,\mathrm{sinc}[\pi(n/t_g + f)t_g],\end{aligned} \tag{H3}$$

such that we now need to minimize

$$\sum_{n=1}^{3} w_n \int_{f_{l,n}}^{f_{h,n}} \left| \sum_{i=1}^{3} c_n \hat{g}_n(f) \right|^2. \tag{H4}$$

We can rewrite Eq. (H4) as

$$\sum_{n,m} c_n c_m A_{nm} = c^T A c, \tag{H5}$$

where $A_{nm} = \sum_{i=1}^{3} w_i \int_{f_{i,l}}^{f_{i,h}} \hat{g}_n(\mathrm{f})\hat{g}_n^*(\mathrm{f})df$ and c is a vector containing the coefficients c_n. The minimization problem now becomes equivalent to solving the matrix equation

$$\begin{bmatrix} A + A^T & -b \\ b^T & 0 \end{bmatrix} \times (c^T, \mu)^T = (0^T, \pi)^T, \tag{H6}$$

where $b = (1, \ldots, 1)^T \in \Re^{N\times 1}$, and we have introduced the Lagrange multiplier μ. The pulse shapes presented in the main text are found using the solutions to this matrix equation.

[1] T. I. Andersen, N. Astrakhantsev, A. H. Karamlou, J. Berndtsson, J. Motruk, A. Szasz, J. A. Gross, A. Schuckert, T. Westerhout, Y. Zhang *et al.*, Thermalization and criticality on an analog-digital quantum simulator, arXiv:2405.17385.

[2] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of molecular spectra on a quantum processor with an error-resilient algorithm, Phys. Rev. X **8**, 011021 (2018).

[3] A. Miessen, Daniel J. Egger, I. Tavernelli, and G. Mazzola, Benchmarking digital quantum simulations above hundreds of qubits using quantum critical dynamics, PRX Quantum **5**, 040320 (2024).

[4] Y. Zhang, Yan Q. Zhu, J. Xu, W. Zheng, D. Lan, G. Palumbo, N. Goldman, Shi L. Zhu, X. Tan, Z. D. Wang, and Y. Yu, Exploring parity magnetic effects through quantum simulation with superconducting qubits, Phys. Rev. Appl. **21**, 034052 (2024).

[5] E. Kim, X. Zhang, Vinicius S. Ferreira, J. Banker, Joseph K. Iverson, A. Sipahigil, M. Bello, A. González-Tudela, M. Mirhosseini, and O. Painter, Quantum electrodynamics in a topological waveguide, Phys. Rev. X **11**, 011015 (2021).

[6] T. Orell, M. Zanner, Mathieu L. Juan, A. Sharafiev, R. Albert, S. Oleschko, G. Kirchmair, and M. Silveri, Collective bosonic effects in an array of transmon devices, Phys. Rev. A **105**, 063701 (2022).

[7] I. T. Rosen, S. Muschinske, C. N. Barrett, D. A. Rower, R. Das, D. K. Kim, B. M. Niedzielski, M. Schuldt, K. Serniak, M. E. Schwartz *et al.*, Flat-band (de)localization emulated with a superconducting qubit array, arXiv:2410.07878.

[8] I. T. Rosen, S. Muschinske, C. N. Barrett, A. Chatterjee, M. Hays, M. A. DeMarco, A. H. Karamlou, D. A. Rower, R. Das, D. K. Kim, B. M. Niedzielski, M. Schuldt, K. Serniak, M. E. Schwartz, J. L. Yoder, J. A. Grover, and W. D. Oliver, A synthetic magnetic vector potential in a 2D superconducting qubit array, Nat. Phys. **20**, 1881 (2024).

[9] R. Acharya *et al.*, Quantum error correction below the surface code threshold, Nature **638**, 920 (2025).

[10] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, Nature **605**, 669 (2022).

[11] J. F. Marques, B. M. Varbanov, M. S. Moreira, H. Ali, N. Muthusubramanian, C. Zachariadis, F. Battistel, M. Beekman, N. Haider, W. Vlothuizen, A. Bruno, B. M. Terhal, and L. DiCarlo, Logical-qubit operations in an error-detecting surface code, Nat. Phys. **18**, 80 (2022).

[12] H. Putterman, K. Noh, C. T. Hann, G. S. MacCabe, S. Aghaeimeibodi, R. N. Patel, M. Lee, W. M. Jones, and H. e. a. Moradinejad, Hardware-efficient quantum error correction via concatenated bosonic qubits, Nature **638**, 927 (2025).

[13] J. Koch, Terri M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, Phys. Rev. A **76**, 042319 (2007).

[14] R. Barends *et al.*, Diabatic gates for frequency-tunable superconducting qubits, Phys. Rev. Lett. **123**, 210501 (2019).

[15] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Demonstration of two-qubit algorithms with a superconducting quantum processor, Nature **460**, 240 (2009).

[16] S. A. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Manenti, M. P. da Silva, R. Sinclair, and E. e. a. Acala, Parametrically activated entangling gates using transmon qubits, Phys. Rev. Appl. **10**, 034050 (2018).

[17] J. M. Chow, J. M. Gambetta, A. W. Cross, S. T. Merkel, C. Rigetti, and M. Steffen, Microwave-activated conditional-phase gate for superconducting qubits, New J. Phys. **15**, 115012 (2013).

[18] Jerry M. Chow, A. D. Córcoles, Jay M. Gambetta, C. Rigetti, B. R. Johnson, John A. Smolin, J. R. Rozen, George A. Keefe, Mary B. Rothwell, Mark B. Ketchen, and M. Steffen, Simple all-microwave entangling gate for fixed-frequency superconducting qubits, Phys. Rev. Lett. **107**, 080502 (2011).

[19] Bradley K. Mitchell, Ravi K. Naik, A. Morvan, A. Hashim, J. M. Kreikebaum, B. Marinelli, W. Lavrijsen, K. Nowrouzi, David I. Santiago, and I. Siddiqi, Hardware-efficient microwave-activated tunable coupling between superconducting qubits, Phys. Rev. Lett. **127**, 200502 (2021).

[20] S. Sheldon, E. Magesan, Jerry M. Chow, and Jay M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, Phys. Rev. A **93**, 060302(R) (2016).

[21] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, Phys. Rev. B **81**, 134507 (2010).

[22] S. Krinner, S. Lazar, A. Remm, C. K. Andersen, N. Lacroix, G. J. Norris, C. Hellings, M. Gabureac, C. Eichler, and A. Wallraff, Benchmarking coherent errors in controlled-phase gates due to spectator qubits, Phys. Rev. Appl. **14**, 024042 (2020).

[23] N. Sundaresan, I. Lauer, E. Pritchett, E. Magesan, P. Jurcevic, and Jay M. Gambetta, Reducing unitary and spectator errors in cross resonance with optimized rotary echoes, PRX Quantum **1**, 020318 (2020).

[24] Michele C. Collodo, J. Herrmann, N. Lacroix, C. K. Andersen, A. Remm, S. Lazar, Jean C. Besse, T. Walter, A. Wallraff, and C. Eichler, Implementation of conditional phase gates based on tunable *ZZ* interactions, Phys. Rev. Lett. **125**, 240502 (2020).

[25] Y. Ye, S. Cao, Y. Wu, X. Chen, Q. Zhu, S. Li, F. Chen, M. Gong, C. Zha, and H.-L. Huang, Realization of high-fidelity controlled-phase gates in extensible superconducting qubits design with a tunable coupler, Chin. Phys. Lett. **38**, 100301 (2021).

[26] F. Marxer, A. Vepsäläinen, Shan W. Jolin, J. Tuorila, A. Landra, C. Ockeloen-Korppi, W. Liu, O. Ahonen, A. Auer, L. Belzane *et al.*, Long-distance transmon coupler with CZ-gate fidelity above 99.8%, PRX Quantum **4**, 010314 (2023).

[27] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, Daniel L.Campbell, Terry P. Orlando, S. Gustavsson, and William D. Oliver, Tunable coupling scheme for implementing high-fidelity two-qubit gates, Phys. Rev. Appl. **10**, 054062 (2018).

[28] Y. Sung, L. Ding, J. Braumuller, A. Vepsalainen, B. Kannan, M. Kjaergaard, A. Greene, Gabriel O. Samach, C. McNally, D. Kim, A. Melville, Bethany M. Niedzielski, Mollie E. Schwartz, Jonilyn L. Yoder, Terry P.Orlando, S. Gustavsson, and William D. Oliver, Realization of high-fidelity CZ and ZZ-free iSWAP gates with a tunable coupler, Phys. Rev. X **11**, 021058 (2021).

[29] Eyob A. Sete, N. Didier, Angela Q. Chen, S. Kulshreshtha, R. Manenti, and S. Poletto, Parametric-resonance entangling gates with a tunable coupler, Phys. Rev. Appl. **16**, 024050 (2021).

[30] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[31] K. C. Miao, M. McEwen, J. Atalaya, D. Kafri, L. P. Pryadko, A. Bengtsson, A. Opremcak, K. J. Satzinger, Z. Chen, P. V. Klimov *et al.*, Overcoming leakage in quantum error correction, Nat. Phys. **19**, 1780 (2023).

[32] J. F. Marques, H. Ali, B. M. Varbanov, M. Finkel, H. M. Veen, S. L. M. van der Meer, S. Valles-Sanclemente, N. Muthusubramanian, M. Beekman, N. Haider, B. M. Terhal, and L. DiCarlo, All-microwave leakage reduction units for quantum error correction with superconducting transmon qubits, Phys. Rev. Lett. **130**, 250602 (2023).

[33] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, Npj Quantum Inf. **7**, 14 (2021).

[34] K. Heya, M. Malekakhlagh, S. Merkel, N. Kanazawa, and E. Pritchett, Floquet analysis of frequency collisions, Phys. Rev. Appl. **21**, 024035 (2024).

[35] A. Morvan, L. Chen, Jeffrey M. Larson, David I. Santiago, and I. Siddiqi, Optimizing frequency allocation for fixed-frequency superconducting quantum processors, Phys. Rev. Res. **4**, 023079 (2022).

[36] S. Ganjam, Y. Wang, Y. Lu, A. Banerjee, C. U. Lei, L. Krayzman, K. Kisslinger, C. Zhou *et al.*, Surpassing millisecond coherence in on chip superconducting quantum memories by optimizing materials and circuit design, Nat. Commun. **15** (2024).

[37] Alexander P. Read, Benjamin J. Chapman, C. U. Lei, Jacob C. Curtis, S. Ganjam, L. Krayzman, L. Frunzio, and Robert J. Schoelkopf, Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity, Phys. Rev. Appl. **19**, 034064 (2023).

[38] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, Fluxonium: Single Cooper-pair circuit free of charge offsets, Science **326**, 113 (2009).

[39] Long B. Nguyen, Yen H. Lin, A. Somoroff, R. Mencia, N. Grabon, and Vladimir E. Manucharyan, High-coherence fluxonium qubit, Phys. Rev. X **9**, 041041 (2019).

[40] A. Somoroff, Q. Ficheux, Raymond A. Mencia, H. Xiong, R. Kuzmin, and Vladimir E. Manucharyan, Millisecond coherence in a superconducting qubit, Phys. Rev. Lett. **130**, 267001 (2023).

[41] F. Wang, K. Lu, H. Zhan, L. Ma, F. Wu, H. Sun, H. Deng, Y. Bai, F. Bao, X. Chang *et al.*, High-coherence fluxonium qubits manufactured with a wafer-scale-uniformity process, Phys. Rev. Appl. **23**, 044064 (2025).

[42] N. Earnest, S. Chakram, Y. Lu, N. Irons, R. K. Naik, N. Leung, L. Ocola, D. A. Czaplewski, B. Baker, J. Lawrence *et al.*, Realization of a Λ System with metastable states of a capacitively shunted fluxonium, Phys. Rev. Lett. **120**, 150504 (2018).

[43] M. F. S. Zwanenburg, S. Singh, E. Y. Huang, F. Yilmaz, T. V. Stefanski, J. Hu, P. Kumaravadivel, and C. K. Andersen, Single-qubit gates beyond the rotating-wave approximation for strongly anharmonic low-frequency qubits, Phys. Rev. Research **7**, 043290 (2025).

[44] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, Amir H. Karamlou, Thomas M. Hazard, K. Azar, David K. Kim, Bethany M. Niedzielski, A. Melville, Mollie E. Schwartz, Jonilyn L. Yoder, Terry P.Orlando, S. Gustavsson, Jeffrey A. Grover, K. Serniak, and William D.Oliver, High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler, Phys. Rev. X **13**, 031035 (2023).

[45] Q. Ficheux, Long B. Nguyen, A. Somoroff, H. Xiong, Konstantin N. Nesterov, Maxim G. Vavilov, and Vladimir E. Manucharyan, Fast logic with slow qubits: Microwave-activated controlled-Z gate on low-frequency fluxoniums, Phys. Rev. X **11**, 021026 (2021).

[46] E. Dogan, D. Rosenstock, L. Le Guevel, H. Xiong, Raymond A. Mencia, A. Somoroff, Konstantin N. Nesterov, Maxim G. Vavilov, Vladimir E. Manucharyan, and C. Wang, Two-fluxonium cross-resonance gate, Phys. Rev. Appl. **20**, 024011 (2023).

[47] Konstantin N. Nesterov, Ivan V. Pechenezhskiy, C. Wang, Vladimir E. Manucharyan, and Maxim G. Vavilov, Microwave-activated controlled-*Z* gate for fixed-frequency fluxonium qubits, Phys. Rev. A **98**, 030301(R) (2018).

[48] H. Zhang, C. Ding, D. K. Weiss, Z. Huang, Y. Ma, C. Guinn, S. Sussman, S. P. Chitta, D. Chen, Andrew A. Houck *et al.*, Tunable inductive coupler for high-fidelity gates between fluxonium qubits, PRX Quantum **5**, 020326 (2024).

[49] X. Ma, G. Zhang, F. Wu, F. Bao, X. Chang, J. Chen, H. Deng, R. Gao, X. Gao, L. Hu *et al.*, Native approach to controlled-*Z* gates in inductively coupled fluxonium qubits, Phys. Rev. Lett. **132**, 060602 (2024).

[50] Ilya A Simakov, Grigoriy S. Mazhorin, Ilya N. Moskalenko, Nikolay N. Abramov, Alexander A. Grigorev, Dmitry O. Moskalev, Anastasiya A. Pishchimova, Nikita S. Smirnov, Evgeniy V. Zikiy, Ilya A. Rodionov, and Ilya S. Besedin, Coupler microwave-activated controlled-phase gate on fluxonium qubits, PRX Quantum **4**, 040321 (2023).

[51] I. N. Moskalenko, I. A. Simakov, N. N. Abramov, A. A. Grigorev, D. O. Moskalev, A. A. Pishchimova, N. S. Smirnov, E. V. Zikiy, I. A. Rodionov, and I. S. Besedin, High fidelity two-qubit gates on fluxoniums using a tunable coupler, Npj Quantum Inf. **8**, 130 (2022).

[52] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly-nonlinear qubits, Phys. Rev. Lett. **103**, 110501 (2009).

[53] E. Hyyppä, A. Vepsäläinen, M. Papič, C. F. Chan, S. Inel, A. Landra, W. Liu, J. Luus, F. Marxer, C. Ockeloen-Korppi *et al.*, Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes, PRX Quantum **5**, 030353 (2024).

[54] Austin G. Fowler, M. Mariantoni, John M. Martinis, and Andrew N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).

[55] R. Versluis, S. Poletto, N. Khammassi, B. Tarasinski, N. Haider, D. J. Michalak, A. Bruno, K. Bertels, and L. DiCarlo, Scalable quantum circuit and control for a superconducting surface code, Phys. Rev. Appl. **8**, 034021 (2017).

[56] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff, Repeated quantum error detection in a surface code, Nat. Phys. **16**, 875 (2020).

[57] Google Quantum AI, Suppressing quantum errors by scaling a surface code logical qubit, Nature **614**, 676 (2023).

[58] F. Yilmaz, S. Singh, M. F. S. Zwanenburg, J. Hu, T. V. Stefanski, and C. K. Andersen, Energy participation ratio analysis for very anharmonic superconducting circuits, arXiv:2411.15039.

[59] Z. K. Minev, Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, M. H. Devoret *et al.*, Energy-participation quantization of Josephson circuits, Npj Quantum Inf. **7**, 131 (2021).

[60] E. Magesan, Jay M. Gambetta, and J. Emerson, Characterizing quantum gates via randomized benchmarking, Phys. Rev. A **85**, 042311 (2012).

[61] A. Silva and E. Greplova, Hands-on introduction to randomized benchmarking, SciPost Phys. Lect. Notes **97** (2025).

[62] D. C. McKay, Christopher J. Wood, S. Sheldon, Jerry M. Chow, and Jay M. Gambetta, Efficient *Z* gates for quantum computing, Phys. Rev. A **96**, 022330 (2017).

[63] J. Kelly *et al.*, Optimal quantum control using randomized benchmarking, Phys. Rev. Lett. **112**, 240504 (2014).

[64] D. Puzzuoli, S. F. Lin, M. Malekakhlagh, E. Pritchett, B. Rosand, and C. J. Wood, Algorithms for perturbative analysis and simulation of quantum dynamics, J. Comput. Phys. **489**, 112262 (2023).
[65] https://doi.org/10.4121/cec2d36b-c31b-4c70-a66f-65d3ef7bd514 (2025).
[66] https://github.com/AndersenQubitLab/fast_two_qubit_gate (2025).
[67] P. Groszkowski and J. Koch, scqubits documentation, https://scqubits.readthedocs.io/en/latest (2021).
[68] T. V. Stefanski, F. Yilmaz, E. Y. Huang, M. F. S. Zwanenburg, S. Singh, S. Wang, L. J. Splitthoff, and C. K. Andersen, Improved fluxonium readout through dynamic flux pulsing, APL Quantum **2**, 026103 (2025).
[69] A. Bruno, G. de Lange, S. Asaad, K. L. van der Enden, N. K. Langford, and L. DiCarlo, Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates, Appl. Phys. Lett. **106**, 182601 (2015).
[70] M. V. Altoé, A. Banerjee, C. Berk, A. Hajr, A. Schwartzberg, C. Song, M. Alghadeer, S. Aloni, Michael J. Elowson, J. M. Kreikebaum, Ed. K. Wong, Sinead M. Griffin, S. Rao, A. WeberBargioni, Andrew M. Minor, David I.Santiago, S. Cabrini, I. Siddiqi, and D. F. Ogletree, Localization and mitigation of loss in Niobium superconducting circuits, PRX Quantum **3**, 020312 (2022).
[71] N. Muthusubramanian, M. Finkel, P. Duivestein, C. Zachariadis, S. L. M. van der Meer, H. M. Veen, M. W. Beekman, T. Stavenga, A. Bruno, and L. DiCarlo, Wafer-scale uniformity of Dolan-bridge and bridgeless Manhattan-style Josephson junctions for superconducting quantum processors, Quantum Sci. Technol. **9**, 025006 (2024).
[72] D. C. Zanuz, Q. Ficheux, L. Michaud, A. Orekhov, K. Hanke, A. Flasby, M. Bahrami Panah, G. J. Norris, and M. e. a. Kerschbaum, Mitigating losses of superconducting qubits strongly coupled to defect modes, Phys. Rev. Appl. **23**, 044054 (2025).
[73] A. Dunsworth, A. Megrant, C. Quintana, Z. Chen, R. Barends, B. Burkett, B. Foxen, Y. Chen, B. Chiaro, A. Fowler *et al.*, Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits, Appl. Phys. Lett. **111**, 022601 (2017).
[74] C. Zhang, T.-L. Wang, L.-L. Guo, X.-Y. Yang, X.-X. Yang, P. Duan, Z.-L. Jia, W.-C. Kong, and G.-P. Guo, Characterization of tunable coupler without a dedicated readout resonator in superconducting circuits, Appl. Phys. Lett. **122**, 022601 (2023).
[75] H. Zhang, S. Chakram, T. Roy, N. Earnest, Y. Lu, Z. Huang, D. K. Weiss, J. Koch, and David I. Schuster, Universal fast-flux control of a coherent, low-frequency qubit, Phys. Rev. X **11**, 011010 (2021).