








Experimental advances with the QICK (Quantum Instrumentation Control Kit) for superconducting quantum hardware

Chunyang Ding ¹, Martin Di Federico ², Michael Hatridge,³ Andrew Houck,⁴ Sebastien Leger,¹ Jeronimo Martinez ⁴,
 Connie Miao ¹, David Schuster I,¹ Leandro Stefanazzi,² Chris Stoughton,² Sara Sussman ², Ken Treptow,²
 Sho Uemura ², Neal Wilcer,² Helin Zhang,⁵ Chao Zhou,³ and Gustavo Cancelo ^{2,*}

¹Department of Physics and Applied Physics, Stanford University, Stanford California 94305, USA

²Fermi National Accelerator Laboratory, Batavia Illinois 60510, USA

³Department of Physics & Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA

⁴Department of Electrical Engineering, Princeton University, Princeton, New Jersey 08544, USA

⁵Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA



(Received 18 December 2023; accepted 25 February 2024; published 20 March 2024)

The Quantum Instrumentation Control Kit (QICK) is a standalone open-source qubit controller that was first introduced in 2022. In this follow-up work, we present recent upgrades to the QICK and the experimental use cases they uniquely enabled for superconducting qubit systems. These include multiplexed signal generation and readout, mixer-free readout, predistorted fast flux pulses, and phase-coherent pulses for parametric operations, including high-fidelity parametric entangling gates. We explain in detail how the QICK was used to enable these experiments.

DOI: [10.1103/PhysRevResearch.6.013305](https://doi.org/10.1103/PhysRevResearch.6.013305)

I. INTRODUCTION

Achieving high-fidelity qubit control and readout requires significant RF engineering and digital signal processing. To control superconducting qubits, for example, signals range from DC to upwards of 10 GHz, with complex envelopes. Signals must jump in frequency while maintaining phase coherence and meeting very precise timing requirements. The total control pulse sequence must be significantly faster than the coherence time of any qubit in the system (typically microseconds), and the latency between any two pulses must be controllable (zero latency is often required). In addition, low-latency conditional logic is required for feedback and feedforward. Each qubit typically needs several dedicated lines for control and readout, and hundreds of qubits must be operated in parallel to perform a nontrivial quantum computation [1,2].

The Quantum Instrumentation Control Kit (QICK) is an open-source system which uses the AMD-Xilinx RFSoc to deliver these functions. The RFSoc [3,4] is a system-on-chip that incorporates high-speed DACs and ADCs and has been adopted for qubit control by both the research community [5,6] and commercial industry [7]. RFSoc-based open-source qubit controllers [8,9] are a growing alternative to commercial solutions [10,11].

The first-generation QICK was introduced in Ref. [9] as an open-source superconducting qubit controller running on

an RFSoc evaluation board. The QICK firmware is a scaffold with a modular timed processor and signal generation/readout blocks that the user can combine in many configurations. The QICK firmware and software can be downloaded from a public GitHub repository [12,13]. The original RFSoc chip supported by QICK was the Gen 1 ZU28DR running on the AMD-Xilinx ZCU111 evaluation board [14]. In this paper, we describe QICK running on the Gen 3 ZU49DR RFSoc chip and its evaluation boards [15,16].

The Gen 3 ZCU216 evaluation board is of particular interest because of its 16 DACs running at 9.85 GS/s. The large number of high-speed DACs makes the ZCU216 useful for controlling multiqubit systems. In addition, the ZCU216 also has 16 ADCs running at 2.5 GS/s. For this next-generation hardware, QICK firmware was developed that enables direct, mixer-free generation of pulses for qubit drives and multiplexed readout up to 10 GHz. Additionally, the 100 ps time resolution enables precisely defined qubit control pulses. Phase coherence is maintained across all ZCU216 channels, reducing the need for external triggers and synchronized clocks across different equipment.

Figure 1 shows a typical control loop for measuring a superconducting qubit system using the QICK running on the ZCU216 evaluation board. This scheme eliminates the warm analog components used in superconducting qubit control besides amplification and filtering. A QICK RF companion board for the ZCU216 has been designed and characterized and is currently being field tested with qubits. Due to the large analog bandwidth of the ZCU216, its QICK RF companion board is correspondingly simpler than that of the ZCU111 [9]. As a cost-effective or educational alternative, QICK can run on the RFSoc4x2 evaluation board which is sold with special academic pricing. The RFSoc4x2 has 2 DACs running at 9.85 GS/s and 4 ADCs running at 5 GS/s.

*cancelo@fnal.gov

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




FIG. 1. A typical superconducting qubit control loop using the QICK on the ZCU216 board. Drive pulses are filtered then amplified before being sent to the fridge. The readout pulse is filtered before being sent to the fridge. After the fridge, the readout pulse is amplified before being directly read into the ADC.

The new capabilities we describe in this paper were also facilitated by an upgraded PYTHON software library [12]. The QICK software is based on the open-source Xilinx PYNQ (PYTHON productivity on Zynq) operating system [4,17], which is distributed as disk images from AMD-Xilinx [18] and the QICK team [19]. Every QICK firmware block has a corresponding PYNQ driver that abstracts the hardware details. These drivers are automatically attached to the firmware block without any user intervention and are maintained as part of the QICK library. On top of the driver abstraction, the QICK library has a powerful parser and automatic connectivity detection algorithm that creates data structures to efficiently handle all the available resources, such as different types of signal generator and readout blocks that all run at different speeds.

Several experiments have been already published using the QICK to control superconducting quantum hardware [20–28]. In this paper, we present recent experimental use cases that were uniquely enabled by the QICK. We highlight some key capabilities (Sec. II) and relevant examples: multiplexed signal generation and readout (Sec. III), predistorted fast flux pulses (Sec. IV), phase-coherent parametric control (Sec. V), and phase-sensitive parametric entangling gates (Sec. VI). We conclude in Sec. VII with remarks on further improvements planned for the QICK.

II. SIGNAL GENERATION AND READOUT

QICK has an extensive library of “signal generator” and “readout” firmware blocks which connect to the DACs and ADCs of the FPGA. These blocks are compatible with any of the supported RFSoc boards and any DAC or ADC sampling rate supported by the FPGA, and can be combined as needed for the requirements of different experiments. In addition, we have released compiled and tested versions of the firmware, enabling researchers to use high-level PYTHON code to immediately control superconducting qubit systems.

A. Multiplexed signal generation and readout

Multiple signals can be frequency-multiplexed on a single DAC output running at up to 9.85 GS/s. Figure 2 shows one approach (used in Sec. III), where independent DDS channels are digitally summed before the DAC. For multiplexing large




FIG. 2. Multiplexed signal generation. Frequency multiplexing is achieved by adding two or more single-frequency signal generators before the DAC module.

numbers of tones, a different approach using a polyphase filter bank (PFB) [29] is possible.

Figure 3 shows a multiplexed readout, also used in Sec. III. The ADC data stream has a bandwidth equal to the Nyquist frequency; a PFB applies an array of bandpass filters to split this spectrum evenly into channels. Each channel has its own DDS oscillator which can be used to demodulate a signal that falls in its frequency range. In the current implementation with eight channels and four outputs, each channel has a width of 1/16 the sampling frequency and up to four channels can be read out simultaneously. In future work, we will implement a larger number of RF multiplexed output and input channels.

B. Full-speed and interpolated envelopes

Waveform envelopes, such as Gaussian, DRAG, triangular, or user-provided, represent general waveform shapes which the signal generator can use to parametrically create pulses with arbitrary frequency, amplitude, and phase.

In the standard “full-speed” signal generator, the envelope sample rate is equal to the DAC sample rate. This enables precise control of the pulse shape and sub-nanosecond time resolution, as used in Sec. IV. Oscillating envelopes can be used for fast chirps or detuned pulses which can use the full DAC bandwidth.

We have also developed an interpolated signal generator where the envelope rate is 1/16 the DAC rate, and time-domain interpolation is used to upsample the envelope to the DAC rate [30]. This uses both the envelope memory and the FPGA logic more efficiently, optimizing resource use for applications with lower-bandwidth envelopes. The interpolated

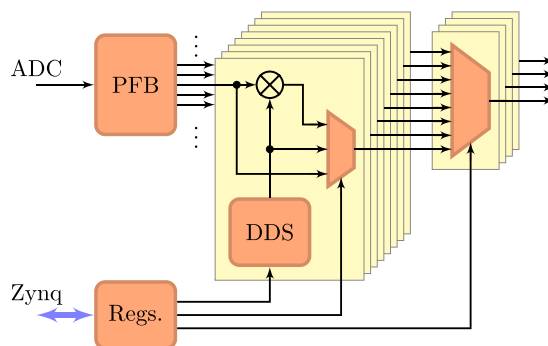


FIG. 3. Multiplexed readout. A polyphase filter bank (PFB) digitally demultiplexes the ADC samples into 8 channels with 50% overlap, to avoid gain losses over the entire bandwidth.

generators were used to implement a QICK firmware where all 16 RF-DACs were enabled.

C. Phase coherence

All generators and readouts use digital mixers, where a complex carrier oscillator is produced by direct digital synthesis (DDS) and is multiplied by an envelope (for a signal generator) or the ADC data stream (for a readout). Because the DDS oscillators are purely numerical and driven by a common FPGA clock, their relative phases are completely predictable. This scheme allows QICK signal generators to preserve phase coherence across multiple pulses of the same frequency while frequency hopping [9]. We have now added the capability (used in Secs. V and VI) to apply a synchronous phase reset to multiple signal generators and/or readouts, which can be used to set a fixed relative phase between pulses of differing frequencies.

D. Mixer-free readout

The ZCU216 DACs have a maximum sampling frequency of 9.85 GS/s, and the ADCs have a maximum sampling frequency of 2.5 GS/s. However, this does not impose a hard limit on the frequencies of signals that can be used; both the DACs and ADCs are limited only by their analog bandwidths. The RFSoc DACs can be configured for normal mode (zero-order hold, optimal for the first Nyquist zone) or “mix mode” (boosted power in the second and third Nyquist zones).

As shown in Fig. 1, the QICK readout input is capable of using direct sampling to simplify the analog electronics and avoid the use of any analog mixers [30,31]. While most of the experiments presented in this paper used existing readout amplifier chains that had already been optimized for analog mixers (Secs. III, IV, and VI), the experiment in Sec. V implemented mixer-free readout. To perform mixer-free readout with the QICK, an analog bandpass filter should be applied to the input signal to remove noise from other Nyquist zones, and a preamplifier should be used to compensate for the reduced ADC gain at higher frequencies [shown in Fig. 4(a)]. Figure 4(b) shows that this prescription performs well: the noise near the signal frequency is dominated by phase noise, and further away is set by the white noise floor. In the worst case measured (100 Hz separation from a 9 GHz signal), the phase noise is still below -70 dBc, which is low enough for typical quantum measurements.

There are many advantages in doing RF up and downconversion in the digital domain compared to using RF analog mixers (i.e., either IQ or DSB analog mixers). As we will demonstrate in Secs. V and VI, one of the main benefits is long term frequency and phase stability. Another benefit is the elimination of expensive and power-hungry analog components such as analog mixers, LO synthesizers (with power output requirements of ~ 15 dBm), filters for large LO feedthrough rejection, etc. Furthermore, analog IQ mixers require two DAC (ADC) outputs (inputs), which are an expensive commodity in quantum control electronics. Analog mixers are subject to drifts and DC offsets and require periodic calibrations. The ability of having multiple outputs (inputs) (e.g., 16 DACs, 16 ADCs per FPGA) running, individually,




FIG. 4. (Top) ZCU216 ADC analog gain as a function of analog input frequency. The transfer function is measured in 0.5 GHz steps at a constant input power of -7 dBm. The ADC sample rate is 2.4572 GHz. (Bottom) ZCU216 ADC phase noise measured at -1dBFS of power, in the range from 4 to 9 GHz f_c and at delta frequencies from 100 Hz to 10 MHz from f_c .

in a wide range of frequencies up to 10 GHz, all synchronized, without requiring periodic calibration is of paramount importance. Frequency and phase jumps can be manipulated back and forth in nanoseconds with no error since they are a digital number. The only disadvantage of RF digital up and downconversion is that it requires more FPGA logic.

III. FOUR-QUBIT SIMULTANEOUS READOUT USING MULTIPLEXED SIGNAL GENERATION AND READOUT

The multiplexed signal generator and readout described in Sec. II A makes experiments more hardware-efficient and less prone to calibration and user error. To illustrate this, we introduce an experiment realized in the Schuster Laboratory at Stanford.

The measured system consists of four capacitively coupled fixed-frequency transmon qubits. In this experiment,

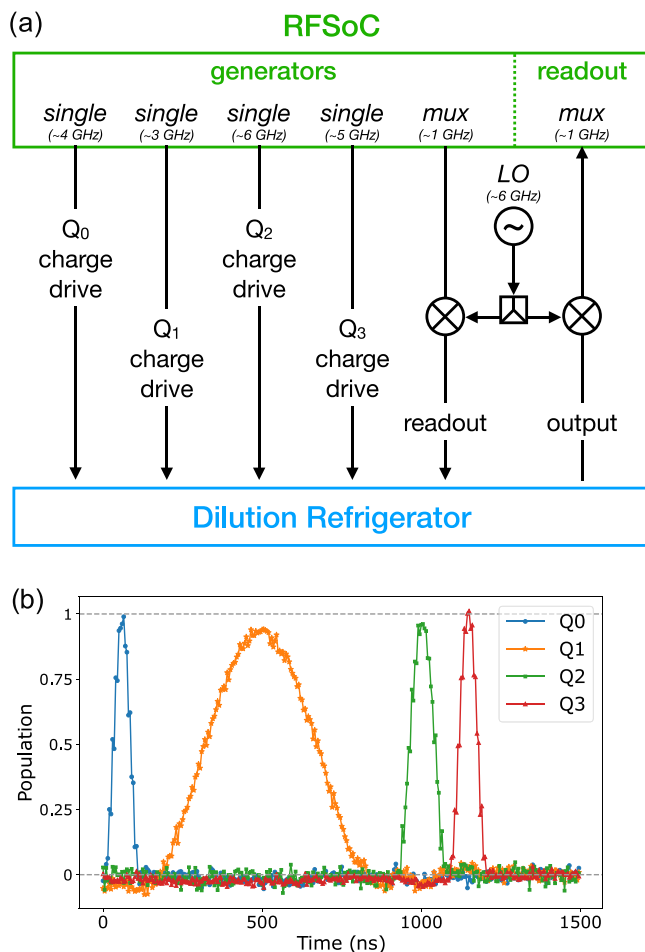


FIG. 5. Wiring schematic and qubit population measured for four-qubit simultaneous readout. (a) We use one full-speed generator per qubit to drive each qubit’s charge line, plus one multiplexed generator and one multiplexed readout channel (mixed up/down with an external LO) to readout all four qubits simultaneously. We thus are able to perform simultaneous four qubit control and readout with just one RFSoc board. (b) An example of a pulse sequence where four transmon qubits are driven sequentially with two consecutive π pulses. The qubit populations are sampled simultaneously at 5 ns intervals over the span of the gate sequence, with the readout handled by just 1 DAC and 1 ADC on the QICK board.

we needed simultaneous qubit readout for several purposes. First, it allowed us to track qubit populations over the time span of a sequence of gates that act on one or more qubits [e.g., Fig. 5(b)], which provides an easy-to-interpret visual representation of the success of a complex protocol. Second, simultaneous readout was necessary to perform quantum state tomography on sets of two or more qubits. Tomography requires simultaneous single-shot readout of all qubits involved in the tomography in order to categorize shots into the proper states.

In Fig. 5(b), we show qubit measurements during a simple protocol in which we performed two consecutive π pulses on each of the four qubits. At each point along the time axis, we measured the population of all qubits, where the y axis is scaled by using the distance between the resonator I/Q peaks when its respective qubit is in $|e\rangle$ versus $|g\rangle$.

The room temperature measurement setup used to realize the previous measurement is shown in Fig. III(a). Importantly, we were able to independently control each of the four qubits and perform simultaneous readout using just five DAC channels and one ADC, all on one QICK board, plus mixing with one external LO. The five DAC channels were respectively four standard generators used to drive each qubit and one multiplexed signal generator used to measure all of them. In theory each of our four resonators could have been measured with just the RFSoc board alone. However, there were a few constraints that we needed to satisfy that led us to introduce an external mixer and LO: (1) The physical resonator frequencies were 6805, 5791, 7697, and 6966 MHz for qubits 0–3, respectively, which spanned a range of 1904 MHz. (2) The multiplexed signal generator used in this implementation sampled at $f_{\text{DAC}} = 6881.28$ MHz, so each of the four DDS’s had a range of $f_{\text{DAC}}/4 = 1720$ MHz - less than the span of our resonator frequencies. The outputted frequencies were specified as $f_{\text{mix}} + [f_0, f_1, f_2, f_3]$, with each $f_{\text{mix}} + f_i$ limited by f_{DAC} , and each f_i had to fall between $(-f_{\text{DAC}}/8, f_{\text{DAC}}/8)$. (3) The multiplexed readout (sampling at $f_{\text{ADC}} = 2457.6$ MHz) required that each of the four frequencies fell in a different frequency bin of width $f_{\text{ADC}}/16 = 153.6$ MHz. To satisfy all of these constraints, we decided to set the RFSoc output at lower frequencies ($950 + [-70, -816, 822, 91] \text{ MHz}$) and mixed it up with an LO (5925 MHz) from a Signalcore 5511A. This choice of multiplexed and LO frequencies allowed us to span all of our readout frequencies by taking the positive sideband from the external mixer for qubits 0, 2, and 3 and the negative sideband for qubit 1. The ability to use the multiplexed signal generator and readout thus significantly reduced the quantity of wiring and splitters needed, though it also reduced our flexibility in designing the filtering and amplification wiring for each readout frequency individually.

IV. PREDISTORTED FAST FLUX PULSES

In the Houck laboratory at Princeton, four frequency-tunable transmon qubits coupled in a ring configuration were prepared to explore the dynamics of particles with and without a synthetic magnetic field [21]. Here, qubits play the roles of lattice sites and microwave excitations play the role of particles modeling the Bose-Hubbard Hamiltonian. The configuration of qubits is of a plaquette of a lattice whose energy bands contain one flat band. Under the addition of a synthetic magnetic field, all bands of the lattice become flat and all single-particle dynamics becomes localized. This is due to destructive interference arising from a combination of the Aharonov-Bohm effect and a particular lattice geometry [32]. One experimental challenge is that flat-band states are highly sensitive to disorder in the lattice site (qubit) frequencies. Therefore any time-varying shifts in qubit frequency need to be minimized.

The experimental sequence used to characterize dynamics consists of three steps: state initialization, time evolution, and readout. In each step, the qubits are biased to different frequencies by threading magnetic flux through the Josephson junction loop, modifying the qubit’s effective inductance. For state initialization, qubits are detuned from each other to be individually addressable. After initializing individual qubits

to their excited states, the qubits are diabatically tuned onto resonance to start the time evolution. We require that this frequency ramp happens much faster than a characteristic tunneling time and that the qubit frequency is stable for the duration of the time evolution of a few microseconds. Finally, the qubits are detuned, freezing the dynamics, and the states of each qubit are measured using a dispersively coupled cavity.

Ideally, the qubit frequency is changed instantaneously by sending a step-function pulse down the control lines in the dilution fridge. However, the shape of the fast-flux pulses generated at room temperature will be distorted by imperfections such as reflections and frequency-dependent attenuation as it travels down the fridge lines. On chip, the square step-function pulse arrives misshapen, greatly slowing down the frequency tuning of the qubit. These distortions can be corrected by “predistorting” the pulse shape on the control hardware such that the qubit receives a sharp and flat flux pulse. The ZCU216 board’s RF-DAC outputs running at 6.88 GS/s were used to calibrate and precisely specify “predistorted” fast flux pulses with 145 ps resolution.

The fast-flux pulses experience distortions on different timescales that are addressed individually. First, the qubit frequencies need to be shifted much faster than the characteristic tunneling time of qubits. For typical coupling rates of tens of MHz, qubits must be placed on resonance within a few nanoseconds. The second timescale requires the qubit frequency to remain constant during the evolution time (μs timescale) of the experiment. We compensate for both using different schemes.

For the short-time distortions, we use an experiment inspired by T_2 Ramsey to infer the frequency shift of the qubit during the fast flux pulse similarly to that done in Ref. [33]. The pulse sequence is given in Fig. 6(a). We first initialize the qubit in the superposition state $\frac{1}{2}(|0\rangle + |1\rangle)$ with a $X_{\pi/2}$ pulse. We then apply a fast flux pulse which detunes the qubit relative to the drive frequency. The qubit accumulates a phase, ϕ , which we can measure by applying a second $X_{\pi/2}$ or $Y_{\pi/2}$ pulse before measuring the qubit population:

$$e^{i\phi(t)} = \frac{X(t) + iY(t)}{\sqrt{X^2 + Y^2}}, \quad (1)$$

in which ϕ is directly proportional to the detuning, Δ , between the qubit and drive:

$$\phi(t) = \int_0^t \Delta(\tau) d\tau. \quad (2)$$

By differentiating Eq. (1), we can measure $\Delta(t)$ and map out the qubit frequency response. In this scheme, we note that the initial qubit frequency is biased to a “sweet spot” where the qubit frequency is first-order insensitive to external flux. This is done to suppress any distortions when the qubit returns to the initial frequency that could alter the accumulated phase. The measured qubit population after the $X_{\pi/2}$ or $Y_{\pi/2}$ pulses is seen in Fig. 6(b). We convert the slope at each time step into the detuning between the qubit and drive frequency, plotted in red in Fig. 6(c). This data is then fit to a sum of exponentials to generate a transfer function, which we invert to define a predistorted pulse. The qubit response to the predistorted step

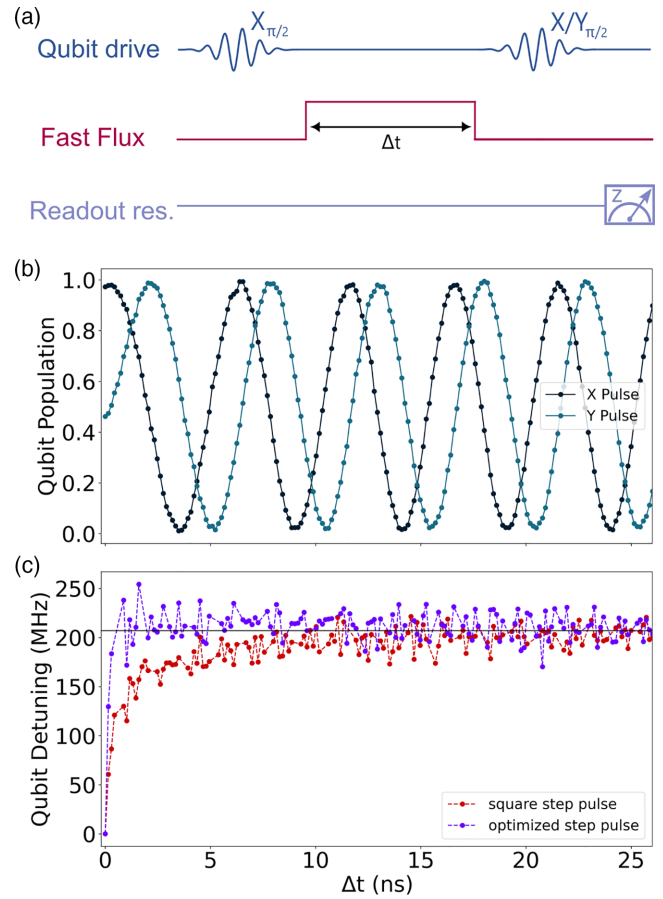


FIG. 6. Procedure for calibrating for short-time step-pulse distortions. (a) Pulse sequence consisting of a step pulse of variable length between two $\pi/2$ pulses. During the fast flux pulse, the qubit frequency is detuned relative to the drive frequency and accumulates a relative phase. (b) Measured qubit population after applying the second $X_{\pi/2}$ pulse (black) or $Y_{\pi/2}$ pulse (blue). Rate of phase accumulation maps directly to the detuning between the qubit and drive. (c) Detuning between qubit and drive, inferred from (b), for a square step-pulse (red) and optimized step-pulse postcalibration (purple).

pulse is shown in Fig. 6(c) where it approaches the target frequency within a nanosecond.

To calibrate for long-time distortions, we do a spectroscopy experiment similar to that in the supplement of Ref. [34]. We send a step pulse on the flux line and then send a short Gaussian pulse with $\sigma = 7$ ns to measure the qubit frequency for variable delay times. The amplitude is weaker than that needed to drive a π pulse. In comparison to the previous scheme, this method measures the qubit frequency with higher precision but suffers from poor time resolution limited by the qubit frequency drifting while the spectroscopy drive pulse is applied. The pulse sequence is shown in Fig. 7(a). We take spectroscopy data after different delay times and fit each time step to a Lorentzian function. The fitted qubit response to the square step pulse is seen in Fig. 7(b). The data is similarly fit to a sum of exponentials and inverted to define the predistorted step pulse. We find that the predistorted pulse allows for the qubit frequency to be stable within 0.1% for the duration of the experiment.

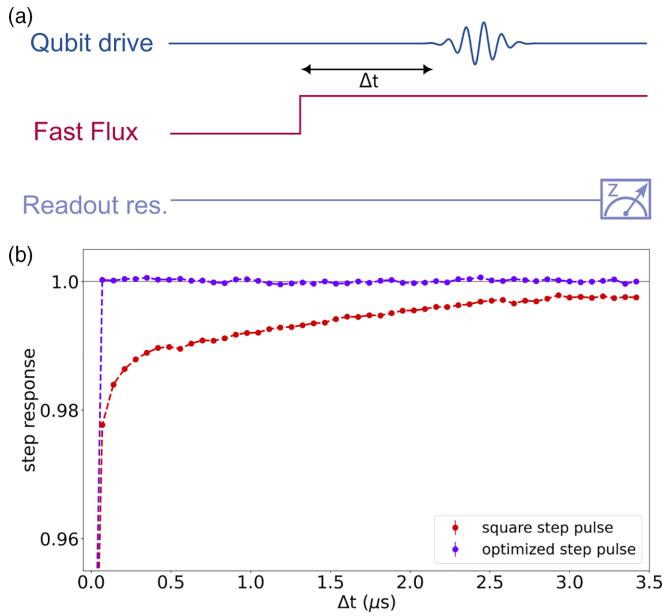


FIG. 7. Procedure for calibrating for long-time step-pulse distortions. (a) Pulse sequence consisting of a fast-flux step pulse, a qubit drive pulse with variable delay, and measurement of the qubit state via the coupled resonator. By sweeping the qubit drive frequency, the qubit frequency can be fitted for different delay times. (b) Qubit response to a square step pulse (red) and optimized step pulse (purple). Step response of 1 is normalized by dividing the fitted and target qubit frequencies.

Using the calibration from these two schemes, we define a pre-distorted pulse that is a sum of four exponentials, where

the short time and long time calibrations each define two of these terms. We find that including additional terms did not noticeably improve the pulse performance. The amplitude of the pulse normalized to the amplitude of a square step pulse is shown in Fig. 8(a). The pre-distorted pulse overshoots the final amplitude to rapidly bring the qubit onto resonance; this amplitude remains slightly elevated for long times which stabilizes the qubit frequency for the duration of the experiment.

We evaluated the calibrated pulses by running a time evolution experiment between two qubits. Starting with the qubits detuned from each other, we initialized one qubit on either the first or second excited state. The qubits were then placed onto resonance using the fast flux pulses. Resonant qubit dynamics are characterized by full contrast oscillations: the excitations “swap” between the qubits at a rate given by their coupling strength with full population transfer. When qubits are detuned from each other, the frequency of swaps increases but the excitations do not have complete state transfer.

In Figs. 8(b) and 8(c), we plot the population of the qubit with an initial excitation as a function of a fast-flux amplitude during the time evolution section of the experiment. We first detune one of the qubits by 200 MHz for state initialization. Using fast-flux pulses, we step the qubit to different fast-flux amplitudes shown in the y axis for varying times. We note that we biased the qubits frequencies to be resonant when the fast-flux amplitude is close to zero. The upper panels show the swaps when a single excitation is initialized on one qubit, while the lower panels has two excitations initialized on one qubit. Two excitations swap via a second-order process yielding a significantly slower oscillation frequency. The swaps using the square step pulse [Fig. 8(b)] show that the slowest frequency oscillations—when qubit frequencies

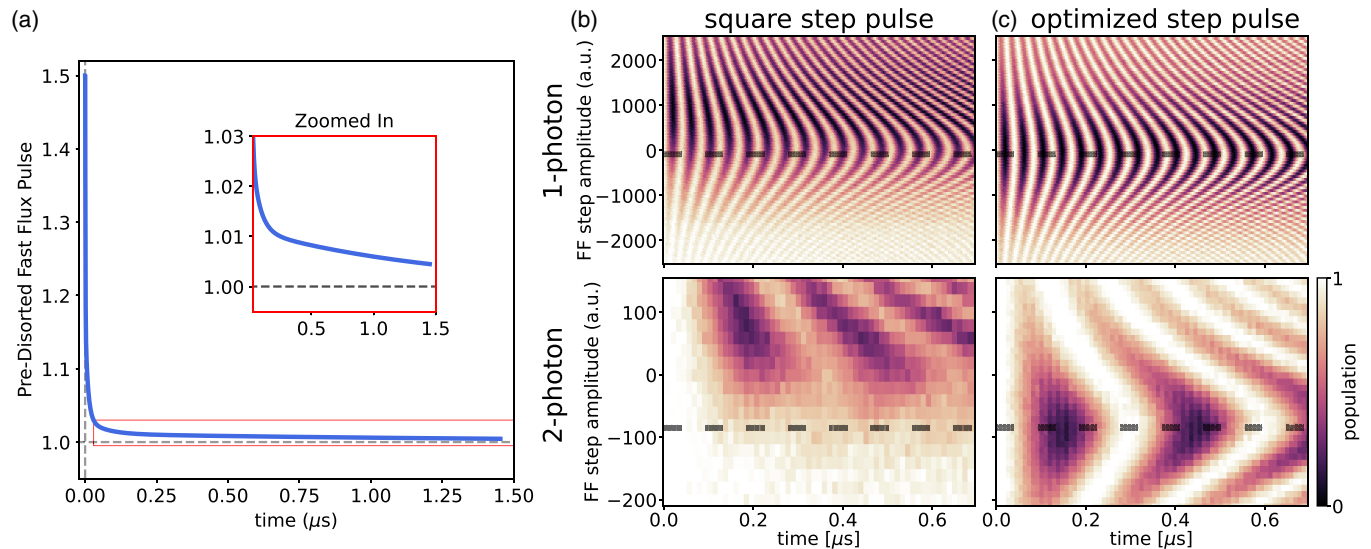


FIG. 8. Optimized step pulse and two-qubit swaps. (a) Amplitude of the pre-distorted step pulse amplitude that results in a more optimized step response when reaching the qubit, normalized by a square step pulse amplitude. Inset shows a zoom-in of the amplitude dominated by long-time distortions. (b) (Top) Swapping of qubit excitation between two qubits, where the one qubit is initialized in the first excited state. After initialization, a step function is applied to the fast-flux line of that qubit with varying amplitude (“FF step amplitude”, y axis) to bring it onto resonance with the second qubit that is biased there. Precalibration, the swaps are not complete and they are not symmetric with respect to the sign of the FF step amplitude. (lower) Same, but for a swap of the second excited state. (c) Same as (b), but with the calibrated fast-flux pulse. In (b) and (c), dashed line indicates the fast flux amplitude where the two qubits should be resonant and serves as a guide for the eye. Figures modified from the supplement of Ref. [21] with the permission of the authors.

are resonant—vary as a function of time, indicating frequency drifts during the experiment. Additionally, the maximum population contrast does not overlap with the slowest frequency oscillations, indicating that the qubit does not reach the target frequency quickly enough. These discrepancies are significantly improved with the optimized step pulse [Fig. 8(c)].

The 145 ps time resolution of the QICK allowed us to characterize the time distortions in fast flux pulses necessary for state initialization as well as starting and freezing dynamics of multiqubit systems. Using this analysis, we compensated for these distortions using arbitrary waveform generation enabled with the board. The high temporal resolution of QICK can compensate for experiments with larger coupling rates and stricter requirements which will be crucial in our future experiments. Finally, we note that one section of the project required measuring correlations between multiple qubits. Doing so utilized the multiplexed readout enabled by QICK (Sec II A) to measure the states of two qubits simultaneously.

V. PHASE-COHERENT PARAMETRIC QUANTUM OPERATION

Parametrically induced interactions are widely used for constructing quantum gates in superconducting quantum systems. In these systems, a central nonlinear coupler is typically coupled to one or more logical quantum modes such as qubits or cavities [35–39]. The desired gate dynamics between logical modes are activated by pumping the coupler mode at specific off-resonance frequencies that create the desired Hamiltonian from a higher-order (typically third- or fourth-order nonlinearity). Such parametric processes have the advantage that they can be easily applied among multiple fixed-frequency modes, with the frequency of the drive determining which parametric interaction is activated. However, during the parametric interaction, the parametric pumps also imprint extra phases onto the quantum states, similarly to the manner in which changing the phase of a resonant qubit drive can alter the angle over which the qubit is rotated. Consequently, most parametric gates possess extra phase elements tied to the phase of the parametric pump, which places strong phase coherence requirements on the control electronics that generate the pumps.

The Hatlab at the University of Pittsburgh has been conducting experiments on parametrically driven quantum state routers [37] and quantum modules [40,41]. Figure 9(a) shows the coupling schematic of our four-qubit quantum module. This device consists of four transmon qubits ($Q_1 - Q_4$, with associated photon annihilation operators q_{1-4}) that are coupled to a central Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) coupler (S , with annihilation operator s) [42]. The SNAIL mode can be flux biased to a point with a third-order nonlinearity and negligible fourth-order nonlinearity, providing us with the initial nonlinear term in the system Hamiltonian: $g_3(s + s^\dagger)^3$, where g_3 is the strength of the SNAIL's third-order nonlinearity. As a result of the dispersive coupling between the SNAIL mode and each qubit mode, the third-order nonlinearity is shared among the qubits, leading to all possible third-order terms of the three modes. For the purposes of creating parametric exchange gates, we focus on the subset of terms which form the interaction Hamil-




FIG. 9. Coupling scheme and device schematic of a parametrically controlled four-qubit module. (a) Coupling scheme of the four-qubit module. A central SNAIL mode (S) is coupled to four transmon qubit modes $Q_1 - Q_4$, each qubit is also dispersively coupled to a readout resonator mode $R_1 - R_4$. (b) Device schematic. A central aluminum tube for hosting a SNAIL chip is connected to four perpendicular channels, each is designed to host one qubit chip with strip line readout resonators. In the parametric measurement experiment described in the text, only one qubit chip and a SNAIL chip were used.

tonian

$$H_{\text{int}}/\hbar = \sum_{i \neq j} g_{ijs} (q_i^\dagger q_j s^\dagger + q_i q_j^\dagger s), \quad (3)$$

where g_{ijs} is the three-wave-mixing coefficient among qubit modes i, j and the SNAIL mode. Based on these terms, we can selectively activate the desired exchange interaction between two qubit modes, say Q_1 and Q_2 , by pumping the SNAIL mode at the frequency $f_p = f_2 - f_1 + \delta$, where f_i is the mode frequency of Q_i , δ is the pump detuning added due to the relative AC Stark shift between qubits 1 and 2, and we assume $f_2 > f_1$. For a given pump phase ϕ_p , the effective Hamiltonian can be written as

$$H_{\text{eff}}/\hbar = g_2^{\text{eff}} (q_1^\dagger q_2 e^{i\phi_p} + q_1 q_2^\dagger e^{-i\phi_p}), \quad (4)$$

where $g_2^{\text{eff}} = g_{ijs} \sqrt{n_p}$ is the effective exchange rate and n_p is the pump strength expressed in units of photons. Such exchange interaction between two qubits creates a continuous family of gates, described $i\text{SWAP}^\alpha$, which can be used to implement a variety of universal two-qubit gates in quantum circuits.

In the above gate operation, the impact of the pump phase can be understood by considering the following scenario: Q_1 is initially prepared to a superposition state with initial phase ϕ_1 , while Q_2 is in the ground state, i.e. $|\psi_0\rangle = 1/\sqrt{2}(|g\rangle + e^{i\phi_1}|e\rangle) \otimes |g\rangle$. After a complete $i\text{SWAP}$ gate with pump phase ϕ_p , the final state of the qubits will become $|\psi\rangle = |g\rangle \otimes 1/\sqrt{2}(|g\rangle + e^{i(\phi_1 + \phi_p)}|e\rangle)$. To check the phase of the resulting state, we apply a $\pi/2$ pulse with phase ϕ_2 on Q_2 and measure its σ_z expectation value. The measurement result will then be $\langle \sigma_{z,2} \rangle = -\cos(\phi_1 + \phi_p - \phi_2)$. Note that each of the three phases involved here are originally carried by the three microwave drive channels at frequencies f_1, f_p , and f_2 , with the qubit drive frequencies $f_1, f_2 \approx 3-6$ GHz, and the pump frequency $f_p \approx 0.5-2$ GHz. Thus, to ensure that the measurement results stay the same between repetitions; the three drive channels must stay phase coherent with each other under the constraint

$$[\phi_1(t) + \phi_p(t) - \phi_2(t)]|_{t=T_N} = \text{constant}, \quad (5)$$

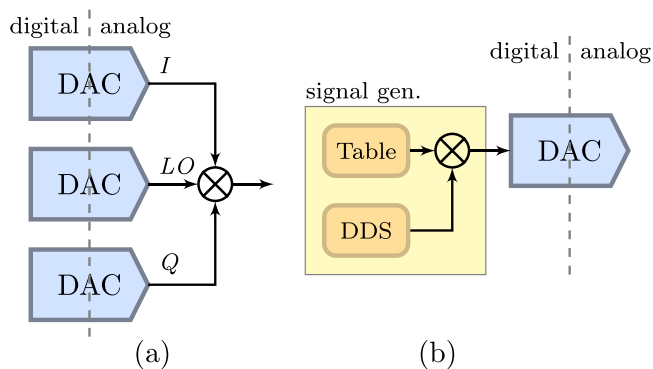


FIG. 10. Typical circuits for synthesising a microwave drive for superconducting qubit control. The possible common subsequent components (e.g., amplifiers, dividers) are omitted. (a) Analog mixer based up-conversion circuit. Two DAC channels are connected to the I and Q ports of an analog mixer, whose local oscillator (LO) is provided using a continuous signal generator with a third DAC output. (b) Direct pulse generation with DDS. Complex IQ mixing is performed in the digital domain, and one DAC output provides the required pulse. Here, only an extra low- or band- pass filter is needed after the DDS DAC to filter out the image tones in other Nyquist zones.

where t_N is the start time of experiment repetition N .

This constraint can be satisfied with different methods, depending on the microwave control circuit setup being used. Figure 10(a) shows an analog up-conversion circuit that is typically used to generate control pulses for superconducting qubits. In this setup, two output channels of a relatively low frequency AWG (typically with bandwidth of ~ 500 MHz) are connected to the I and Q ports of an analog IQ mixer, whose local oscillator (LO) is provided by a high frequency signal generator. In this setup, the phase of each AWG pulse can be set to arbitrary value ϕ^I , while the generator phase is continuously evolving and cannot be reset in real time. The total phase of each drive channel now comes from two parts, i.e.,

$$\phi(t) = 2\pi * \int_{-\infty}^t f^L(t) dt + \phi^I, \quad (6)$$

where $f^L(t)$ is the frequency of the signal generator. To satisfy the phase coherence constraint in Eq. (5) for ideal, stable generators, we can choose the generator frequencies (f_i^L) such that $f_1^L + f_p^L - f_2^L = 0$ to remove the absolute time dependent part of the total phase, and then choose the AWG pulse frequencies (f_i^I) accordingly, such that $f_i^I = f_i - f_i^L$. However, for experiments that requires averaging over long time periods, this constraint places a strong requirement on the relative stability between all the signal generators and the AWG channels, which may be challenging for conventional microwave electronics. Although small-scale experiments may circumvent this issue by generating LOs for pump signals through analog mixing of the qubit drives' LOs, such an approach significantly complicates the control circuit for large-scale quantum processors.

Alternately, QICK offers a much simplified way to generate the control pulses. As in Fig. 10(b), the pulse for each control channel can be generated using DDS, eliminating the need for additional analog mixer and generator hardware.

Crucially, the QICK firmware offers the functionality of resetting the phase of all the DDS channels simultaneously at well-defined times. Thus the absolute time dependent parts of the drive phases in Eq. (5) are completely removed, which makes the phase coherence condition automatically satisfied. Since all the drive phases and frequencies are defined digitally on a single FPGA chip, even if the absolute frequency of each channel drifts over time, their relative phases will always stay locked, which also makes long-term relative phase stability easily realizable. Furthermore, as needed, one can directly source a LO channel using the ZCU216 board, enabling coherent reset.

To have full control over the device described in Fig. 9(a), we will need one qubit drive channel plus one readout resonator drive channel for each qubit, and there could be 6 possible SNAIL drive frequencies for the parametric two-qubit gates. With the limited instantaneous bandwidth available in low frequency AWGs [as in Fig. 10(a)], all these drive channels will need to be synthesized individually, which requires a total of 14 signal generators plus IQ mixers, and 28 AWG channels. With a high-frequency DDS, the whole device can be controlled with only 14 DAC channels on a single ZCU216 board. This can be even further simplified with frequency-multiplexed digital generators. Therefore the DDS pulse generation provided by the QICK-controlled RF-SoC boards not only provides stable phase coherent control, but also significantly reduces the complexity of microwave hardware setup.

As an example demonstration of the phase stability of QICK-controlled parametric quantum operation, we present the results from a parametric readout experiment conducted in a subset of our four-qubit module. For this particular experiment, only one qubit and a SNAIL are used, as depicted in Fig. 9(b).

The parametric measurement process for this experiment is activated by applying two microwave pumps on the SNAIL mode at the difference and sum frequencies between the SNAIL and qubit mode. The difference frequency pump will activate the photon exchange (conversion) interactions between the qubit and SNAIL mode, similar to the *i*SWAP interaction described earlier with Eq. (4). The sum frequency pump will turn on the two-photon transition (gain) process [43] that, instead of exchanging photons, will jointly excite or destroy photons in the two modes. By turning on both pumps simultaneously and tuning their relative amplitudes to have matched interaction strength, we can activate the effective Hamiltonian:

$$H_{\text{eff}}/\hbar = g(q^\dagger s e^{i\phi_c} + q s^\dagger e^{-i\phi_c} + q^\dagger s^\dagger e^{-i\phi_g} + q s e^{i\phi_g}) \\ = g(\cos(\phi_m)\sigma_x - \sin(\phi_m)\sigma_y)(s^\dagger e^{-i\phi_s} + s e^{i\phi_s}), \quad (7)$$

where ϕ_g and ϕ_c are the phases of the gain and conversion pumps respectively, and $\phi_m = (\phi_g - \phi_c)/2$, $\phi_s = (\phi_g + \phi_c)/2$. By controlling the two pump phases ϕ_g and ϕ_c , the qubit can be projected along an arbitrary direction on the XY plane of the Bloch sphere, and creates a qubit state dependent coherent state in the SNAIL mode. This process effectively performs a measurement on the qubit along the projection axis, and the measured result can be collected by demodulating the signal coming out from the SNAIL port.

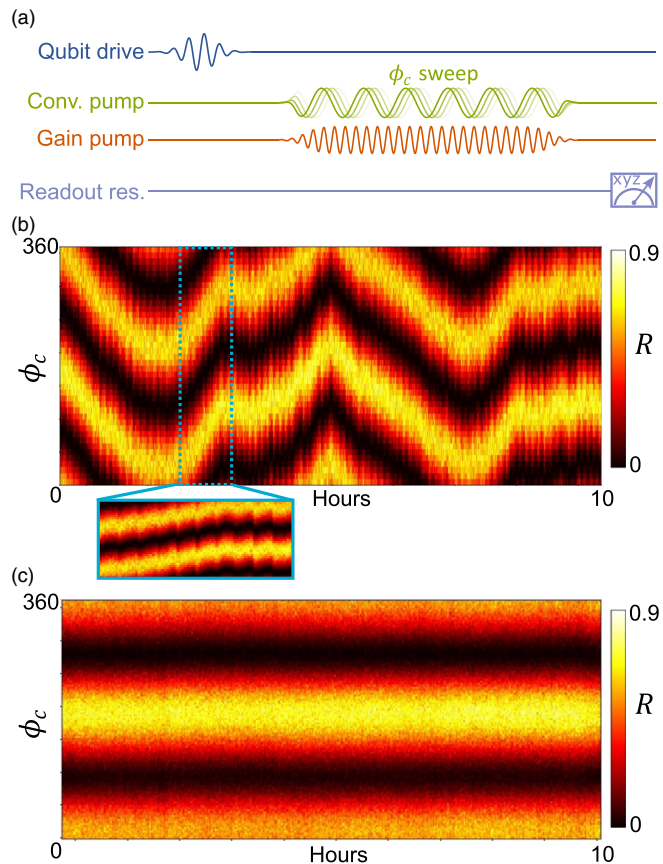


FIG. 11. Pulse sequence and long-time phase stability data for calibrating the measurement direction in the parametric measurement experiment. (a) Pulse sequence for parametric measurement direction calibration. The qubit is first prepared in a $|X_+\rangle$ state with a $\pi/2$ pulse, then the gain and conversion pumps are applied simultaneously on the SNAIL mode to perform the parametric measurement. Finally, conventional state tomography is performed on the qubit using the readout resonator. The phase of the conversion pump (ϕ_c) is swept to align the parametric measurement direction with the $|X_+\rangle$ direction. (b) Phase stability data obtained from running the ϕ_c sweep experiment using a mixer-based control circuit for a duration of 10 hours. The inset shows zoomed-in data over one hour, which exhibits a 6 minute period drift caused by imperfections in the phase locked loop of the signal generators, while long-time data also show a larger overall phase drift over several hours due to temperature fluctuations. In contrast, (c) show the same phase stability test performed using a QICK-controlled RFSoc board in which the phase remained stable throughout the entire 10-hour experiment.

Figure 11(a) shows the pulse sequence used to calibrate the direction of the parametric measurement. The qubit is first prepared to the $|X_+\rangle = (|g\rangle + |e\rangle)/\sqrt{2}$, then the gain and conversion pumps are applied simultaneously. After the pumps are off, a full qubit state tomography is performed using the readout resonator. When the parametric measurement direction is aligned with the $|X_+\rangle$ or $|X_-\rangle = (|g\rangle - |e\rangle)/\sqrt{2}$ direction, the length of the qubit Bloch vector $R = \sqrt{\langle\sigma_x\rangle^2 + \langle\sigma_y\rangle^2 + \langle\sigma_z\rangle^2}$ will be preserved during the parametric measurement. On the other hand, when the qubit is prepared in the perpendicular states $|Y_{\pm}\rangle = (|g\rangle \pm i|e\rangle)/\sqrt{2}$, the qubit will collapse to the mixed state $1/2 |X_+\rangle \langle X_+| +$

$1/2 |X_-\rangle \langle X_-|$ with $R = 0$ after the measurement. The measurement direction is simply calibrated by sweeping the phase of one of the pumps, say ϕ_c . Similarly to the *i*SWAP case, here the qubit drive and both pump channels must stay phase coherent under a slightly different constraint:

$$[\phi_q(t) - (\phi_g(t) - \phi_c(t))/2]_{t=t_N} = \text{constant}, \quad (8)$$

where ϕ_q is the phase of the qubit drive pulse. Moreover, the relative phase must stay stable over a long time for experiments that requires a large number of average repetitions, e.g., measurement efficiency characterization [44]. Figure 11(b) shows the result of running the pump phase calibration experiment for 10 hours using an analog mixer-based setup. Although all the signal generators and AWG channels used were locked to an external 10 MHz Rubidium clock, we still observed a ~ 6 minute period phase drift due to the imperfections of the phase-lock-loop in the signal generators, as well as an overall slow drift we attribute to changes in room temperatures over time [45,46]. Such drift makes all experimental studies extremely tedious, as re-calibration is required every few minutes. More importantly, as the advancements of longer coherence qubits [26,47–49] and the emergence of novel quantum error correction schemes [50–52] gradually allowing the execution of longer quantum circuits, it is becoming more crucial that the phase drifts in classical control electronics do not set a limit on the duration of quantum circuits we aim to execute. In contrast, Fig. 11(c) shows the same experiment performed using a QICK-controlled RFSoc board, where all the drives were generated using direct digital synthesis with phase reset on each repetition. The measurement phase remained stable over the entire 10-hour experiment, providing clear evidence of the relative phase stability of the RFSoc and its excellent suitability for parametric quantum operations.

In general, the relative phase coherence between different drive channels is essential for performing parametric quantum operations. The requirements for phase coherence can vary depending on the specific parametric process being implemented, and can be expressed as specific constraints on the relative phases between the drive channels [e.g., Eqs. (5) and (8)]. While it is possible to construct custom analog interferometer circuits with high-stability microwave generators to meet these constraints, using DDS-based pulse generation with a well-defined phase reset function is a more efficient and convenient solution. Because it offers precise control over the frequencies and phases of the generated signals, and also greatly reduces hardware requirements as well as the complexity of experimental setup. Note that the phase reset functionality available in the QICK firmware is vital here as it erases the historically accumulated phases, thereby eliminating the need for long-term high-precision phase stability in control electronics, which could be hard to realize.

VI. Phase-sensitive parametric entangling gates

In a Schuster laboratory experiment conducted at the University of Chicago, we used a tunable, fluxonium-like coupler [53] to galvanically couple two low-frequency, heavy-fluxonium qubits [26,54]. A key feature of our design is that for a particular DC bias, the tunable coupler has a zero-coupling “off-position,” where the *XX*, *ZZ* qubit-qubit




FIG. 12. Circuit diagram of the inductively-coupled fluxonium qubits experiment. Qubit A is in red, qubit B is in blue, and the tunable coupler is in green. We use long (> 200) chains of Josephson junctions to realize the kinetic inductance in our qubit, and have shorter, shared chains to provide galvanic coupling between each individual fluxonium with the coupler. Each loop (for the qubits as well as for the coupler) has a dedicated flux bias line, for both DC biasing and low-frequency RF drives.

coupling terms are nearly nulled. A circuit diagram is shown in Fig. 12, and the effective Hamiltonian can be written as

$$H_{\text{eff}} = - \sum_{\mu=a,b} \frac{\omega_{\mu}}{2} \sigma_z^{\mu} - \Omega_{\mu} \sigma_x^{\mu} + J \sigma_x^a \sigma_x^b + \zeta \sigma_z^a \sigma_z^b, \quad (9)$$

where $\sigma_x^{\mu}, \sigma_z^{\mu}$ are the qubit Pauli operators in the basis of symmetric and antisymmetric wavefunctions. When biased at the coupler “off-position,” experimental data confirm that the residual ZZ couplings is less than 100 Hz. At that DC bias, we dynamically activate the XX coupling by RF driving at either the sum or difference of the qubit frequencies.

When galvanically coupling our fluxonium qubits, we have both intrinsic crosstalk due to the construction of our circuit, as well as a significant geometric crosstalk as the qubit loops are adjacent to one another. We measure this DC crosstalk matrix to have up to 25% crosstalk between flux bias lines. To realize a parametric interaction between our qubits using the coupler, we perform an RF drive at the sum and difference of the qubit frequencies. However, this effectively induces a detuned AC-Stark shift on each individual qubit. Because the optimal drive amplitude of the parametric oscillation is dependent on the effective qubit frequencies, such AC-Stark shifts change the frequency of the coupler drive at which there is maximum Rabi amplitude contrast. From other single-qubit measurements, we determine that the bare frequencies of our qubits are $\omega_a = 48.4$ and $\omega_b = 61.8$ MHz, meaning that the optimal coupler drive for the bSWAP interaction [53] should be at $\omega_a + \omega_b = 110.2$ MHz. However, when performing a frequency-length Rabi experiment using the coupler drive, we see that the center of this Rabi chevron has been shifted down by ≈ 7 MHz, see Fig. 13(a).

In general, such crosstalk can be canceled using the additional degrees of freedom afforded by having individual flux lines for each qubit loop. We implement this correction with simultaneous flux crosstalk cancellation pulses, played at the same time as driving our coupler [55]. Along the $\Phi_{\text{ext},a}, \Phi_{\text{ext},b}$ bias lines, we use pulses with the same shape, frequency, and length as our two-qubit gate drive but with calibrated amplitudes and relative phase offsets. Because these compensation parameters are independent to the coupler drive, we can measure the Rabi oscillation of the coupler while sweeping the phase and amplitude of our compensation pulses for each parameter of both qubit A and B’s flux bias lines. We show




FIG. 13. Compensation pulses showing a RF crosstalk calibration and reduction (a) Rabi measurement with variable length pulses of the coupler drive, in the neighborhood of the sum frequency $\omega_a + \omega_b$ of the qubits. The colorscale here indicates the population present in qubit A. In the absence of crosstalk cancellation pulses, the maximum Rabi amplitude oscillation is 7 MHz detuned from the bare qubit frequencies due to a large AC-Stark shift. (b) The calibration sweep for one cancellation pulse, where the relative phase between the qubit drive and the coupler drive is adjusted. One can see the maximum contrast when the compensation pulse has roughly the opposite phase of the coupler drive. (c) Same Rabi measurement as (a), but with applied compensation pulses on both $\Phi_{\text{ext},a}, \Phi_{\text{ext},b}$. Note that the maximum Rabi amplitude oscillation is now centered at exactly 110.2 MHz, which is what we expect given our single qubit measurements. Data for figures modified from the supplement of Ref. [26] with the permission of the authors.

the phase sweep of a cancellation pulse in Fig. 13(b), revealing that the optimal relative phase between the cancellation channels and the coupler channel is $\approx 180^\circ$.

By applying the cancellation pulses and performing the same Rabi experiment as before, we see oscillations between $|00\rangle$ and $|11\rangle$ with maximum Rabi contrast at exactly the sum of qubit frequencies $\omega_a + \omega_b$, see Fig. 13(c). This indicates that there is no more off-resonant drive on the qubits and that our compensation pulses effectively eliminated flux crosstalk from the coupler flux drive.

With our calibrated crosstalk cancellation pulses and our ability to realize generic XX interactions using the RF drive on our coupler, we can realize a $|gg\rangle \leftrightarrow |ee\rangle$ oscillation. The Hermitian matrix for such a gate can be written as

$$\sqrt{\phi_b \text{SWAP}} = \begin{pmatrix} \cos \theta & 0 & 0 & i e^{i\phi_D} \sin \theta \\ 0 & e^{i\phi_{01}} & 0 & 0 \\ 0 & 0 & e^{i\phi_{10}} & 0 \\ i e^{i(\phi_{11} - \phi_D)} \sin \theta & 0 & 0 & e^{i\phi_{11}} \cos \theta \end{pmatrix}, \quad (10)$$

where ϕ_D is the phase of the coupler drive, and $\phi_{01}, \phi_{10}, \phi_{11}$ are phases due to the frequency shift of levels while the drive is on, which have relationship $\phi_{11} = \phi_{01} + \phi_{10} + \phi_{zz}$.




FIG. 14. Calibration of coupler and single qubit phase in preparation of a \sqrt{b} SWAP gate (a) The pulse sequence used for calibration. This sequence is constructed to measure the drive induced phase ϕ_{11} , and each block is repeated 402 times in order to amplify small errors. (b) The measured phase relationship while sweeping ϕ_A phase gates (Z_A). Notice that this sequence amplifiers errors, such that a change of $<1^\circ$ of ϕ_A results in a 10% change in the state population at the maximally sensitive locations. We find that the optimal phase angle is at 357° , corresponding to $\phi_{01} + \phi_{10} = 3^\circ$. Figures modified from the supplement of Ref. [26] with the permission of the authors.

From this matrix, we can see that control of the coupler drive phase, along with the application of single qubit Z gates, allows us to exactly realize a \sqrt{b} SWAP gate. (A similar matrix can be written for the \sqrt{i} SWAP oscillation). Furthermore, we recognize that both the phase of individual qubits and the phase of the coupler are set by the individual DAC channels. Thus the relative phase between our channels directly leads to an effective rotation angle within the $fSim$ family of gates. Conversely, if there was noise in either the absolute or relative phases of our channels, we would not be able to realize high fidelity parametric gates.

In order to execute pure \sqrt{i} SWAP and \sqrt{b} SWAP gates, we need consistent, individual control of the relative phases of the various DAC channels, which the QICK enables. To determine the appropriate phases with very high accuracy, we amplify the error and make measurements. The sequence that is used to do this is shown in Fig. 14(a), inspired by similar protocols in Ref. [56]. We play a series of \sqrt{b} SWAP gates with single-qubit phase (Z) gates interleaved between each pair of them, and sweep the rotation angle of all the phase gates (Z_A) simultaneously. When we play $4n + 2$ \sqrt{b} SWAP gates, we can only get a full π rotation when $\phi_A + \phi_B = -\phi_{11}$. Thus we measure the drive induced phase ϕ_{11} with a ϕ_A sweep, as shown in Fig. 14(b). Using similar error amplification sequences, we can also measure the drive induced phase ϕ_{01} and ϕ_{10} . We compensate these phases by setting qubit A Z gate phase $\phi_A = -\phi_{10}$, qubit B Z gate phase $\phi_B = -\phi_{01}$, and obtain a \sqrt{b} SWAP gate with correct phases.

After calibrating our \sqrt{b} SWAP gate, we use it to prepare a Bell state $(|gg\rangle + i|ee\rangle)/\sqrt{2}$ (Fig. 15). To determine the state preparation fidelity, we choose to perform quantum state tomography (QST). This involves the measurement of our state in an overcomplete set of basis states, and subsequently performing maximal likelihood estimation to find the most likely




FIG. 15. The density matrix of a Bell state prepared with our parametric gate, found with quantum state tomography (QST). (a) We initialize our qubit to the gg state and perform a single application of the \sqrt{b} SWAP gate. This prepares the Bell state, $(|gg\rangle + i|ee\rangle)/\sqrt{2}$. Since performing QST requires averaging over thousands of measurements in different basis, a pure two-qubit density matrix implies a stable phase relation between QICK DAC channels. We measure a purity and state fidelity that are limited primarily by state preparation and measurement errors to be 95%. (b) The density matrix for the ideal Bell state. Figures modified from the supplement of Ref. [26] with the permission of the authors.

corresponding density matrix (subject to physical constraints). In this experiment, we average 50 000 times for measurement along each of the nine basis vectors. If there were unstable relative phase between the QICK DAC channels, this would be revealed as a decrease in the purity of density matrix. In the worst-case scenario, where a completely random relative phase exists between control channels for each shot, this would result in a completely mixed density matrix. However, by using the QICK's phase reset functionality to synchronize the phases between the QICK DAC channels, we find a state purity of 95%, indicating consistent phase preparation.

Using these compensation pulses, we can perform full benchmarking of the \sqrt{b} SWAP gate, with a continuous cross-entropy benchmarking result over the course of 25 hours. Here, we show the variance in gate fidelity over this period in Fig. 16. During this period, no adjustments were made to any of the pulse parameters, nor to the crosstalk cancellation parameters. We know that our gate is sensitive to phase [see Fig. 14(b)] and frequency [see Fig. 13(c)], while amplitude errors would correspond to under- or over-rotation of the gate. Therefore the consistency of the cross-entropy benchmarking results can be used as a metric to quantify the stability of these three parameters between our QICK DAC channels.

We find that not only does the cross-entropy benchmarking show a very high two-qubit gate fidelity, where the fidelity of the \sqrt{b} SWAP gate is greater than 99.9%, we find that there are very few fluctuations in the gate fidelity over the span of 25 hours. As shown in Fig. 16, each individual cross-entropy




FIG. 16. Variance in cross entropy benchmarking over 20 consecutive hours. To benchmark our gate, we execute continuous cross entropy benchmarking sequences on our system. Each individual sequence is up to 600 layers of randomly selected two-qubit Clifford gates, interleaved with the reference \sqrt{b} SWAP gate. For each cross-entropy benchmark sequence, we find an average fluctuation in gate fidelity to be $3\sigma \approx 0.02\%$. Throughout the entire sequence, the average fidelity measured did not drift by more than $\pm 0.03\%$. Data used for figures modified from Ref. [26] with the permission of the authors.

benchmarking sequence has a statistical fit infidelity on the order of 0.02%. Furthermore, there is no significant drift in the average gate fidelity, with a maximum fluctuation of no more than 0.03%. In our error analysis, we attribute the majority of the gate infidelity to qubit relaxation channels, not through coherent errors that can be attributed to the QICK DAC channels.

VII. CONCLUSION

The experimental use cases presented here show how the QICK can be applied to control tasks relevant to superconducting qubit systems. Multiplexed readout is used in virtually all scaled-up planar superconducting qubit systems, including those from academic labs [51,57]. Mixer-free readout simplifies the control loop and can extend the control benefits gained from direct digital synthesis [58]. Predistorted fast flux pulses are useful for any control scheme that includes flux-tunable qubits or couplers [54,59]. Parametric quantum operations are generally advantageous in qubit architectures as they simultaneously allow large coupling strengths and large qubit-drive detunings [46,60].

Ongoing work on the QICK is branching into several directions. The QICK team has developed new laboratory tools running on the RFSoc to improve the control and characterization of superconducting hardware [61,62]. The team is also

developing QICK for different qubit platforms such as atomic qubits, spin qubits, and color centers. Preliminary work used the standard QICK firmware [63,64], but the QICK team has since begun creating firmware tailored for these new platforms.

Longer term work on the QICK includes integrating a redesigned timed-processor that will improve the QICK's speed and modularity. This new timed-processor will allow the QICK to scale up to multiboard systems.

ACKNOWLEDGMENTS

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics, with support from its QUANTISED program and from National Quantum Information Science Research Centers, Quantum Science Center. This work was funded in part by EPiQC, an NSF Expedition in Computing, under Grant No. CCF-1730449. The contribution of C.Z. and M.H. was supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under Contract Number DE-SC0012704. The work of J.M. was supported by the NSF Quantum Leap Challenge Institute for Robust Quantum Simulation 2120757 and ARO MURI W911NF-15-1-0397. J.M. also acknowledges support from NSF GRFP DGE-2039656. The contributions of S.L., C.M., and D.I.S. were supported by the Air Force Office of Scientific Research (AFOSR) Multidisciplinary Research Program of the University Research Initiative (MURI) Grant No. W911NF2010177. C.M. also acknowledges support from NSF GRFP Grant No. DGE-2146755. S.L., C.M., and D.S. also thank MIT Lincoln Laboratories for device fabrication and packaging. The contributions of C.D., H.Z., and D.I.S. were supported by the Army Research Office under Grant No. W911NF1910016, and was partially supported by the University of Chicago Materials Research Science and Engineering Center, which is funded by the National Science Foundation under award number DMR-1420709. Devices were fabricated in the Pritzker Nanofabrication Facility at the University of Chicago, which receives support from Soft and Hybrid Nanotechnology Experimental (SHyNE) Resource (NSF ECCS-1542205), a node of the National Science Foundation's National Nanotechnology Coordinated Infrastructure. The authors thank National Quantum Information Science Research Centers Q-NEXT under Contract No. DE-AC02-06CH11357, SQMS under Contract No. DE-AC02-07CH11359, and C2QA members who participated in discussions.

- [1] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, *J. Appl. Phys.* **132**, 160902 (2022).
- [2] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton *et al.*, Suppressing

quantum errors by scaling a surface code logical qubit, *Nature (London)* **614**, 676 (2023).

- [3] Zynq UltraScale+ RFSoc, <https://www.xilinx.com/products/silicon-devices/soc/rfsoc.html> (2023).
- [4] L. Crockett, D. Northcote, and R. W. Stewart, Software Defined Radio with the Zynq UltraScale+ RFSoc (2023).
- [5] R. Gebauer, N. Karcher, and O. Sander, A modular rfsoc-based approach to interface superconducting quantum bits, in *2021*

- International Conference on Field-Programmable Technology (ICFPT)* (2021), pp. 1–9.
- [6] K. H. Park, Y. S. Yap, Y. P. Tan, C. Hufnagel, L. H. Nguyen, K. H. Lau, S. Efthymiou, S. Carrazza, R. P. Budoyo, and R. Dumke, ICARUS-Q: A scalable RFSoc-based control system for superconducting quantum computers, [arXiv:2112.02933](https://arxiv.org/abs/2112.02933).
- [7] Intermodulation Products Website, <https://intermodulation-products.com/products/microwave-platforms> (2021).
- [8] Y. Xu, G. Huang, N. Fruitwala, A. Rajagopala, R. K. Naik, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, Qubic 2.0: An extensible open-source qubit control system capable of mid-circuit measurement and feed-forward, [arXiv:2309.10333](https://arxiv.org/abs/2309.10333).
- [9] L. Stefanazzi, K. Treptow, N. Wilcer, C. Stoughton, C. Bradford, S. Uemura, S. Zorzetti, S. Montella, G. Cancelo, S. Sussman, A. Houck, S. Saxena, H. Araldi, A. Agrawal, H. Zhang, C. Ding, and D. I. Schuster, The QICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors, *Rev. Sci. Instrum.* **93**, 044709 (2022).
- [10] Quantum Machines Website, <https://www.quantum-machines.co/products/opx1000/> (2023).
- [11] Zurich Instruments Website, <https://www.zhinst.com/americas/en> (2021).
- [12] QICK GitHub repository, <https://github.com/openquantumhardware/qick> (2023).
- [13] QICK documentation website, <https://qick-docs.readthedocs.io/en/latest/> (2023).
- [14] Xilinx and Avnet ZCU111 evaluation board user guide, https://www.xilinx.com/support/documentation/boards_and_kits/zcu111/ug1271-zcu111-eval-bd.pdf (2020).
- [15] Xilinx and Avnet ZCU216 evaluation board user guide, https://www.xilinx.com/support/documentation/boards_and_kits/zcu216/ug1390-zcu216-eval-bd.pdf (2020).
- [16] Xilinx RFSoc-PYNQ website, <http://www.rfsoc-pynq.io/> (2023).
- [17] PYTHON productivity for Zynq, <https://pynq.io> (2023).
- [18] PYNQ - PYTHON productivity for Zynq - Downloadable PYNQ images, <http://www.pynq.io/board.html> (2023).
- [19] ZCU216 PYNQ image GitHub repository, <https://github.com/sarafs1926/ZCU216-PYNQ> (2023).
- [20] J. Bryon, D. K. Weiss, X. You, S. Sussman, X. Croot, Z. Huang, J. Koch, and A. A. Houck, Time-dependent magnetic flux in devices for circuit quantum electrodynamics, *Phys. Rev. Appl.* **19**, 034031 (2023).
- [21] J. G. C. Martinez, C. S. Chiu, B. M. Smitham, and A. A. Houck, Flat-band localization and interaction-induced delocalization of photons, *Sci. Adv.* **9**, ead71195 (2023).
- [22] S. Xie, L. Stefanazzi, C. Wang, C. Pena, R. Valivartha, L. Narvaez, G. Cancelo, K. Kapoor, B. Korzh, M. Shaw, P. Spentzouris, and M. Spiropulu, Entangled Photon Pair Source Demonstrator using the Quantum Instrumentation Control Kit System (2023).
- [23] A. Anferov, K.-H. Lee, F. Zhao, J. Simon, and D. I. Schuster, Improved coherence in optically-defined niobium trilayer junction qubits, [arXiv:2306.05883](https://arxiv.org/abs/2306.05883).
- [24] J. M. Cruz-Martinez, M. Robbiati, and S. Carrazza, Multi-variable integration with a variational quantum circuit, [arXiv:2308.05657](https://arxiv.org/abs/2308.05657).
- [25] S. Efthymiou, A. Orgaz-Fuertes, R. Carobene, J. Cereijo, A. Pasquale, S. Ramos-Calderer, S. Bordoni, D. Fuentes-Ruiz, A. Candido, E. Pedicillo, M. Robbiati, Y. P. Tan, J. Wilkens, I. Roth, J. I. Latorre, and S. Carrazza, Qibolab: an open-source hybrid quantum operating system, [arXiv:2308.06313](https://arxiv.org/abs/2308.06313).
- [26] H. Zhang, C. Ding, D. K. Weiss, Z. Huang, Y. Ma, C. Guinn, S. Sussman, S. P. Chitta, D. Chen, A. A. Houck, J. Koch, and D. I. Schuster, Tunable inductive coupler for high fidelity gates between fluxonium qubits, [arXiv:2309.05720](https://arxiv.org/abs/2309.05720).
- [27] M. Odeh, K. Godeneli, E. Li, R. Tangirala, H. Zhou, X. Zhang, Z.-H. Zhang, and A. Sipahigil, Non-markovian dynamics of a superconducting qubit in a phononic band gap, [arXiv:2312.01031](https://arxiv.org/abs/2312.01031).
- [28] A. Anferov, S. P. Harvey, F. Wan, J. Simon, and D. I. Schuster, Superconducting qubits above 20 GHz operating over 200 mK, [arXiv:2402.03031](https://arxiv.org/abs/2402.03031).
- [29] F. J. Harris, *Multirate Signal Processing for Communication Systems* (CRC Press, 2022).
- [30] L. H. Crockett, D. Northcote, and R. W. Stewart, eds., *Software Defined Radio with Zynq Ultrascale+ RFSoc* (Strathclyde Academic Media, 2023).
- [31] D. M. Joan Proakis, in *Digitel Signal Processing*, 3rd ed. (Prentice Hall, 1996), Chap. 9, pp. 742–746.
- [32] J. Vidal, R. Mosseri, and B. Douçot, Aharonov-bohm cages in two-dimensional structures, *Phys. Rev. Lett.* **81**, 5888 (1998).
- [33] J. Braumüller, A. H. Karamlou, Y. Yanay, B. Kannan, D. Kim, M. Kjaergaard, A. Melville, B. M. Niedzielski, Y. Sung, A. Vepsäläinen, R. Winik, J. L. Yoder, T. P. Orlando, S. Gustavsson, C. Tahan, and W. D. Oliver, Probing quantum information propagation with out-of-time-ordered correlators, *Nat. Phys.* **18**, 172 (2022).
- [34] R. Ma, B. Saxberg, C. Owens, N. Leung, Y. Lu, J. Simon, and D. I. Schuster, A dissipatively stabilized Mott insulator of photons, *Nature (London)* **566**, 51 (2019).
- [35] G. S. Paraoanu, Microwave-induced coupling of superconducting qubits, *Phys. Rev. B* **74**, 140504(R) (2006).
- [36] M. Reagor, C. B. Osborn, N. Tezak, A. Staley, G. Prawiroatmodjo, M. Scheer, N. Alidoust, E. A. Sete, N. Didier, M. P. da Silva *et al.*, Demonstration of universal parametric entangling gates on a multi-qubit lattice, *Sci. Adv.* **4**, eaao3603 (2018).
- [37] C. Zhou, P. Lu, M. Praquin, T. C. Chien, R. Kaufman, X. Cao, M. Xia, R. S. Mong, W. Pfaff, D. Pekker, and M. Hatridge, Realizing all-to-all couplings among detachable quantum modules using a microwave quantum state router, *npj Quantum Inf.* **9**, 54 (2023).
- [38] B. J. Chapman, S. J. de Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida *et al.*, A high on-off ratio beamsplitter interaction for gates on bosonically encoded qubits, [arXiv:2212.11929](https://arxiv.org/abs/2212.11929).
- [39] N. Goss, A. Morvan, B. Marinelli, B. K. Mitchell, L. B. Nguyen, R. K. Naik, L. Chen, C. Jünger, J. M. Kreikebaum, D. I. Santiago *et al.*, High-fidelity qutrit entangling gates for superconducting circuits, *Nat. Commun.* **13**, 7481 (2022).
- [40] E. McKinney, M. Xia, C. Zhou, P. Lu, M. Hatridge, and A. K. Jones, Co-designed architectures for modular superconducting quantum computers, in *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, 2023), pp. 759–772.
- [41] E. McKinney, C. Zhou, M. Xia, M. Hatridge, and A. K. Jones, Parallel driving for fast quantum computing under speed limits, [arXiv:2302.01252](https://arxiv.org/abs/2302.01252).

- [42] N. Frattini, U. Vool, S. Shankar, A. Narla, K. Sliwa, and M. Devoret, 3-wave mixing Josephson dipole element, *Appl. Phys. Lett.* **110**, 222603 (2017).
- [43] S. Poletto, J. M. Gambetta, S. T. Merkel, J. A. Smolin, J. M. Chow, A. D. Córcoles, G. A. Keefe, M. B. Rothwell, J. R. Rozen, D. W. Abraham, C. Rigetti, M. Steffen, Entanglement of two superconducting qubits in a waveguide cavity via monochromatic two-photon excitation, *Phys. Rev. Lett.* **109**, 240505 (2012).
- [44] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin *et al.*, Quantum back-action of an individual variable-strength measurement, *Science* **339**, 178 (2013).
- [45] H. Ball, W. D. Oliver, and M. J. Biercuk, The role of master clock stability in quantum information processing, *npj Quantum Inf.* **2**, 16033 (2016).
- [46] M. Xia, C. Zhou, C. Liu, P. Patel, X. Cao, P. Lu, B. Mesits, M. Mucci, D. Gorski, D. Pekker, and M. Hatridge, Fast superconducting qubit control with sub-harmonic drives, [arXiv:2306.10162](https://arxiv.org/abs/2306.10162).
- [47] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon *et al.*, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nat. Commun.* **12**, 1779 (2021).
- [48] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim *et al.*, High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler, *Phys. Rev. X* **13**, 031035 (2023).
- [49] S. Ganjam, Y. Wang, Y. Lu, A. Banerjee, C. U. Lei, L. Krayzman, K. Kisslinger, C. Zhou, R. Li, Y. Jia *et al.*, Surpassing millisecond coherence times in on-chip superconducting quantum memories by optimizing materials, processes, and circuit design, [arXiv:2308.15539](https://arxiv.org/abs/2308.15539).
- [50] V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. Brock, A. Ding, L. Frunzio *et al.*, Real-time quantum error correction beyond break-even, *Nature (London)* **616**, 50 (2023).
- [51] 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 (London)* **605**, 669 (2022).
- [52] K. S. Chou, T. Shemma, H. McCarrick, T.-C. Chien, J. D. Teoh, P. Winkel, A. Anderson, J. Chen, J. Curtis, S. J. de Graaf *et al.*, Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements, [arXiv:2307.03169](https://arxiv.org/abs/2307.03169).
- [53] D. K. Weiss, H. Zhang, C. Ding, Y. Ma, D. I. Schuster, and J. Koch, Fast high-fidelity gates for galvanically-coupled fluxonium qubits using strong flux modulation, *PRX Quantum* **3**, 040336 (2022).
- [54] H. Zhang, S. Chakram, T. Roy, N. Earnest, Y. Lu, Z. Huang, D. K. Weiss, J. Koch, and D. I. Schuster, Universal fast-flux control of a coherent, low-frequency qubit, *Phys. Rev. X* **11**, 011010 (2021).
- [55] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, *Phys. Rev. A* **93**, 060302(R) (2016).
- [56] M. Ganzhorn, G. Salis, D. J. Egger, A. Fuhrer, M. Mergenthaler, C. Müller, P. Müller, S. Paredes, M. Pechal, M. Werninghaus, and S. Filipp, Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform, *Phys. Rev. Res.* **2**, 033447 (2020).
- [57] B. Saxberg, A. Vrajitoarea, G. Roberts, M. G. Panetta, J. Simon, and D. I. Schuster, Disorder-assisted assembly of strongly correlated fluids of light, *Nature (London)* **612**, 435 (2022).
- [58] W. D. Kalfus, D. F. Lee, G. J. Ribeill, S. D. Fallek, A. Wagner, B. Donovan, D. Riste, and T. A. Ohki, High-fidelity control of superconducting qubits using direct microwave synthesis in higher nyquist zones, *IEEE Trans. Quantum Eng.* **1**, 1 (2021).
- [59] J. Stehlik, D. M. Zajac, D. L. Underwood, T. Phung, J. Blair, S. Carnevale, D. Klaus, G. A. Keefe, A. Carniol, M. Kumph, M. Steffen, and O. E. Dial, Tunable coupling architecture for fixed-frequency transmon superconducting qubits, *Phys. Rev. Lett.* **127**, 080505 (2021).
- [60] X. Y. Jin, K. Cicak, Z. Parrott, S. Kotler, F. Lecocq, J. Teufel, J. Aumentado, E. Kapit, and R. W. Simmonds, Versatile parametric coupling between two statically decoupled transmon qubits, [arXiv:2305.02907](https://arxiv.org/abs/2305.02907).
- [61] Real Life Applications of Polyphase Filter Banks, https://www.dsponlineconference.com/session/Real_Life_Applications_of_Polyphase_Filter_Banks (2023).
- [62] QICK-tools repo for the 4x2 board, https://github.com/leoeltipo/rfsoc4x2_public (2023).
- [63] QICK 2023 Workshop website, <https://indico.fnal.gov/event/57557/> (2023).
- [64] A. M. Mounce, J. D. Henshaw, L. Basso, R. Cong, P. Kehayias, J. M. Kitzman, J. Pollanen, F. F. Balakirev, and M. P. Lilly, Open-source FPGA control for nitrogen-vacancy quantum sensing, *Bull. Am. Phys. Soc.* (2023).