



PAPER

Quantum SWAP gate realized with CZ and iSWAP gates in a superconducting architecture

OPEN ACCESS

RECEIVED

20 December 2024

REVISED

16 April 2025

ACCEPTED FOR PUBLICATION

3 July 2025





















PUBLISHED

22 July 2025

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Christian Krizan^{1,*} , Janka Biznářová¹ , Liangyu Chen¹ , Emil Hogedal¹ , Amr Osman¹ ,
Christopher W Warren¹ , Sandoko Kosen¹ , Hang-Xi Li¹ , Tahereh Abad¹ , Anuj Aggarwal¹,
Marco Caputo² , Jorge Fernández-Pendás¹ , Akshay Gaikwad¹ , Leif Grönberg² ,
Andreas Nylander¹ , Robert Rehammar¹ , Marcus Rommel¹ , Olga I Yuzepovich¹ ,
Anton Frisk Kockum¹ , Joonas Govenius² , Giovanna Tancredi¹  and Jonas Bylander¹ 

¹ Department of Microtechnology and Nanoscience, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

² VTT Technical Research Centre of Finland Ltd, QTF Centre of Excellence, Espoo FI-02044 VTT, Finland

* Author to whom any correspondence should be addressed.

E-mail: krizan@chalmers.se and christian.krizan@gmail.com

Keywords: superconducting microwave devices, quantum information science, qubit

Abstract

It is advantageous for any quantum processor to support different classes of two-qubit quantum logic gates when compiling quantum circuits, a property that is typically not present in existing platforms. In particular, access to a gate set that includes support for the CZ-, iSWAP-, and SWAP-type families of gates renders conversions between these gate families unnecessary during compilation, as any two-qubit Clifford gate can be executed using at most one two-qubit gate from this set, plus additional single-qubit gates. We experimentally demonstrate that a SWAP gate can be decomposed into one iSWAP gate followed by one CZ gate, affirming a more efficient compilation strategy over the conventional approach that relies on three iSWAP or three CZ gates to replace a SWAP gate. Our implementation makes use of a superconducting quantum processor design based on fixed-frequency transmon qubits coupled together by a parametrically modulated tunable transmon coupler, extending this platform's native gate set so that any two-qubit Clifford unitary matrix can be realized using no more than two two-qubit gates and single-qubit gates.

1. Introduction

The available gate set is an important consideration for practical quantum computation on the hardware platform of choice. A complete single-qubit gate set, supplemented with an entangling two-qubit operation, is a sufficient resource for universal quantum computation [1, 2].

In practice, however, different types of two-qubit gates can be advantageous depending on the particularities of the hardware and its circuit architecture [3], such as the feasible qubit connectivity [3–6], the presence of residual interactions in the Hamiltonian leading to coherent errors, and the dominant noise channels leading to incoherent errors [7–9]. For superconducting qubits, the circuit-QED architecture offers flexibility in addressing such design trade-offs for quantum processors [10].

When executing quantum algorithms, it is often desirable to have easy access to a larger gate set than the bare minimum to enable efficient quantum circuit compilation [3] even though combinations of two-qubit and single-qubit gates can, in principle, perform conversions between two-qubit gate families [5, 11, 12]. The Clifford group contains four such families of locally inequivalent two-qubit gates [12]: CNOT-like gates such as the CZ gate, iSWAP-like gates, SWAP-like gates, and identity [11, 13–15]. Gates are said to be locally inequivalent when one cannot be converted into another using only single-qubit rotations, which typically feature higher fidelity than two-qubit operations and execute $\sim 10\times$ faster. An example of a conversion between a CNOT gate and an iSWAP gate is shown in figure 1.

In devices made for noisy intermediate-scale quantum (NISQ) [16] computation, quantum circuits run within their constituent qubits' coherence time and without error correction. Here for example, the CZ gate

is advantageous in some NISQ algorithms based on minimizing the energy of the Ising Hamiltonian [17], which includes an oftentimes residual and always-on ZZ interaction that commutes with the CZ unitary matrix.

In another example, the iSWAP family is naturally suited for quantum chemistry calculations where the qubits represent fermionic states; here, iSWAP plus single-qubit gates form the fSWAP interaction, which respects fermionic anti-commutation relations. This interaction has known decompositions using four CNOT-like gates [18], although two CNOT gates is theoretically sufficient [19] as $f\text{SWAP} \in \text{SO}(4)$.

However, iSWAP gates do not commute with ZZ terms [20], which is why errors of such always-on interactions may accumulate. Even though the CZ gate is implemented via a Hamiltonian that also does not commute with ZZ terms, this gate may have its pulse parameters adjusted to omit static ZZ terms [20, section V.D]. Other methods of mitigating residual ZZ coupling remains an active study for parametric CZ gates, such as carefully engineering gate pulse parameters to target ZZ-free coordinates in parameter space [21].

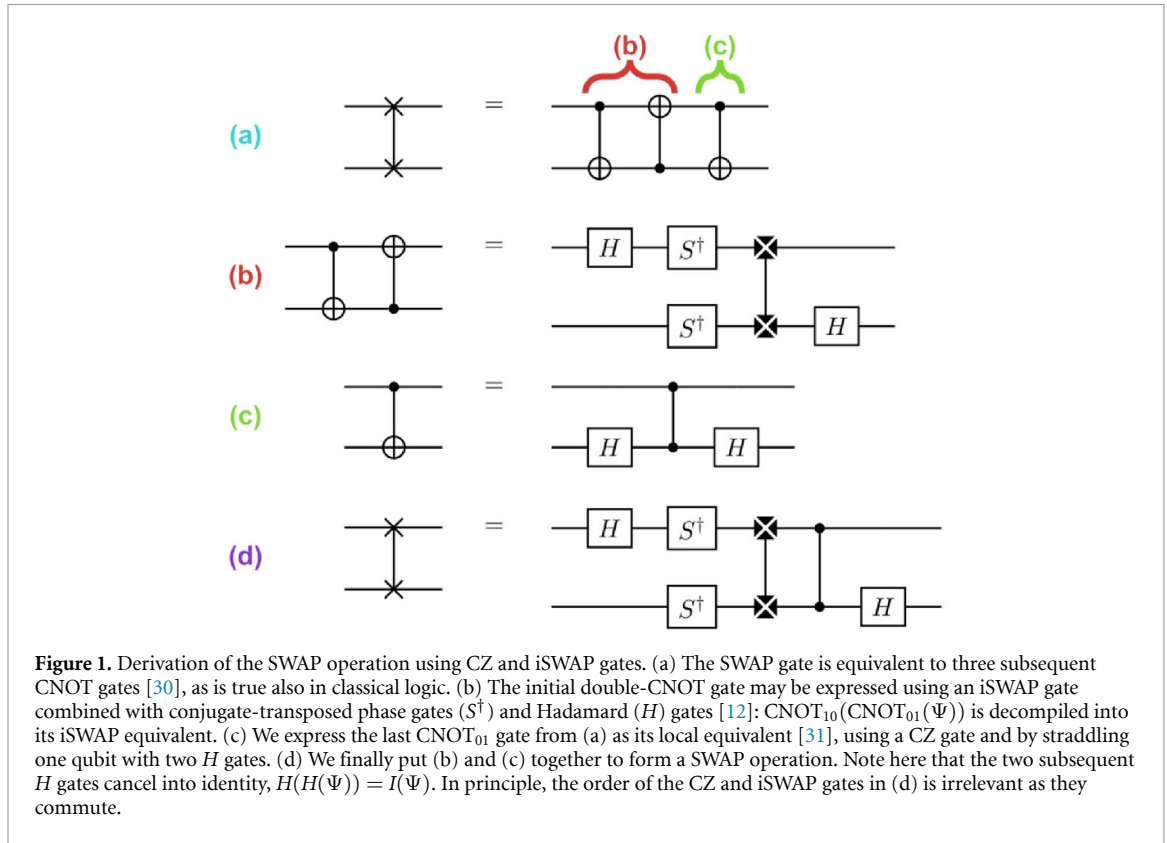
In the context of fault-tolerant quantum computation, proposals for subcells of the error-correcting surface code use the CNOT gate, locally equivalent to the CZ gate, as their fundamental two-qubit gate [22]. On the other hand, in surface-code circuits, the iSWAP gate may reduce the coupling per qubit and boosts resilience to leakage compared to CNOT-like gates [5].

Superconducting qubits afford versatility in the implementation of two-qubit gates as we can engineer interactions between states, and have access to a rich level structure. One prominent implementation is the parametrically modulated coupler, which, with transmon qubits, has been used to demonstrate high-fidelity two-qubit gates [20, 21, 23–27] and to run quantum algorithms [17]. The device used in the investigation reported here comprises a frequency-tunable transmon coupler that is connected to two fixed-frequency transmon qubits, see figure 2. Microwave modulation of the coupler's frequency mediates interactions between the qubits. In this system, we can address the different transitions in the two-qubit manifold, which enables native execution of CZ and iSWAP gates, similar to previous examples of simultaneous support for parametric CZ and iSWAP gates [20, 21, 23].

By means of microwave pulses, the CZ gate can be implemented by driving the $|11\rangle \rightarrow |20\rangle$ (or $|11\rangle \rightarrow |02\rangle$) transition; the iSWAP gate relies on $XX+YY$ exchange interaction [28], activated by driving the $|01\rangle \rightarrow |10\rangle$ transition; and the bSWAP uses the $XX-YY$ exchange interaction, which can be realized by driving the two-photon $|00\rangle \rightarrow |11\rangle$ transition [24, 28, 29]. We call these gates native to this system, whereas for example CNOT is not native, since we compose it from CZ and single-qubit gates, shown in figure 1(c).

One important gate that is not native to most [3, 32] quantum processors is the SWAP gate. This gate is useful in systems without long-range or all-to-all couplings between qubits, which is typical for superconducting quantum processor architectures. The SWAP gate allows a two-qubit operation between two distant qubits to be performed by repeatedly swapping the states [4, 6] within a chain of qubits, until the desired qubit states come adjacent to one another [6]. Moreover, the SWAP-like family of gates is the only two-qubit gate type that is guaranteed to avoid creating entanglement for any separable input state [33]. Notably, the SWAP gate is native to quantum-dot-based quantum processors through isotropic Heisenberg exchange coupling [34, 35]. Native support for the SWAP-like family of gates is rare in other platforms, but has been shown in recent advancements alongside support for iSWAP-like gates and the B gate, in a different superconducting device platform [36]. Whenever the SWAP gate is not natively supported, it can be realized by either three CZ gates [37, 38], three iSWAP gates [39], or three CNOT gates in sequence [30, figure 1].

In this paper, we show experimentally that we can implement SWAP operations more simply by combining one CZ gate and one iSWAP gate in sequence, with the addition of single-qubit gates. We note that CZ and iSWAP gates commute, so that they can be applied in any order—and even simultaneously, which would further reduce the total gate time, compared to the longer concatenation required by the three consecutive CZ, iSWAP, or CNOT gates. Our work represents an improvement at the circuit compilation level; we augment the native gate set without changing the architecture. Our implementation reduces the length of compiled quantum algorithms produced by contemporary circuit synthesis tools where SWAP operations are realized as trivial solutions of three CNOT-like gates, or three iSWAP-like gates. Such enhancements are particularly crucial considering NISQ-era gate-error constraints given the high number of consecutive gates that must be implemented during the course of an algorithm [3, 4, 6, 16]. Besides NISQ, an augmented gate set may become useful in quantum error correction, where encoding and decoding segments can benefit from circuits designed to minimize the impact of SWAP gates [40].



2. SWAP through combined CZ and iSWAP gates

This section derives the proposed SWAP operation, composed of an iSWAP gate followed by a CZ gate. The role of the iSWAP gate is to exchange state information, albeit with unwanted added phases for some input states, which are conditionally corrected by the CZ gate.

Our quantum circuits assume right-handed rotations about Bloch sphere axis $N \in [X, Y, Z]$, where $|0\rangle$ is located at the north pole. We denote *increasing* a phase as a counter-clockwise Z -axis rotation of the Bloch vector when viewed from $|0\rangle$. The basis elements of our two-qubit computational states are given as $[|00\rangle, |01\rangle, |10\rangle, |11\rangle]^T$, and accordingly, we define our two-qubit gates as the following unitary operations:

$$U_{\text{CZ}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (1)$$

$$U_{\text{iSWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2)$$

$$U_{\text{SWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3)$$

We perform U_{SWAP} using our equivalent quantum circuit derived in figure 1. The quantum circuit in figure 1(d) uses conjugate-transpose phase gates (S^\dagger) and Hadamard (H) gates, defined as

$$S^\dagger = \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}, \quad (4)$$

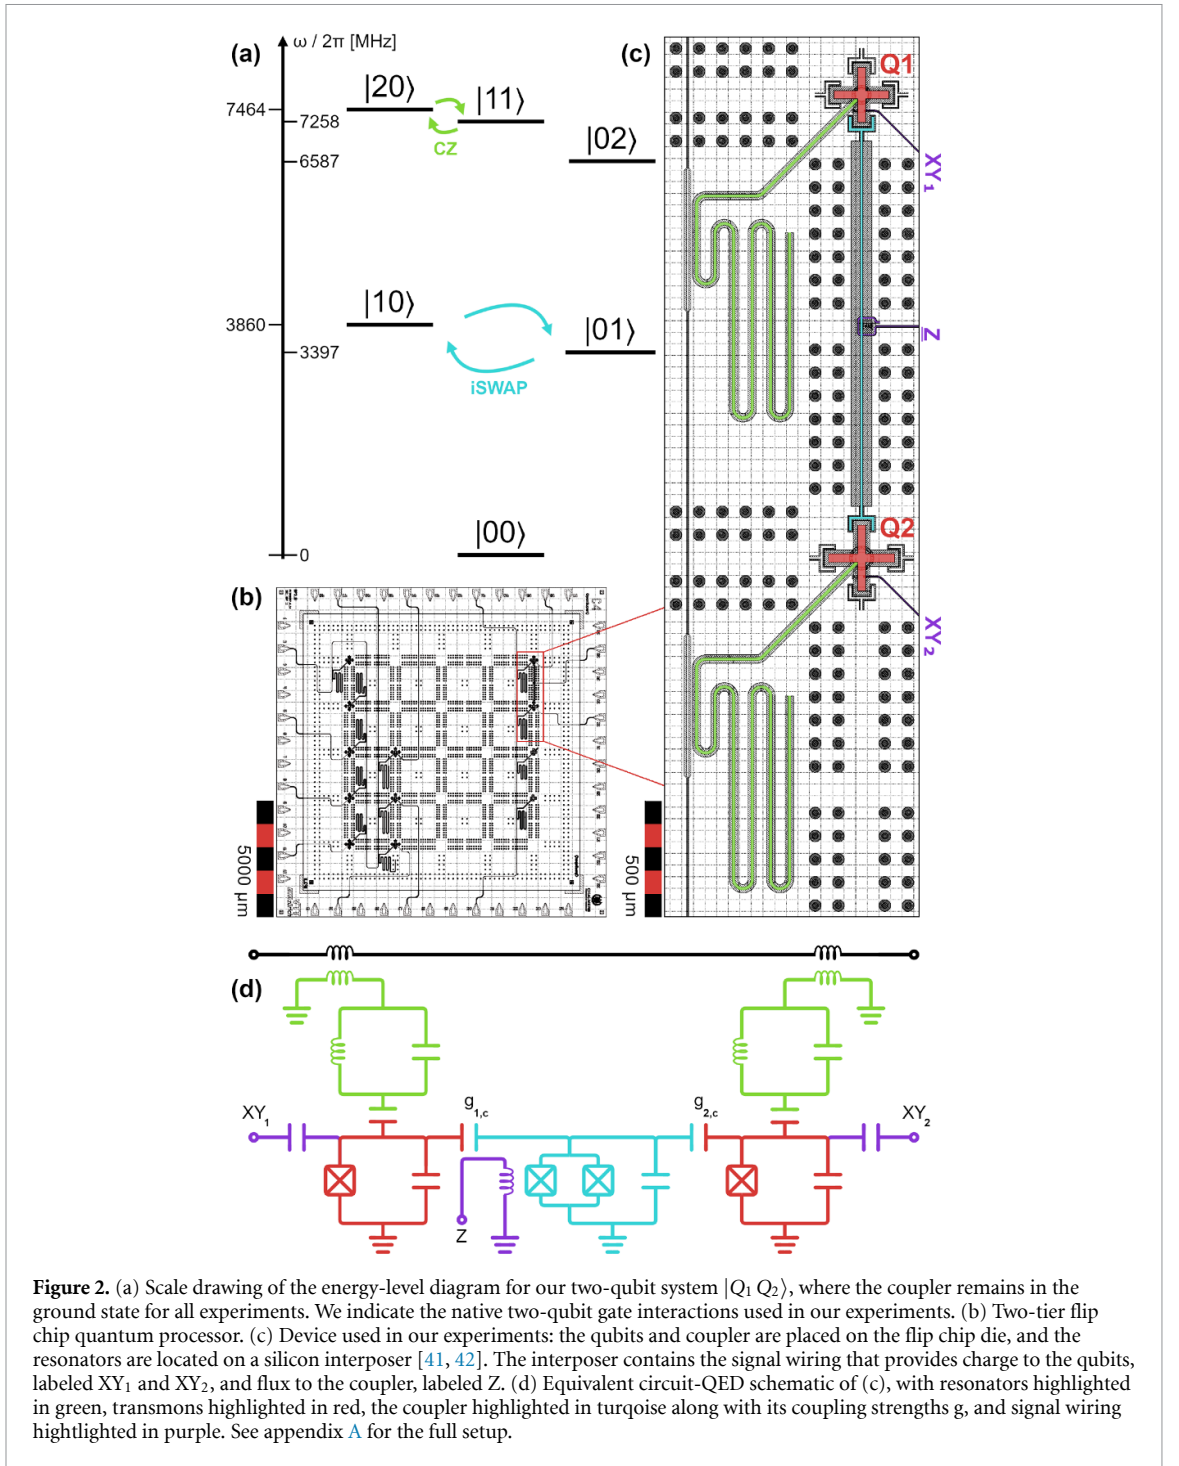


Figure 2. (a) Scale drawing of the energy-level diagram for our two-qubit system $|Q_1 Q_2\rangle$, where the coupler remains in the ground state for all experiments. We indicate the native two-qubit gate interactions used in our experiments. (b) Two-tier flip chip quantum processor. (c) Device used in our experiments: the qubits and coupler are placed on the flip chip die, and the resonators are located on a silicon interposer [41, 42]. The interposer contains the signal wiring that provides charge to the qubits, labeled XY₁ and XY₂, and flux to the coupler, labeled Z. (d) Equivalent circuit-QED schematic of (c), with resonators highlighted in green, transmons highlighted in red, the coupler highlighted in turquoise along with its coupling strengths g , and signal wiring highlighted in purple. See appendix A for the full setup.

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad (5)$$

where the S^\dagger gates are implemented as virtual-Z gates, explained in section 4.

3. Device

For the experimental demonstration, we used two fixed-frequency transmon qubits [43] of a flip chip quantum processor shown in figure 2, based on Chalmers' 3D architecture [41, 42]. We drive single-qubit state transitions using individual 'XY' control lines and two-qubit transitions via a parametrically modulated, flux-tunable transmon coupler [25] using a 'Z' control line.

The transmons are dispersively coupled to coplanar waveguide resonators for state readout, visible in figures 2(c) and (d). The qubits and couplers are located on one chip, while the resonators and control lines are located on a control-chip interposer [41]. The resonators share a common feedline, enabling multiplexed readout of the qubits. The coupler has no dedicated readout resonator.

As our qubits are fixed in frequency, we supply microwave pulses to drive transitions between the energy levels as illustrated in figure 2(a). The $|0\rangle \rightarrow |1\rangle$ transition frequencies of the transmons are $f_1 = 3.860$ GHz and $f_2 = 3.397$ GHz. In these experiments, we bias the coupler with a static magnetic flux at $\Phi = -0.336 \Phi_0$ [25], putting its eigenfrequency at $f_c = 4.863$ GHz, i.e. more than 1 GHz above both qubits. A complete setup description is available in appendix A, the system's Hamiltonian in appendix B, and additional system parameters in table C1 of appendix C.

Our readout state assignment is configured to distinguish between states $|0\rangle$, $|1\rangle$, and $|2\rangle$. The readout frequency is chosen such that the population spacing in the complex readout plane is maximized. The readout amplitude was then chosen to maximize state-assignment fidelity. The readout duration was not optimized; we selected a readout duration of $2.3 \mu\text{s}$ to ensure that we would integrate over the readout trace differences corresponding to the different qubit states [44]. Appendix D shows additional details of the readout setup, as well as how we used two-qubit confusion matrices to post-process our experimental data [45, 46].

Our two-qubit gates rely on parametric exchange coupling [25, 47, 48], which enables the coupling of widely detuned, fixed-frequency qubits [25], although with a known drawback of being slow compared to alternatives such as cross-resonance gates [49]. We consider CZ and iSWAP gates as shown in figure 2(a): to execute iSWAP gates, we apply a resonant pulse at the frequency $\omega_{10} - \omega_{01}$, causing a transition between states $|01\rangle$ and $|10\rangle$. To execute CZ gates, we drive a resonant pulse at the frequency $\omega_{20} - \omega_{11}$, causing a brief excursion of the qubit state outside of the computational subspace (and back again) for a gate duration τ_{CZ} ; the completed round-trip adds a geometric phase of π . We perform SWAP operations by compiling Hadamard, S^\dagger , iSWAP, and CZ gates, as was derived in section 2.

4. Two-qubit gate characterization

We characterize the two-qubit gate operations in a simplified way in a Ramsey interference experiment, see figure 3. Conventional Ramsey interferometry measures the accrued phase over time t of a single-qubit superposition state, $|0\rangle + e^{i\varphi(t)}|1\rangle$. Instead, we use a fixed, short wait time τ between the \sqrt{X} initialization and analysis pulses, leading to a phase φ_τ . We emulate phase accrual by adding a virtual-Z gate V_z which offsets the phase of each subsequent pulse by φ .

In the quantum circuits at the top of figure 3, the upper rail represents the control qubit q_c , and the lower rail the target qubit q_t . Here, we use notation $|q_c q_t\rangle$. We verify that the CZ gate results in the correct qubit populations and phases by using the conditional Ramsey sequence shown in figures 3(a) and (d); likewise, we characterize the iSWAP and SWAP operations by using the cross-Ramsey sequence, in figures 3(b), (c), (e) and (f). In all cases, we first prepare q_t in state $|\bar{i}\rangle$ by using a \sqrt{X} gate, i.e. the Bloch vector is oriented towards $-Y$. We then execute the two-qubit gate of interest, apply the virtual-Z, rotate back by again using a \sqrt{X} gate, and read out in the Z-basis. We verify that the phase of q_t has accrued by the correct amount, depending on the input state $|0\rangle$ or $|1\rangle$ of the control qubit q_c .

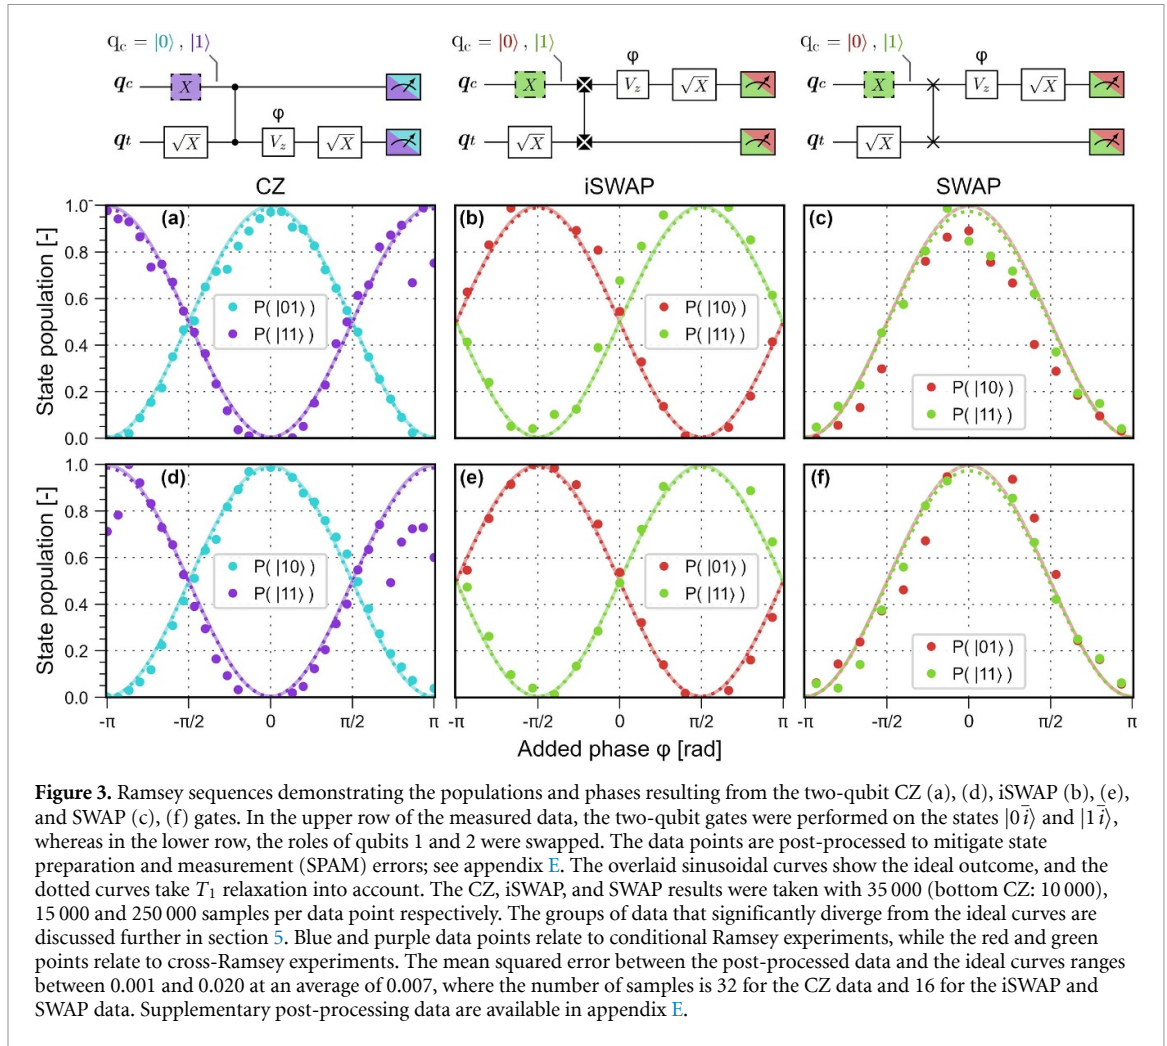
For all experiments, we sweep the phase φ using virtual-Z gates, which determines the Bloch vector longitude to a unique phase coordinate. Consider the CZ gate as an example: the sweep verifies that the state q_t after the CZ gate is $\cos \frac{\pi}{2} |0\rangle + e^{i\varphi} \sin \frac{\pi}{2} |1\rangle$ and not $\cos \frac{\pi}{2} |0\rangle + e^{i(\pi-\varphi)} \sin \frac{\pi}{2} |1\rangle$.

4.1. CZ by conditional Ramsey sequence

As seen in equation (1), the conditional-Z gate adds π phase onto q_t , conditional on q_c being in $|1\rangle$. In figure 3(a) and (d), we see that the phase of q_t accrues by π if we execute a CZ gate when q_c is in $|1\rangle$. Similarly, we see that no phase accrual occurs if we execute a CZ gate while q_c is in $|0\rangle$. Both of these conditions are necessary for verifying a CZ gate.

4.2. iSWAP by cross-Ramsey sequence

At phase $\varphi = 0$ in figures 3(b) and (d), the expected outcome after the iSWAP gate in the cross-Ramsey sequence is the two-qubit state $|+0\rangle$ when q_c is initially prepared in $|0\rangle$, and $|-1\rangle$ when q_c is prepared in $|1\rangle$. If the virtual-Z gate adds 0 phase to q_c , then the final \sqrt{X} rotation does nothing to the superposition state, resulting in 0.5 population at state readout. As phase $\varphi = \pi/2$ is added, the Ramsey oscillation approaches $|0\rangle$ for q_c initially prepared in $|0\rangle$ but instead assumes $|1\rangle$ for q_c initially prepared in $|1\rangle$.



4.3. SWAP by cross-Ramsey sequence

Just as with the iSWAP case, the SWAP operation in figures 3(c) and (f) is verified using a cross-Ramsey sequence. The SWAP operation swaps the two states of q_c and q_t with each other, thus q_c is now in state $|\bar{i}\rangle$. The virtual-Z gate adds a phase φ onto the state of q_c . If $\varphi = 0$, the final \sqrt{X} gate will rotate q_c to $|1\rangle$, and will otherwise produce π -periodic Ramsey oscillations between $|0\rangle$ and $|1\rangle$ depending on the phase added. As SWAP gates are non-entangling, any separable input state into the cross-Ramsey sequence involving a SWAP gate will remain separable. One may thus always discuss the single-qubit states throughout the SWAP cross-Ramsey sequence, as was done here.

5. Discussion and outlook

This discussion mainly targets aspects of our post-processing error mitigation, addresses an observed random switching problem with qubit 1, and discusses the various error sources identified including their relative contribution.

Figure 3 features several clusters of data points that deviate from the ideal outcome. These were acquired close in time. We show in appendix C that Q_1 was afflicted by periods where the acquired Ramsey fringes yielded large fit errors such that T_2^* could not be accurately determined. The erratic behavior points to the presence of a two-level-system defect strongly coupled to the qubit, with adverse effects on spectral stability and coherence. From figure 3 we may estimate such sessions to be the largest error impact when they happen, on the order of 0.25 loss in expected probability on readout. Due to the fixed-frequency design of this device's transmons, we do not suspect flux noise significantly contributing to these observed behaviors. However, observing the T_2^* improvement over T_2^* in figure C2 indicates [50] the presence of low-frequency noise.

Our two-qubit gate implementation has relatively long RF pulses, and as such we are interested in a proof-of-principle demonstration of a simplified SWAP decomposition using CZ and iSWAP gates, rather than a complete gate fidelity characterization (which would be performed, e.g. by means of randomized benchmarking). Here we estimate the theoretically achievable gate fidelities. The mean qubit T_1 energy relaxation times are $77 \mu\text{s}$ and $79 \mu\text{s}$ for Q_1 and Q_2 , respectively, which when taking the T_2^* times of $37 \mu\text{s}$ and $33 \mu\text{s}$ into account, yield a theoretical [8] upper limit of 99.98% single-qubit fidelity for 20 ns pulses but only 98.8% two-qubit fidelity for the 890 ns long CZ gate, and 99.2% for the 640 ns long iSWAP gate. The total duration of our SWAP cross-Ramsey circuit implementation is $1.960 \mu\text{s}$, setting an upper fidelity limit of 97.4%. The additional duration beyond the CZ and iSWAP gates consists of single-qubit gates and pauses inserted by the compiler. This duration should still be compared to the longer $2.67 \mu\text{s}$ minimum duration required for compiling a SWAP operation using three CZ gates; despite the comparably long gates, the default SWAP compilation would have been longer.

However, we have previously demonstrated parametric two-qubit gates of 270–350 ns [17, 41, 51] on comparable quantum processors, and the pulse duration could realistically be further decreased down toward 150 ns, which would enable two-qubit gate errors smaller than 0.1%.

In figure D3 we estimate an energy relaxation error of about 3% during the readout pulse, and a thermal population error of about 2%. Furthermore, from the remainder of the classification error in figure D3, we expect between 2–10% error in $|0\rangle$ to stem from gate error, overlap error, and readout-induced mixing. Comparing figures 3(c)–(f) to figure D6, we see an example between applying, and not applying, the SPAM error mitigation described in appendix E; this difference in population gain is of the order of more than 10%, in which we draw the conclusion that the broad category of SPAM errors encompasses the largest source of error in our experiments. Figure E1 further highlight effects pertaining to the two-qubit readout, such as measurement-induced state decay and readout crosstalk.

Our SWAP decomposition can be improved by an approach beyond the scope for this paper: the parametric interactions that drive the CZ and iSWAP gates, as shown in figure 2(a), are in principle executable in parallel. More generally, the CZ gate (equation (1)) commutes with the iSWAP gate (equation (2)), and they can be taken in any order in figure 1(d), or even simultaneously. Tuning up this bichromatic excitation requires frequency adjustments to account for additional dispersive shifts of the coupler and phase compensation offsets [51]. In practice, realizing such a bichromatic SWAP gate on this specific platform would be possible following upgrades to the instrument control software. We note that a similar bichromatic gate approach has been suggested and simulated for combining bSWAP and iSWAP gates [28].

The SWAP sequence in figure 1(d) may also be simplified by omitting the initial and final Hadamard gates. Considering $U_{\text{iSWAP(CZ)}} = \text{CZ}(\text{iSWAP})$,

$$U_{\text{iSWAP(CZ)}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (6)$$

we may use the S^\dagger gates from equation (4) on both qubits to convert the off-diagonal i elements into 1 while also removing the negation on $|11\rangle$,

$$\begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (7)$$

$$\begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}^{\otimes 2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (8)$$

in which we see that equation (8) is identical to the SWAP unitary matrix from equation (3). Since the conjugate transpose phase gates are implemented as zero-duration virtual-Z gates, such a SWAP sequence may be as short as the duration of one iSWAP and one CZ gate.

Finally, our current approach to verifying the function of our CZ, iSWAP, and SWAP operations using Ramsey interferometry can be further improved by applying a tomographically complete input set of states

[52]. Such a verification uses our current approach as a necessary condition for certifying the CZ, iSWAP, and SWAP operations, while also determining these gate operations to a unique solution of the unitary gate matrices in section 2.

6. Conclusion

We proposed and experimentally confirmed a composite two-qubit SWAP gate compiled from one iSWAP gate followed by one CZ gate. We demonstrated this on a subset of our superconducting flip chip-integrated quantum processor using quantum-gate interactions driven by parametric modulation of a coupler.

This extends the available gate set in this architecture to include the iSWAP gate, reducing the need for three CZ gates to just one iSWAP and one CZ gate for each decompiled SWAP gate. We characterized the compound SWAP gate using a cross-Ramsey sequence with two-qubit input states, followed by post-processing to account for SPAM errors.

Going forward, we expect that a sequence of optimized CZ and iSWAP gates with shorter duration can bring the two-qubit gate error down below 0.1%, even with coherence times at the 100 μ s level. We further expect a gate-time reduction by approximately one half by driving both a CZ gate and an iSWAP gate simultaneously.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://dx.doi.org/10.5281/zenodo.11149402> [53].

Acknowledgments

For their discussions and assistance, we thank Per Delsing, Miroslav Dobšíček, Anita Fadavi Roudsari, Göran Johansson, Tong Liu, Simon Pettersson Fors, Hampus Renberg Nilsson, Daryoush Shiri, and Vitaly Shumeiko at Chalmers, as well as Riccardo Borgani and Mats O. Tholén at KTH Stockholm.

We acknowledge financial support from the Knut and Alice Wallenberg Foundation (KAW) through the Wallenberg Center for Quantum Technology (WACQT), and in part from the EU Flagship on Quantum Technology under Project HORIZON-CL4-2022-QUANTUM-01-SGA Grant 101113946 OpenSuperQPlus100. A Frisk Kockum acknowledges support from the Swedish Research Council Grant 2019-03696, and the Swedish Foundation for Strategic Research Grants FFL21-0279 and FUS21-0063.

The devices were fabricated at the Myfab Chalmers Nanofabrication Laboratory. Evaporation of bumps and flip-chip bonding was done at the VTT Technical Research center of Finland.

Appendix A. Experimental setup

Figure A1 illustrates our cryogenic microwave setup used to perform the experiments in this paper.

Appendix B. Energy dynamics

Our device shown in figure 2 constitutes a three-body system of multi-level anharmonic oscillators [43]. The two lowest energy levels of the oscillators define our qubits' computational basis states [30]. The anharmonic oscillator in the middle of figure 2(d) is used to couple together the two qubits on the left and right. We describe our circuit's behavior using the following Hamiltonian:

$$\begin{aligned} \frac{\hat{H}}{\hbar} = & \sum_{i=1}^2 \omega_i \hat{a}_i^\dagger \hat{a}_i + \frac{\eta_i}{2} \hat{a}_i^\dagger \hat{a}_i (\hat{a}_i^\dagger \hat{a}_i - 1) \\ & + \omega_c(\Phi) \hat{b}^\dagger \hat{b} + \frac{\eta_c}{2} \hat{b}^\dagger \hat{b} (\hat{b}^\dagger \hat{b} - 1) \\ & + \sum_{i=1}^2 g_{ic} (\hat{a}_i^\dagger + \hat{a}_i) (\hat{b}^\dagger + \hat{b}) \end{aligned} \quad (\text{B.1})$$

where \hat{a} , \hat{b} , \hat{a}^\dagger , \hat{b}^\dagger are annihilation and creation operators [55], with \hat{a}_i acting on transmon i and \hat{b} acting on the coupler; g_{ic} are the coupling strengths between qubits i and the coupler c ; ω_i are the angular frequencies of all three transmons, where the coupler transmon's frequency is adjustable by a threading magnetostatic flux Φ ; η_i and η_c are the transmons' anharmonicities [43], with suffix c denoting the coupler in particular; and $\hbar = \frac{h}{2\pi}$ is the reduced Planck constant. We assume that the direct qubit–qubit coupling $g_{12} \ll (g_{1,c}, g_{2,c})$. We follow notation $\omega_i = 2\pi f_i$, where ω denotes an angular frequency and f the ordinary frequency.

Appendix C. Parameter estimation

Figures C1 and C2 show acquired decoherence data used for establishing the mean T_1 , T_2^* , and T_2^c metrics. A hard upper limit of 15% fit error was enforced for all fits; data where a fit could not be determined with an accuracy below 15% of the fitted decoherence time T are marked as unreliable and discarded. The fit error is determined from the square root of the diagonal elements of the covariance matrix of the fitted parameters, where the fit equations are

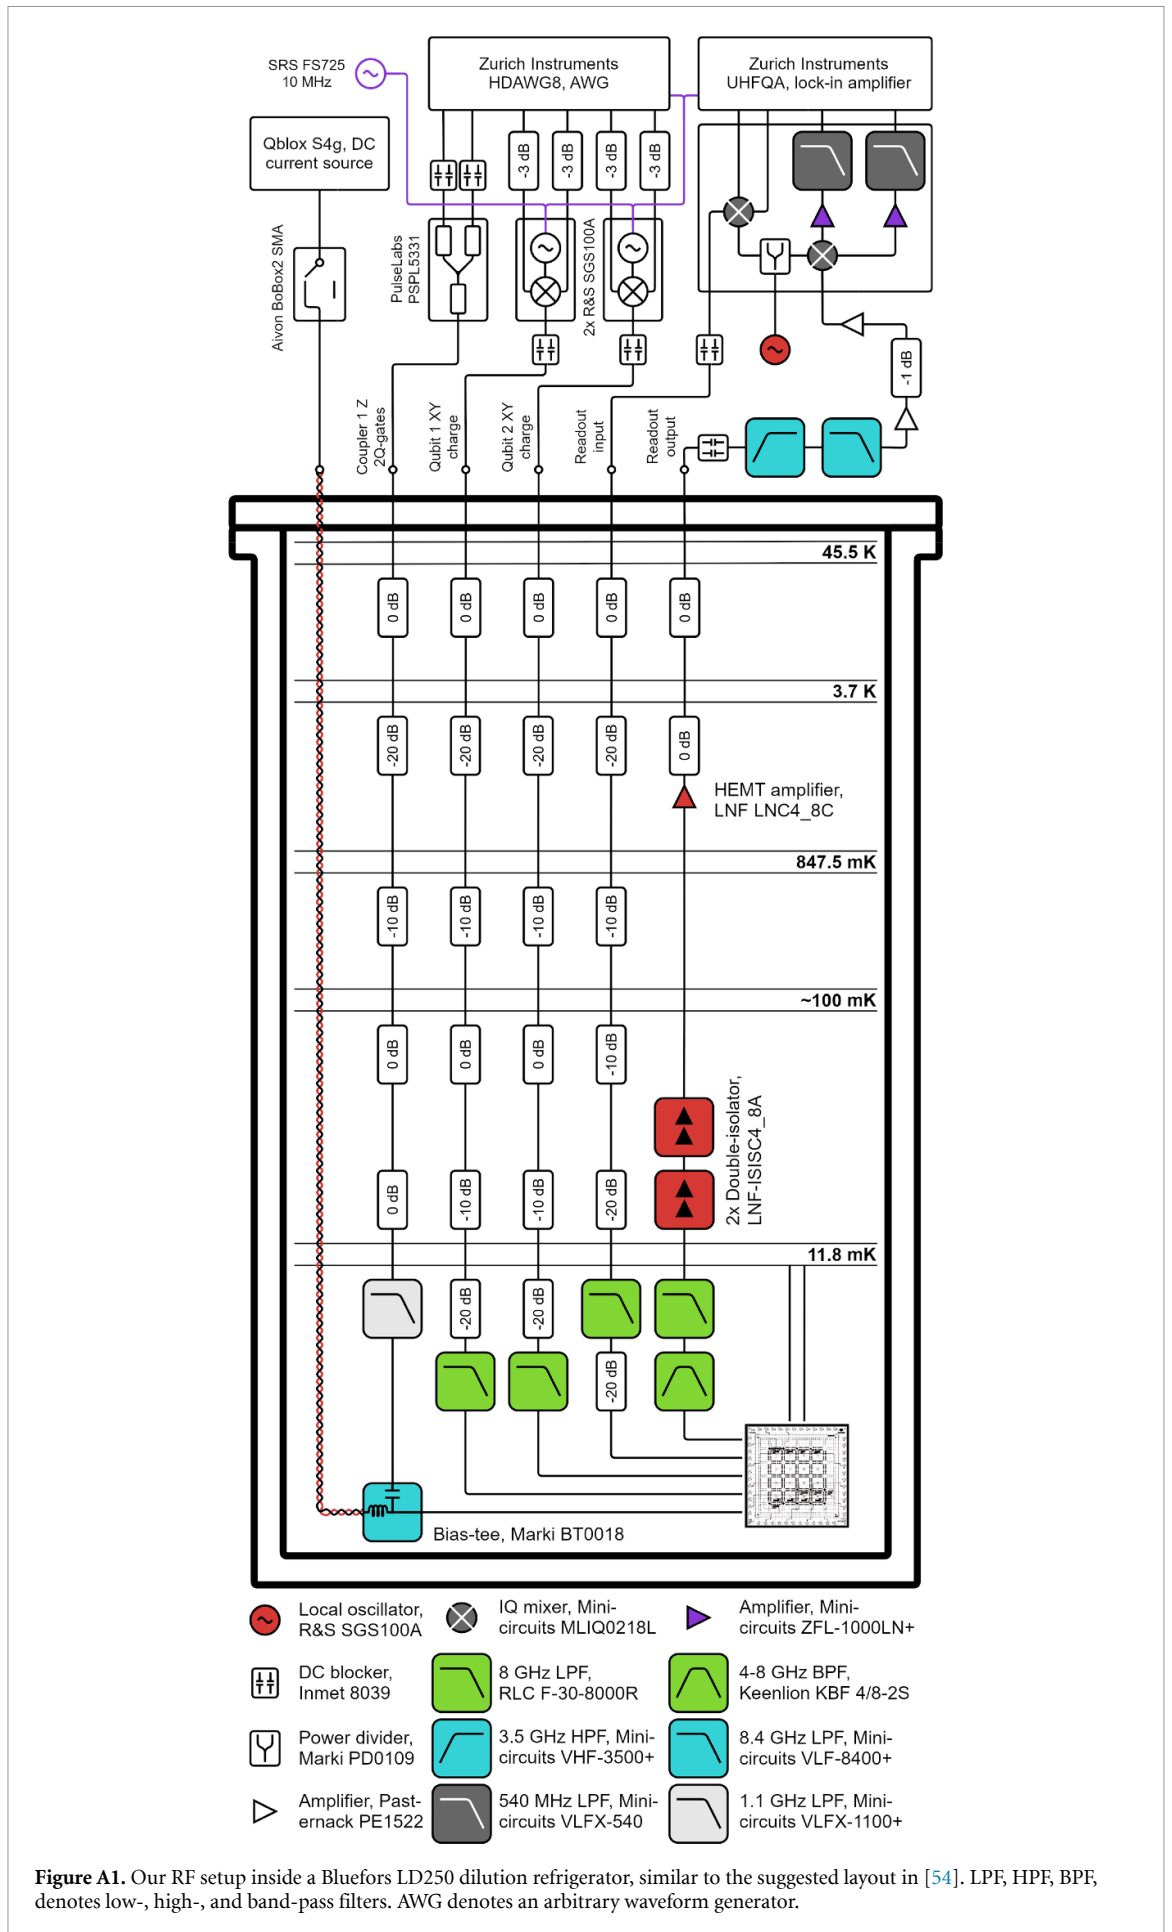
$$Y_{T_1} = A e^{-t/T_1} + m, \quad (\text{C.1})$$

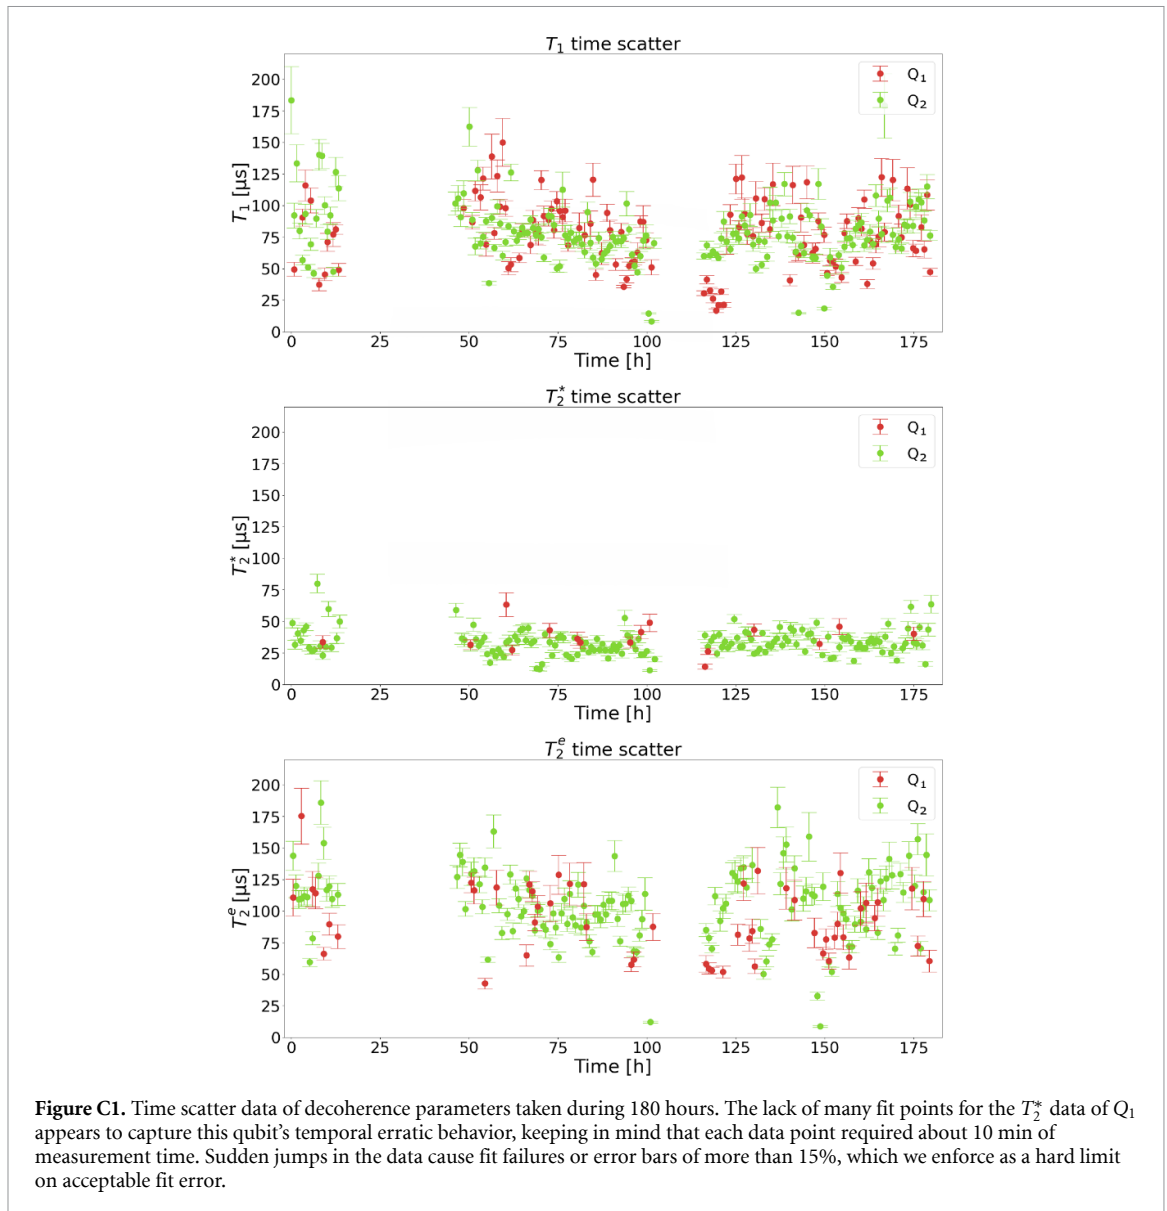
$$Y_{T_2^*} = A e^{-t/T_2^*} \cos(2\pi f t + \varphi) + m, \quad (\text{C.2})$$

$$Y_{T_2^c} = A e^{-t/T_2^c} + m, \quad (\text{C.3})$$

where A is an amplitude scalar; t is the time axis stepped through in the decoherence time measurement; T_1 , T_2^* , T_2^c are decoherence times metrics; f is the oscillation rate of the Ramsey fringe; φ its apparent phase offset; and m its demodulated magnitude offset.

Table C1 lists additional parameter data for our device.





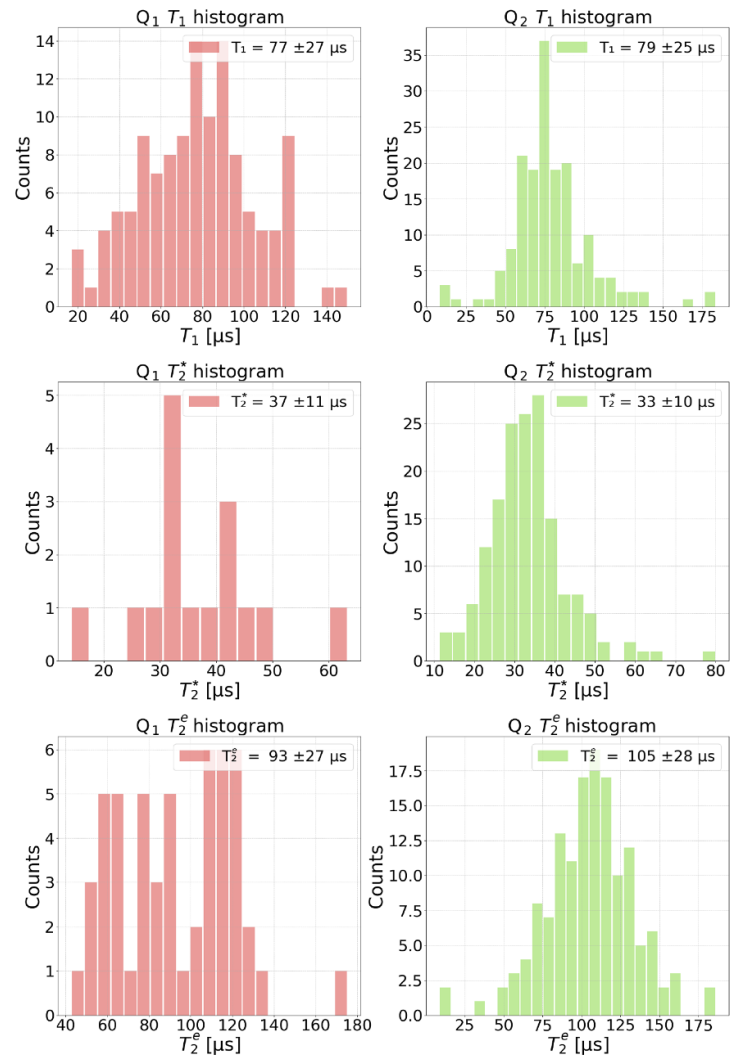


Figure C2. Histogram distribution of decoherence parameters taken during 180 hours. The number of bins K in the histogram were selected according to the work by Doane [59], $K = 1 + \log_2(N) + |\gamma|/\sigma_\gamma$, where N is the number of samples, γ is the third standardized moment, and σ_γ is the standard deviation of γ . This approach attempts to highlight skew in a normally distributed dataset while also avoiding oversmoothing distributions.

Table C1. Physical parameters of our system. The DRAG [56] parameter α is defined using notation from [57, (3.12)]. The error bars reported for the frequency offsets Δ were calculated using simplified error propagation [58],

$$\text{Err}(\Delta f_{\text{CZ}}) = \sqrt{(-1)^2 \cdot \text{Err}(f_{12, Q1})^2 + (+1)^2 \cdot \text{Err}(f_{01, Q2})^2 + (-1)^2 \cdot \text{Err}(f_{\text{CZ}})^2},$$

$$\text{Err}(\Delta f_{\text{iSWAP}}) = \sqrt{(+1)^2 \cdot \text{Err}(f_{01, Q1})^2 + (+1)^2 \cdot \text{Err}(f_{01, Q2})^2 + (-1)^2 \cdot \text{Err}(f_{\text{iSWAP}})^2}. \text{ The error reported for the three-state mean assignment fidelity, } F_{\text{Assign}\{|0\rangle, |1\rangle, |2\rangle\}}, \text{ is the standard error from the mean (S.E.M. = } \sigma/\sqrt{3}\text{), and the variance is } \sigma^2 = (\sum_{i=1}^3 (F_{|i\rangle} - F_{\text{Assign}\{|0\rangle, |1\rangle, |2\rangle\}})^2) / (3 - 1).$$

Parameter	Q_1, R_1	Q_2, R_2	C_1
$f_{\text{RO}, 0\rangle}$	6.48084 GHz	6.25968 GHz	n/a
$f_{\text{RO}, 1\rangle}$	6.48068 GHz	6.25946 GHz	n/a
$f_{\text{RO}, 2\rangle}$	6.48056 GHz	6.25920 GHz	n/a
$f_{\text{RO}, \text{optimal}}$	6.48048 GHz	6.25960 GHz	n/a
$\chi_{ 0\rangle \rightarrow 1\rangle}$	80 kHz	110 kHz	
$\chi_{ 1\rangle \rightarrow 2\rangle}$	60 kHz	130 kHz	
τ_{RO}		2.3 μs	n/a
$F_{\text{Assign}\{ 0\rangle, 1\rangle, 2\rangle\}}$	0.8793 ± 0.0494	0.8873 ± 0.0528	n/a
$f_{ 0\rangle \rightarrow 1\rangle}$	3.86011 GHz ± 156 kHz	3.39741 GHz ± 156 kHz	4.863 GHz
$f_{ 1\rangle \rightarrow 2\rangle}$	3.60404 GHz ± 190 kHz	3.18992 GHz ± 190 kHz	
$\eta / 2\pi$	-256.07 MHz ± 346 kHz	-207.49 MHz ± 346 kHz	
$\tau_{1\text{QB}}$		20 ns	
α_{Motzoi}	$-770.5 \cdot 10^{-3}$ $\pm 13 \cdot 10^{-3}$	$-221.4 \cdot 10^{-3}$ $\pm 13 \cdot 10^{-3}$	
T_1	77 $\mu\text{s} \pm 27 \mu\text{s}$	79 $\mu\text{s} \pm 25 \mu\text{s}$	
T_2^*	37 $\mu\text{s} \pm 10 \mu\text{s}$	33 $\mu\text{s} \pm 10 \mu\text{s}$	
T_2^e	93 $\mu\text{s} \pm 27 \mu\text{s}$	105 $\mu\text{s} \pm 28 \mu\text{s}$	
$f_{\text{LO}, XY}$		3.600000000 GHz	n/a
$f_{\text{LO}, RO}$		6.600000000 GHz	n/a
Φ_{bias}		n/a	$-0.336 \Phi_0$ $\pm 0.003 \Phi_0$
$f_{ 0\rangle \rightarrow 1\rangle, \Phi=0}$	3.8610 GHz	3.3977 GHz	6.9373 GHz
$g_{q,c} / 2\pi$	37 MHz	35 MHz	
CZ			
τ_{CZ}		890 ns	
f_{CZ}		207.24 MHz ± 10 kHz	
Δf_{CZ}		-0.61 MHz ± 246 kHz	
$\varphi_{\text{comp.}}$	+0.61 rad	+1.03 rad	0 rad
iSWAP			
τ_{iSWAP}		640 ns	
f_{iSWAP}		461.51 MHz ± 82 kHz	
Δf_{iSWAP}		+1.19 MHz ± 235 kHz	
$\varphi_{\text{comp.}}$	+2.24 rad	+3.36 rad	+4.86 rad
SWAP			
τ_{SWAP}		1812 ns	
$\varphi_{\text{comp.}}$	-1.07 rad	+1.25 rad	= φ_{iSWAP}

Appendix D. Experimental tune-up procedure

Our single-qubit gate tune-up mainly follows well-established parameter estimation methods used for superconducting quantum computing [57, 60].

Our microwave drive pulses are generated using a Zurich Instruments HDAWG8 arbitrary waveform generator (AWG) operating in 300 MHz mode, 16-bit 2.4 GSa/s. Our single-qubit gates are realized as single sideband (SSB) pulses upconverted in frequency using standalone phase-locked mixers with individual local oscillators (Rohde & Schwarz SGS100A), see the setup shown in figure A1. Note that the coupler tone is not supplied onto the QPU using frequency mixing; the two-qubit gate pulses are generated at baseband on the AWG.

All of our single-qubit operations use a 20 ns Gaussian envelope, whereas for all two-qubit operations this envelope is a flat rectangle with a Gaussian-shaped 10 ns rise time and 10 ns fall time.

Our single-qubit pulses use the Derivative Removal by Adiabatic Gate (DRAG) technique to lessen single-qubit gate coupling to leakage levels [56]. We use DRAG notation from [57, equation (3.12)], in which the Motzoi parameter is given as $-\alpha/\eta$, where α is a scaling parameter that is tuned-up, and η is the transmon anharmonicity in units of angular frequency.

D.1. Selection of repetition rates and local oscillator frequencies

When using an AWG to generate microwave pulses at baseband, and then using frequency mixers to upconvert the pulse carrier into the final qubit frequencies as done in figure A1, we must ensure that the repetition rate of every iteration of the experiment is commensurate with the frequency of the mixers' local oscillators as we describe in [51]. Otherwise, the LO will erroneously gather phase between quantum circuit iterations. This phase error refers to the start of the quantum circuit that is currently being played, which is seen as the reference from which our quantum circuit phases are referred to.

In a scenario where a repetition rate is selected such that it is not phase commensurate with the LO frequency, the CZ and iSWAP tune-up steps where local phase adjustment is determined [20, equations (C3) and (C11)] will still provide measurement results that appear normal for a successful experiment, and a local phase compensation parameter can be extracted from the data. However, for quantum circuits that depend on the phase of the coupler pulse, as is the case when using iSWAP gates, the readout will imply experiment failure by showing maximally mixed states. This erroneous result happens whenever the coupler phase impacts the final state; the phase of the coupler pulse appears to be random, and this random phase information will get lost in the averaging done between readout shots.

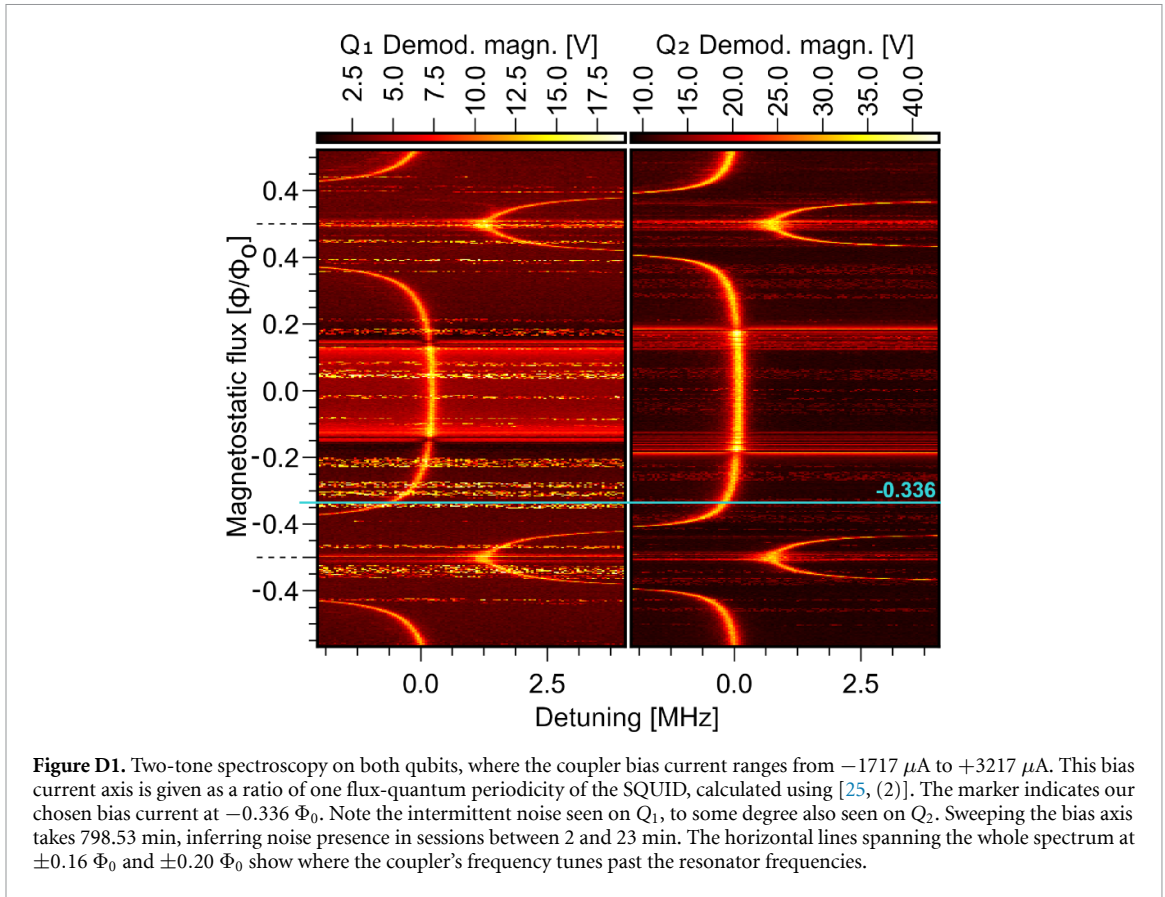
We selected an LO frequency that is a multiple of the inverse of the repetition rate, thus being phase commensurate between measurement iterations [51]. We chose 400 μ s, sharing least common multiple 2500 with $f_{LO} = 3.6$ GHz. At a 400 μ s repetition rate, the qubits would on average be expected to naturally relax to $\exp(-400/78) = 0.6\%$ population between experiment iterations, which is about three times lower than the residual thermal population of the dilution refrigerator.

D.2. Coupler biasing

The coupler is based on a superconducting quantum interference device (SQUID) [61, 62]. A SQUID-based coupler allows us to adjust its transmon frequency by threading its central loop with magnetostatic flux, which we generate from a DC bias. This bias is current-controlled using a Qblox S4g 18-bit current source. Using current-controlled sources compensates for variations in resistance of the feedline, for instance due to temperature or moisture dependence of the room-temperature cabling. Using voltage-controlled sources for SQUID biasing does not stabilize the magnetostatic flux against changes in feedline resistance, thus such errors propagate onto the SQUID biasing.

Since our device does not feature a resonator for reading out the SQUID coupler, we instead perform SQUID parameter estimation via its coupled qubits, specifically two-tone spectroscopy as a function of the applied magnetostatic coupler bias as is shown in figure D1.

We initially select a magnetostatic bias close to $0.5 \Phi_0$ for our tunable coupler, which is nominally identical to $-0.5 \Phi_0$ (see figure D1 and [25, equation (2)]). Next, we proceed with setting up readout state assignment, as will be explained in appendix D.3. We then proceed with two-qubit gate tune-up. Since the readout state assignment has been tuned up, we may fine-tune the coupler flux bias to prevent leakages into undesired states while modulating the coupler with the carrier tone of a two-qubit gate. To find such a flux bias, we first prepare the state $|11\rangle$. Then, we sweep the coupler's $\{\omega_c, \Phi\}$ parameter space while applying a constant carrier tone onto the coupler that equals the $|11\rangle \rightarrow |20\rangle$ transition frequency. We avoid selecting biases where we discern that this modulation induces population transfer into unwanted states. As can be seen in figure D1, the qubits' transition frequencies are expected to shift on the order of a few MHz as a



function of applied flux; it is good practice to be aware of the spectral width of the applied $|11\rangle \rightarrow |20\rangle$ coupler tone, to ensure that the $|11\rangle \rightarrow |20\rangle$ transition is still driven as the qubits move in frequency.

D.3. Readout setup and optimization

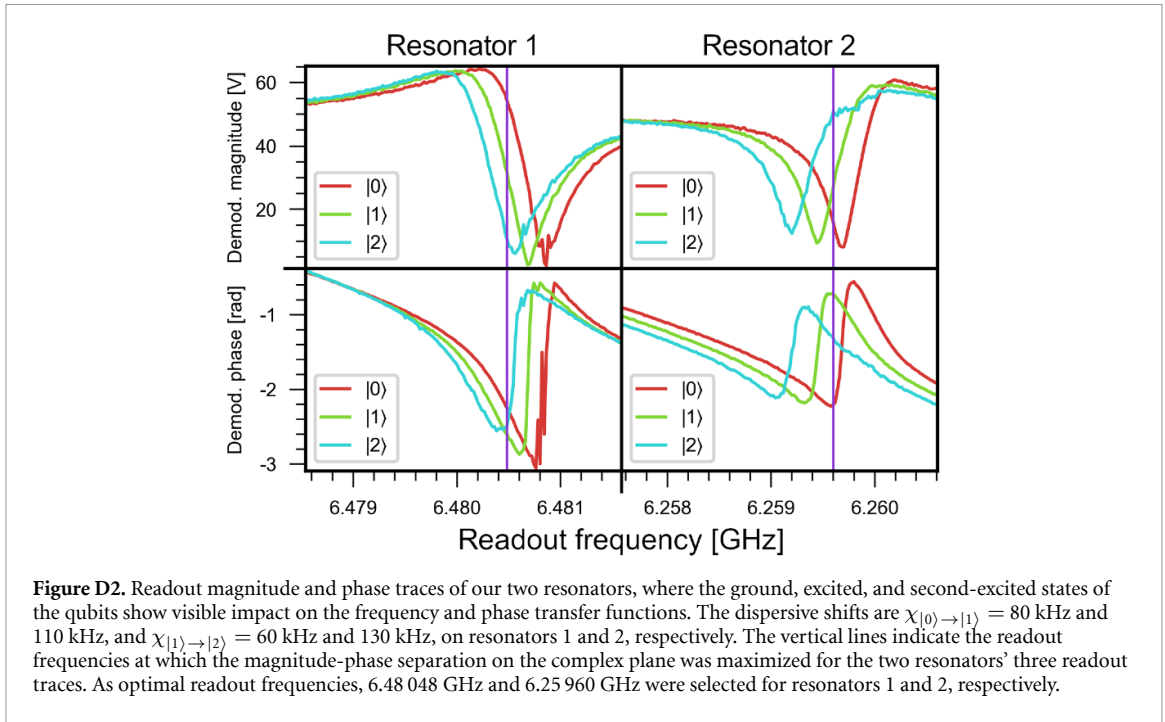
We perform state readout using a Zurich Instruments UHFQA lock-in amplifier, 600 MHz, 12/14-bit ADC/DAC 1.8 GSa/s, facilitated by a combination of microwave devices for frequency conversion as is shown in figure A1. The readout is performed in transmission mode where the S_{21} transmission function is used for readout parameter estimation, as shown in figure D2.

We set up the lock-in amplifier to discard the first 1000 scoped samples, to allow for signal propagation through the setup and resonator ring-up. We collect the following 4096 samples, corresponding to $2.3 \mu\text{s}$, which makes up two time traces corresponding to the in-phase and quadrature (IQ) components of the signal. These two time traces are demodulated to baseband, and integrated, which forms a single complex data point in the IQ plane. After averaging, each data point in figure D2 corresponds to such a data point.

D.3.1. Frequency optimization for state assignment

We proceed with optimizing the frequency of the resonator ring-up tone used for state readout, which we denote as the 'readout frequency'. This frequency optimization is performed by sweeping the readout tone's carrier frequency while monitoring magnitude and phase of the ensuing three frequency responses for $\{|0\rangle, |1\rangle, |2\rangle\}$. To minimize data transfer between the PC and the lock-in amplifier, changing the readout tone sent onto the resonators is done by stepping the LO frequency of the mixer connected to the lock-in amplifier's pilot tone.

Since we have acquired the qubit control pulse parameters for the $|0\rangle \rightarrow |1\rangle$ and $|1\rangle \rightarrow |2\rangle$ transitions of both qubits, we may place the qubits in the three respective states, and trace out the IQ coordinates for every step along the x -axis of figure D2. We analyze the distance between the three respective IQ coordinates for all frequency steps, in order to select an optimal readout frequency. The parameter being optimized for, is the perimeter spanned by the three population centers of the readout states; optimizing for different parameters may better suit certain usage cases, one optimization alternative is to minimize the worst distance between any two states. It is crucial to note that the resonator frequency responses at this time are acquired using averaged measurements; each data point is the result of repeating the same experiment 1,024 times.



D.3.2. Amplitude optimization for state assignment

Using the optimized readout frequency, we instead sweep the IF amplitude of the readout tone, for the three considered states of the two resonators. It is crucial to note that the previous averaging of 1024 is now set to 1, i.e. we perform single-shot readouts from this point onward.

The two qubits are prepared in the ground state by natural relaxation, meaning no state preparation is performed. We then probe the resonators at their optimized frequency 8000 times, with averaging disabled. The result of these 8000 state readouts form a 2D histogram, showing a distinct population in the complex space, which can be seen in figure D3(a). This experiment is repeated, but preparing the qubits in the excited state. The resonator readout 2D histograms now show a population that has moved in the complex space, as is visible in figure D3(b). Similarly, this experiment is repeated for the second-excited state, seen in figure D3(c).

The three populations all feature a population center: a coordinate in the complex plane that corresponds to the average of the 8000 shots performed in order to create that population in the histogram. We now assign each one of the $3 \cdot 8,000$ shots to its closest population center, which creates three distinct zones in the complex plane where any shot that ends up in a certain zone, is assigned a distinct discrete state. Both resonators feature three such zones, seen in figure D3(d).

Since we know which was the intended state for all $3 \cdot 8,000$ shots of this experiment using the optimized readout frequency at some amplitude, we may now also count how many shots missed their target, i.e. how many shots are not assigned to their prepared (intended) state. The results of this counting is shown in figure D3(e). The trace of these matrices is commonly known as the assignment fidelity. Our assignment fidelity is defined from the three states we consider in our readout,

$$F_{\text{Assignment}, R_j} = \frac{|\langle 0 | \Psi_0 \rangle|^2 + |\langle 1 | \Psi_1 \rangle|^2 + |\langle 2 | \Psi_2 \rangle|^2}{3} \quad (\text{D.1})$$

where Ψ_i is the state of the qubit that corresponds to resonator R_j when this qubit was prepared in state $|i\rangle$.

The state-assignment process is repeated for a sweep of 18 different IF amplitudes, ranging from 37 mV to 87 mV for resonator 1, and from 12 mV to 62 mV for resonator 2. We select the readout IF amplitudes which give us the highest assignment fidelities when reading out the two resonators, 82 mV and 45 mV, shown in figure D4. All plots in figure D3 were taken at the optimized frequency and optimized amplitude. We employ no further readout optimization.

From this point onward, whenever we perform single-shot measurements, each readout yields one data point in the complex plane. We assign this data point to its closest population center, see figure D3(d).

We do not employ matched filtering techniques [64] where the demodulated signal is compared to known time traces corresponding to the resonators' frequency responses for the ground, excited, or second excited states of the qubits.

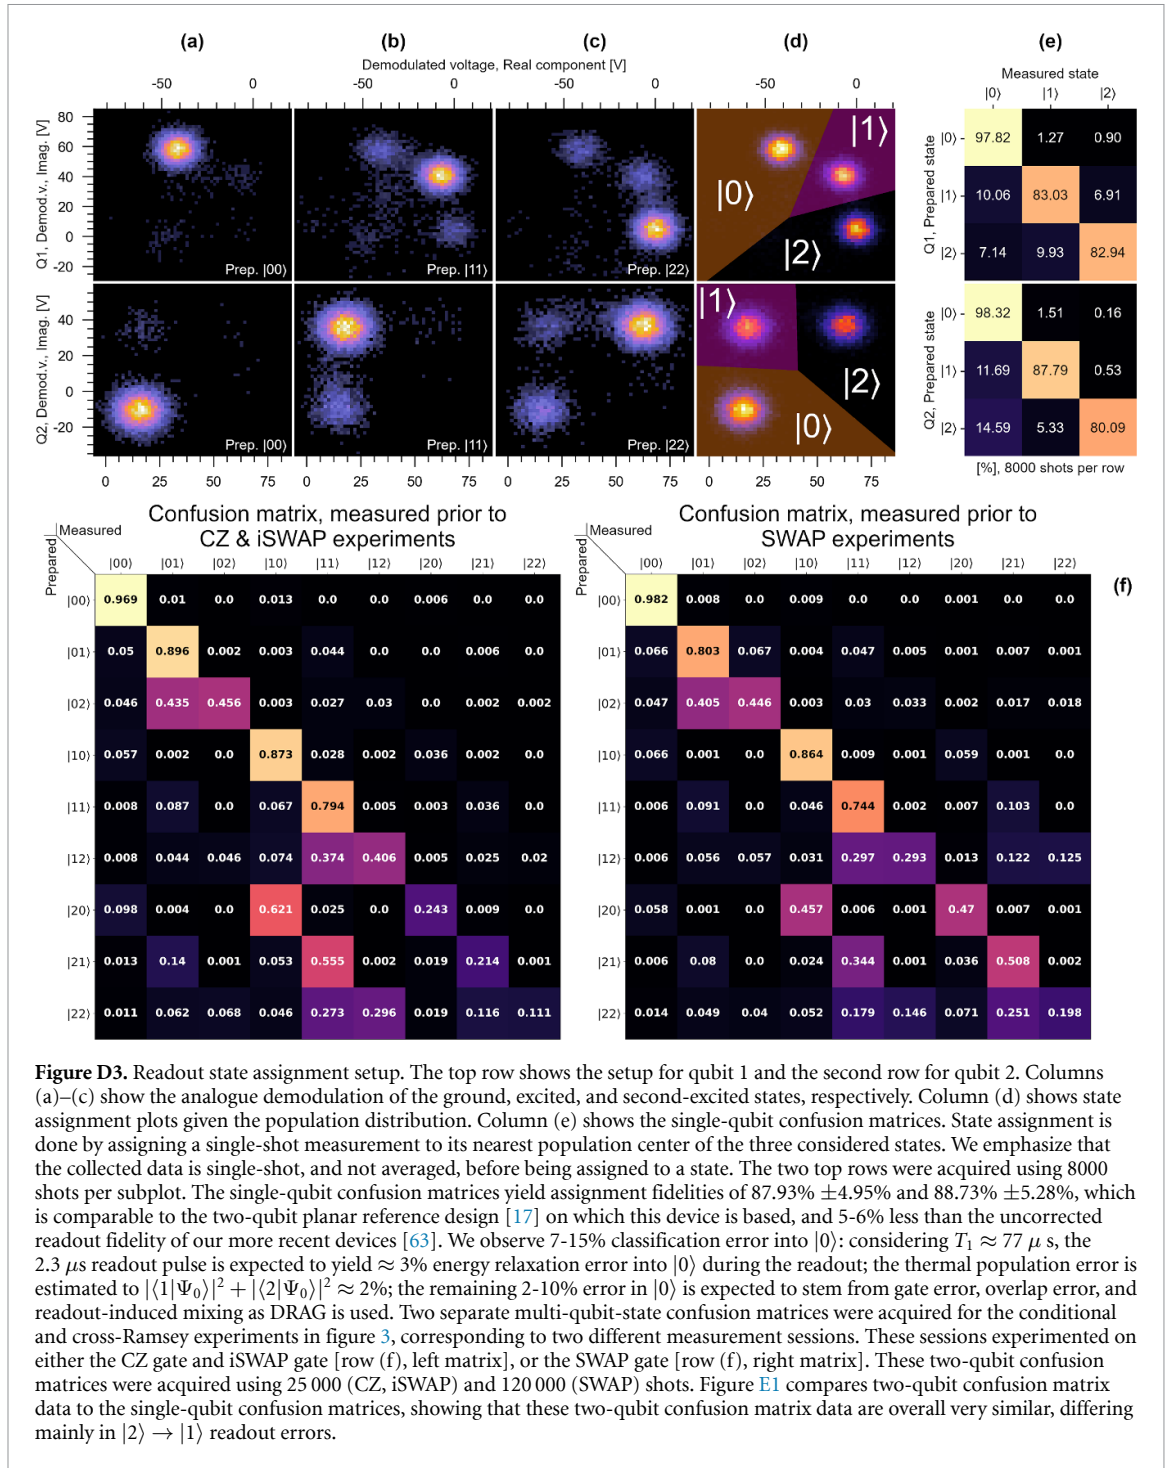


Figure D3. Readout state assignment setup. The top row shows the setup for qubit 1 and the second row for qubit 2. Columns (a)–(c) show the analogue demodulation of the ground, excited, and second-excited states, respectively. Column (d) shows state assignment plots given the population distribution. Column (e) shows the single-qubit confusion matrices. State assignment is done by assigning a single-shot measurement to its nearest population center of the three considered states. We emphasize that the collected data is single-shot, and not averaged, before being assigned to a state. The two top rows were acquired using 8000 shots per subplot. The single-qubit confusion matrices yield assignment fidelities of $87.93\% \pm 4.95\%$ and $88.73\% \pm 5.28\%$, which is comparable to the two-qubit planar reference design [17] on which this device is based, and 5–6% less than the uncorrected readout fidelity of our more recent devices [63]. We observe 7–15% classification error into $|0\rangle$: considering $T_1 \approx 77 \mu\text{s}$, the $2.3 \mu\text{s}$ readout pulse is expected to yield $\approx 3\%$ energy relaxation error into $|0\rangle$ during the readout; the thermal population error is estimated to $|\langle 1|\Psi_0\rangle|^2 + |\langle 2|\Psi_0\rangle|^2 \approx 2\%$; the remaining 2–10% error in $|0\rangle$ is expected to stem from gate error, overlap error, and readout-induced mixing as DRAG is used. Two separate multi-qubit-state confusion matrices were acquired for the conditional and cross-Ramsey experiments in figure 3, corresponding to two different measurement sessions. These sessions experimented on either the CZ gate and iSWAP gate [row (f), left matrix], or the SWAP gate [row (f), right matrix]. These two-qubit confusion matrices were acquired using 25 000 (CZ, iSWAP) and 120 000 (SWAP) shots. Figure E1 compares two-qubit confusion matrix data to the single-qubit confusion matrices, showing that these two-qubit confusion matrix data are overall very similar, differing mainly in $|2\rangle \rightarrow |1\rangle$ readout errors.

D.4. Two-qubit gate tune-up

At this point of the tune-up process, all readout and single-qubit parameters are known, and the DC current bias for the coupler is set. To perform two-qubit gates, we modulate the coupler with a microwave tone giving rise to the magnetic flux

$$\Phi_c(t) = A(t) \cdot \cos(2\pi f_c t + \varphi) + \Phi_{\text{DC}}, \quad (\text{D.2})$$

where f_c is the coupler pulse's carrier frequency, φ is the coupler pulse's phase relative to the start of the quantum circuit, $A(t)$ is the pulse envelope, and Φ_{DC} is the static offset bias which sets the unmodulated frequency of the coupler.

Appendix D.2 explained how Φ_{DC} was set using a current-controlled bias from a current generator. Parameters A , f_c , φ are properties of the IF tone generated by the AWG. As can be seen in figure A1, the AC

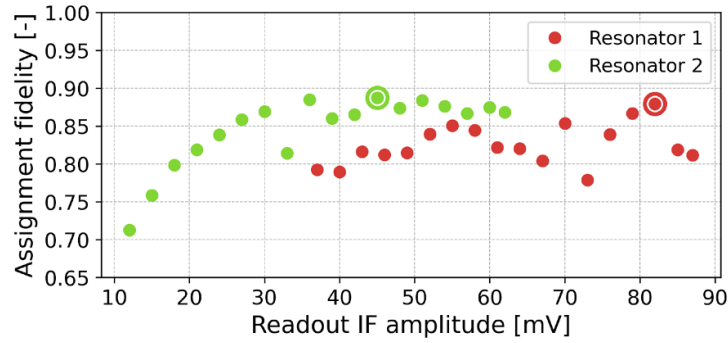


Figure D4. State-assignment fidelity calculated from equation (D.1) as a function of the readout tone IF amplitude. The amplitudes for the two resonators' readout pulses are selected at the highest assignment fidelity, which were $0.8793 \pm 4.95\%$ and $0.8873 \pm 5.28\%$ for Q_1 and Q_2 respectively.

component of equation (D.2) is combined with the bias current at the mixing chamber of the dilution refrigerator.

When performing two-qubit gate tune-up, we are often forced to repeat a subset of the flux bias tune-up of the coupler; we would initially select some coupler bias, then engage in CZ and iSWAP tune-up, knowing that we might have to adjust the coupler bias. Changing the coupler bias requires re-tuning the readout assignment parameter. New state assignment parameters are often similar to those found previously; spectrally, the coupler's transition frequency is typically located ~ 1400 MHz away from the closest considered resonator transition, meaning little change to the readout parameters upon adjustment.

One example of a failure point could be that for some selected CZ amplitude and carrier frequency, population would transfer from $|11\rangle \rightarrow |20\rangle$ at $\frac{\tau_{CZ}}{2}$. But, as τ_{CZ} is swept to determine the CZ gate time, we would not see characteristic population oscillations; instead, we might see a single localized area where population transfers from $|11\rangle \rightarrow |20\rangle$ but does not come back to $|11\rangle$ at time τ_{CZ} . Another example could be that CZ gates can be tuned up, but iSWAP gates would show close to no swapping ability between $|10\rangle \leftrightarrow |01\rangle$.

D.4.1. Parameter estimation

For tuning up CZ and iSWAP gates, our tune-up follows the procedure outlined by Ganzhorn *et al* [20]. We perform conditional Ramsey and cross-Ramsey experiments of these gates to verify their behavior, and ensure that the gates are phase-corrected as can be seen in figure 3.

Our iSWAP drive frequency of 461.51 MHz, e.g. the pulse carrier, lies outside of the 300 MHz baseband bandwidth of the AWG, which results in a significant roll-off in amplitude output from the instrument as compared to what was defined in software. This roll-off does not prevent the iSWAP gate from being tuned up as the AWG is capable of making up for the amplitude lost to roll-off.

In superconducting quantum processor architectures that use fixed-frequency qubits with tunable couplers, local qubit phase adjustments are required after applying two-qubit gates, as the coupler drive pulse itself is dispersively shifting the coupler [20], consequently changing the qubit frequencies as was seen in figure D1. Using virtual-Z gates, we adjust for these local phase offsets of the iSWAP and CZ gates, in a similar manner to [20].

As is shown in [4], using virtual-Z gates for local phase adjustment requires adjusting the coupler drive for the iSWAP gates with $\varphi_c = \varphi_2 - \varphi_1$, where φ_1 and φ_2 are the local qubit phase adjustments done for the iSWAP gate. We do not adjust the phase of the coupler drive for CZ gates, as the CZ gate is invariant to the phase of the coupler pulse [20, 48].

In addition to this coupler phase offset, the coupler drive must be further phase adjusted to account for frame tracking, as the iSWAP gate is not resonant with the qubit frames. The ideal frequency f_{2qf} at which an iSWAP gate is performed, equals the frequency corresponding to the energy difference between $||10\rangle - |01\rangle$. However, as the qubits move in frequency due to the coupler being affected during the coupler drive pulse playing, the tuned-up iSWAP gate carrier frequency $f_{iSWAP} \neq f_{2qf}$. This frequency difference between the coupler drive and the two-qubit frame leads to a phase offset accumulating during the idle time between two subsequent iSWAP gates, which must be compensated for. To experimentally measure f_{2qf} for iSWAP gates, consider the quantum circuit in figure D5.

In the cross-Ramsey interferometry experiments performed for this work, we are only concerned with the coupler phase of the iSWAP gates at two distinct moments: the first moment is after one initial \sqrt{X} gate




Figure D5. Experiment for measuring the two-qubit frame f_{2qf} . The gate element shown as τ is a time delay between the two iSWAP gates. Begin the experiment by assuming some coupler drive frequency f close to the tuned-up $f_{i\text{SWAP}}$. As τ increases, the readout will show a Ramsey-like oscillation dependent on the detuning between f and f_{2qf} . Once f matches the frequency of the two-qubit frame f_{2qf} , the oscillation frequency of the Ramsey-like readout will be 0 Hz. Therefore, a two-dimensional sweep of f and τ reveals f_{2qf} graphically, similar to a two-dimensional Ramsey interferometry.


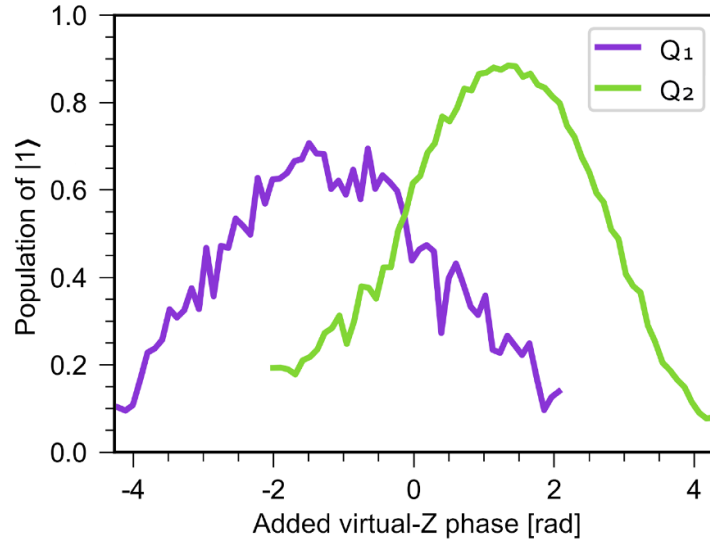


Figure D6. Parameter estimation of the SWAP local phase correction, shown as the population in $|1\rangle$ on a target qubit $q_t = \{Q_1, Q_2\}$ as the virtual-Z phase φ is swept. The control qubit $q_c = \{Q_2, Q_1\}$ is prepared in $|1\rangle$ with an initial X gate. The local qubit phase corrections are selected as -1074.1 mrad and $+1253.0$ mrad for qubits Q_1 and Q_2 respectively. The traces shown are not error-corrected for SPAM errors, hence the comparably low population at the two maxima. Each data point was taken using 30 000 single shots, where states $\{|0\rangle, |1\rangle, |2\rangle\}$ were considered in the state assignment. Recall that the S^\dagger gates are implemented as 0 ns virtual-Z gates.

has played in the iSWAP cross-Ramsey interferometry, and the second moment is after the Hadamard gate has played after preparing $|q_c q_t\rangle = |0\bar{i}\rangle$, or $|1\bar{i}\rangle$, in the SWAP cross-Ramsey interferometry. In the case of the first moment, the frame-tracked coupler phase of the iSWAP gate was acquired from the iSWAP tune-up procedure. In the second moment, we measured the acquired local phase offset that occurred from playing the same iSWAP gate but at a different time in the quantum circuit, i.e. after executing the first Hadamard of the SWAP sequence in figure 1(d) on a prepared input state. Figure D6 demonstrates how the local phase of the qubits was impacted from playing the same iSWAP coupler phase as was done for the iSWAP cross-Ramsey interferometry.

We adjust the local phases after the CZ gate has finished playing with the phase value that that places the investigated qubit as close as possible to $|1\rangle$ in figure D6, as this state is the expected outcome. As such, for our work, we ensure frame-tracked iSWAP gates through the means of local phase updates. However, we emphasize that for a generic implementation using iSWAP gates, the frame tracking could just as well have been done on the coupler drive.

Appendix E. SPAM error mitigation

Let us consider one (any) of the twelve traces of data from figure 3. For each step along the x -axis of an experiment, we measure some data point \vec{y} that contains the observed probability of all nine considered

two-qubit states, (3 states)^(2 qubits):

$$\vec{y} = [|\alpha_{00}|^2, |\alpha_{01}|^2, |\alpha_{02}|^2, |\alpha_{10}|^2, |\alpha_{11}|^2, \dots, |\alpha_{22}|^2]^T \quad (\text{E.1})$$

where α_i is the complex amplitude of the two-qubit state $|i\rangle \in \{0, 1, 2\}^2$ of the system's wave function Ψ at measurement.

Before the conditional and cross-Ramsey experiments, we intentionally prepared Ψ in all nine two-qubit states, and measured the outcome from thousands of iterations as shown in figure D3. From these empirical statistics, we construct a right stochastic [65] matrix:

$$T = \begin{bmatrix} |\langle 00 | \Psi_{00} \rangle|^2 & |\langle 01 | \Psi_{00} \rangle|^2 & \dots & |\langle 22 | \Psi_{00} \rangle|^2 \\ |\langle 00 | \Psi_{01} \rangle|^2 & |\langle 01 | \Psi_{01} \rangle|^2 & & \\ \vdots & & \ddots & \\ |\langle 00 | \Psi_{22} \rangle|^2 & & & |\langle 22 | \Psi_{22} \rangle|^2 \end{bmatrix} \quad (\text{E.2})$$

where each row shows the most likely distribution of measurement outcomes from observing the prepared two-qubit state Ψ_i , given that some two-qubit state $|i\rangle \in \{0, 1, 2\}^2$ was prepared for that row.

Each element of the matrix shows the conditional probability $P(|j\rangle \text{ measured} \mid |i\rangle \text{ prepared})$, where $j, i \in \{0, 1, 2\}^2$. Our experimentally acquired values for matrix T in equation (E.2) was shown in the bottom row of figure D3.

We now assume that our acquired data point \vec{y} is the result of manipulating some 'true value' data point \vec{x} ,

$$\vec{x} = [|\beta_{00}|^2, |\beta_{01}|^2, |\beta_{02}|^2, |\beta_{10}|^2, |\beta_{11}|^2, \dots, |\beta_{22}|^2]^T \quad (\text{E.3})$$

which was the true outcome of the experiment before \vec{x} was manipulated by T , resulting in the actual measured output \vec{y}

$$\vec{y} = T\vec{x}. \quad (\text{E.4})$$

Remember that for every data point in any given trace from figure 3, we acquired a vector $\vec{y} = T\vec{x}$ during the experiment itself. Figure E1 highlights that T must be acquired by preparing and measuring all possible qubit state combinations, and that assembling T from its single-qubit confusion matrices yields different results.

Once we have acquired T , we use an implementation [66] of the STIR algorithm [67] for minimizing the difference

$$\min \left(\frac{1}{2} \cdot |\vec{y} - T\vec{x}|^2 \right). \quad (\text{E.5})$$

Specifically, since \vec{y} and T were found empirically and are thus known, the algorithm finds the vector \vec{x} that minimizes equation (E.5), meaning that \vec{x} is now our best guess as to the non-distorted outcome of the experiments that resulted in \vec{y} when distorted by T , for all data points.

Using notation $|q_c q_t\rangle$: in figure 3, we plotted the calculated $|\beta_{01}|^2$ and $|\beta_{11}|^2$ from equation (E.3) for the CZ conditional Ramsey traces. And similarly for the iSWAP and SWAP cross-Ramsey traces, we plot $|\beta_{10}|^2$ and $|\beta_{11}|^2$. Remember that the roles of the qubits in the top row of figure 3 is $|q_c q_t\rangle = |Q_2 Q_1\rangle$ for the CZ experiments, and $|q_c q_t\rangle = |Q_1 Q_2\rangle$ for the iSWAP and SWAP experiments. In the bottom row of figure 3, the two qubits exchange roles in the quantum circuits.

Table E1 shows numerical results of this best-guess reconstructed $|\beta_{01}|^2$ and $|\beta_{11}|^2$ for the CZ experiments, and $|\beta_{10}|^2$ and $|\beta_{11}|^2$ for the iSWAP and SWAP experiments. We note that the mean squared error (MSE) on average is 0.007 where 0 corresponds to no error between datapoints and the ideal curve; where, the CZ datasets contained 32 samples per trace and the iSWAP and SWAP datasets contained 16 samples per trace. A notable outlier is seen with the $|\beta_{11}|^2$ reconstruction for the CZ trace where $|q_c q_t\rangle = |Q_2 Q_1\rangle$ at MSE = 0.02033, which is twice as large as the second largest MSE. We attribute this outlier to the two notably deviating clusters of data in figure 3(d), $P(|11\rangle)$, which were discussed in section 5.

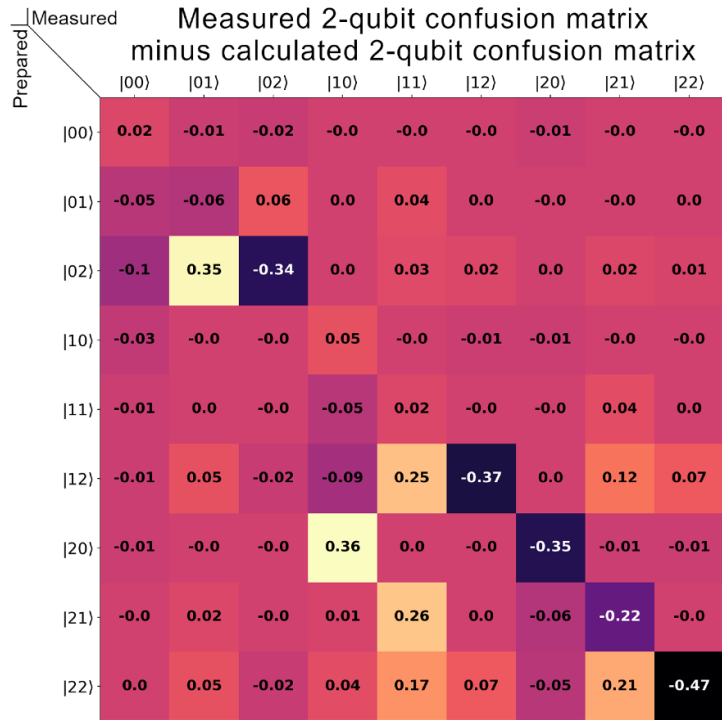


Figure E1. Matrix showing $C_{\text{Measured}} - C_{\text{Constructed}}$, where C_{Measured} is the confusion matrix measured prior to the SWAP experiments from figure D3, and $C_{\text{Constructed}}$ is the two-qubit confusion matrix acquired from combining the measured single-qubit confusion matrices from figure D3. This matrix highlights effects that stem from two-qubit readout, such as measurement-induced state decay and readout crosstalk, as seen below the diagonal. From preparing $|12\rangle$, we see that C_{Measured} yielded 0.19 more population above the diagonal than $C_{\text{Constructed}}$, which we believe hints at measurement-induced transition errors. Apart from $|\langle 12 | \Psi_{22} \rangle|^2$, the remaining six combinations where a qubit could decay $|2\rangle \rightarrow |1\rangle$, all accounted for the largest observed population increases: {0.36, 0.35, 0.26, 0.25, 0.21, 0.17}.

Table E1. Conditional Ramsey and cross-Ramsey data of figure 3. P , swing is the fitted population swing where the ideal swing is from 0.0 to 1.0, i.e. the expected ideal number is 1.0; P , Δ offset is the difference $m_{\text{ideal}} - m$ between the expected ideal offset $m_{\text{ideal}} = 0.5$ and the fitted offset m , where $m_{\text{ideal}} = 0.5$ denotes that the expected population swing is oscillating about 0.5, and P , Δ offset = 0.0 means no difference between the expected ideal and the experimental data; Δ phase is similarly the deviation in the oscillating population's phase, where Δ phase = 0.0 means no difference between the expected phase and the experimental data; and MSE denotes the mean squared error: $\text{MSE} = \frac{\sum (y-p)^2}{N}$, where y is the datapoint, p is its predicted ideal value, and N is the number of datapoints. An MSE of 0 means no error between the datapoints and the ideal prediction.

Operation	Output state	P , swing [—]	P , Δ offset [—]	Δ phase [mrad]	MSE [—]
CZ	$ 01\rangle$	0.987(1)	0.027(11)	37(8)	0.00 117
	$ 11\rangle$	0.967(3)	0.077(15)	-57(15)	0.00 632
CZ	$ 10\rangle$	1.008(13)	-0.006(8)	71(12)	0.00 114
	$ 11\rangle$	0.868(41)	0.150(17)	-16(52)	0.02 033
iSWAP	$ 10\rangle$	1.079(18)	-0.023(6)	33(16)	0.00 548
	$ 11\rangle$	1.064(29)	-0.060(10)	-85(26)	0.00 915
iSWAP	$ 01\rangle$	1.041(5)	0.064(2)	85(4)	0.00 401
	$ 11\rangle$	1.057(9)	-0.038(3)	75(8)	0.00 570
SWAP	$ 10\rangle$	0.849(33)	0.049(16)	-96(42)	0.00 997
	$ 11\rangle$	0.859(44)	0.023(27)	-68(47)	0.00 584
SWAP	$ 01\rangle$	0.954(39)	-0.055(20)	223(44)	0.00 953
	$ 11\rangle$	0.993(23)	-0.050(12)	103(26)	0.00 327

References

- [1] Kitaev A Y, Shen A H, Vyalii M N Senechal L J 1999 Classical and quantum computation (MTsNMO-CheRo) *Classical and Quantum Computation* (American Mathematical Society) ch 8, pp 65–82
- [2] Barenco A, Bennett C H, Cleve R, DiVincenzo D P, Margolus N, Shor P, Sleator T, Smolin J A and Weinfurter H 1995 *Phys. Rev. A* **52** 3457–67
- [3] Leymann F and Barzen J 2020 *Quantum Sci. Technol.* **5** 044007
- [4] Abrams D M, Didier N, Johnson B R, da Silva M P da and Ryan C A 2020 *Nat. Electron.* **3** 744–50
- [5] McEwen M, Bacon D and Gidney C 2023 *Quantum* **7** 1172
- [6] Yan G, Wu W, Chen Y, Pan K, Lu X, Zhou Z, Wang Y, Wang R and Yan J 2024 Quantum circuit synthesis and compilation optimization: overview and prospects (arXiv:2407.00736)
- [7] Iverson J K and Preskill J 2020 *New J. Phys.* **22** 073066
- [8] Abad T, Fernández-Pendás J, Frisk Kockum A and Johansson G 2022 *Phys. Rev. Lett.* **129** 150504
- [9] Abad T, Schattner Y, Frisk Kockum A and Johansson G 2023 *Quantum* **9** 1684
- [10] Blais A, Grimsmo A L, Girvin S M and Wallraff A 2021 *Rev. Mod. Phys.* **93** 025005
- [11] Zhang J, Vala J, Sastry S and Whaley K B 2003 *Phys. Rev. A* **67** 042313
- [12] Crooks G E 2024 Gates, States and Circuits White paper *Tech. Note 014 v0.11.0 Beta* (Normal Computing Corporation) pp 5–69 (available at: <https://threeplusone.com/gates>)
- [13] Sung Y et al 2021 *Phys. Rev. X* **11** 021058
- [14] Lao L, Korotkov A, Jiang Z, Mruczkiewicz W, O'Brien T E and Browne D E 2022 *Quantum Sci. Technol.* **7** 025021
- [15] Huang C et al 2023 *Phys. Rev. Lett.* **130** 070601
- [16] Preskill J 2018 *Quantum* **2** 79
- [17] Bengtsson A et al 2020 *Phys. Rev. Appl.* **14** 034010
- [18] Cervera-Lierta A 2018 *Quantum* **2** 114
- [19] Williams C P 2011 *Quantum Gates* 2nd edn (Springer) ch 2, p 110
- [20] Ganzhorn M et al 2020 *Phys. Rev. Res.* **2** 033447
- [21] Jin X Y, Cicak K, Parrott Z, Kotler S, Lecocq F, Teufel J, Aumentado J, Kapit E and Simmonds R W 2024 Fast, tunable, high fidelity cz-gates between superconducting qubits with parametric microwave control of zz-coupling (arXiv:2305.02907v3)
- [22] Fowler A G, Mariantoni M, Martinis J M and Cleland A N 2012 *Phys. Rev. A* **86** 032324
- [23] Roy T, Li Z, Kapit E and Schuster D I 2023 *Phys. Rev. Appl.* **19** 064024
- [24] Roth M, Ganzhorn M, Moll N, Filipp S, Salis G and Schmidt S 2017 *Phys. Rev. A* **96** 062323
- [25] McKay D C, Filipp S, Mezzacapo A, Magesan E, Chow J M and Gambetta J M 2016 *Phys. Rev. Appl.* **6** 064007
- [26] Xu Y et al 2020 *Phys. Rev. Lett.* **125** 240503
- [27] Stehlik J et al 2021 *Phys. Rev. Lett.* **127** 080505
- [28] Roth M, Moll N, Salis G, Ganzhorn M, Egger D J, Filipp S and Schmidt S 2019 *Phys. Rev. A* **99** 022323
- [29] Poletto S et al 2012 *Phys. Rev. Lett.* **109** 240505
- [30] Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information* 10th edn (Cambridge University Press)
- [31] Zhang J, Vala J, Sastry S and Whaley K B 2004 *Phys. Rev. Lett.* **93** 020502
- [32] Han J X, Zhang J, Xue G M, Yu H and Long G 2024 Protecting logical qubits with dynamical decoupling (arXiv:2402.05604)
- [33] Rezakhani A T 2004 *Phys. Rev. A* **70** 052313
- [34] Ni M, Ma R L, Kong Z Z, Xue X, Zhu S K, Wang C, Li A R, Chu N, Liao W Z, Cao G, Wang G L, Guo G C, Hu X, Jiang H W, Li H O and Guo G P 2023 *Nano Lett.* **25** 3766–72
- [35] Taylor J R, Foulk N L and Das Sarma S 2024 *Phys. Rev. B* **109** 165403
- [36] Wei K X, Lauer I, Pritchett E, Shanks W, McKay D C and Javadi-Abhari A 2024 *PRX Quantum* **5** 020338
- [37] Lukac M, Perkowski M and Kameyama M 2010 *Evolutionary Quantum Logic Synthesis of Boolean Reversible Logic Circuits Embedded in Ternary Quantum Space Using Structural Restrictions (IEEE Congress on Evolutionary Computation)* (IEEE) p 6
- [38] Sangouard N, Lacour X, Guérin S and Jauslin H R 2005 *Phys. Rev. A* **72** 062309
- [39] Venturelli D, Do M, Rieffel E and Frank J 2018 *Quantum Sci. Technol.* **3** 025004
- [40] Mondal A and Parhi K K 2024 *IEEE Circuits Syst. Mag.* **24** 33–51
- [41] Kosen S et al 2022 *Quantum Sci. Technol.* **7** 035018
- [42] Kosen S et al 2024 *PRX Quantum* **5** 030350
- [43] Koch J, Yu T M, Gambetta J, Houck A A, Schuster D I, Majer J, Blais A, Devoret M H, Girvin S M and Schoelkopf R J 2007 *Phys. Rev. A* **76** 042319
- [44] Bianchetti R A 2010 Control and readout of a superconducting artificial atom *PhD Dissertation* ETH Zurich (available at: <http://hdl.handle.net/20.500.11850/152306>)
- [45] Geller M R 2020 *Quantum Sci. Technol.* **5** 03LT01
- [46] Bialczak R C et al 2010 *Nat. Phys.* **6** 409–13
- [47] Niskanen A O, Nakamura Y and Tsai J-S 2006 *Phys. Rev. B* **73** 094506
- [48] Reagor M et al 2018 *Sci. Adv.* **4** eaao3603
- [49] Kubo K and Goto H 2023 *Appl. Phys. Lett.* **122** 064001
- [50] Krantz P, Kjaergaard M, Yan F, Orlando T P, Gustavsson S and Oliver W D 2019 *Appl. Phys. Rev.* **6** 021318
- [51] Warren C W et al 2023 *npj Quantum Inf.* **9** 44
- [52] Baldwin C H, Kalev A and Deutsch I H 2014 *Phys. Rev. A* **90** 012110
- [53] Krizan C 2024 Dataset and code used for Quantum SWAP gate realized with CZ and iSWAP gates in a superconducting architecture *Zenodo* (<https://doi.org/10.5281/zenodo.11149402>)
- [54] Krinner S, Storz S, Kurpiers P, Magnard P, Heinsoo J, Keller R, Lütolf J, Eichler C and Wallraff A 2019 *EPJ Quantum Technol.* **6** 2
- [55] Gerry C and Knight P 2004 Field Quantization *Introductory Quantum Optics* (Cambridge University Press) ch 2, pp 10–42
- [56] Motzoi F, Gambetta J M, Reberntrost P and Wilhelm F K 2009 *Phys. Rev. Lett.* **103** 110501
- [57] Bengtsson A 2020 Quantum information processing with tunable and low-loss superconducting circuits *PhD Dissertation* Chalmers University of Technology, Department of Microtechnology and Nanoscience, Quantum Technology laboratory (available at: <https://research.chalmers.se/en/publication/516196>)
- [58] Ku H H 1966 *J. Res. Natl Bureau Stand. C* **70C** 264
- [59] Doane D P 1976 *Am. Stat.* **30** 181–3

- [60] Chen Z 2018 Metrology of Quantum control and measurement in superconducting qubits *PhD Dissertation* University of California (available at: <https://web.physics.ucsb.edu/~martinisgroup/theses/Chen2018.pdf>)
- [61] Jaklevic R C, Lambe J, Silver A H and Mercereau J E 1964 *Phys. Rev. Lett.* **12** 159–60
- [62] Clarke J and Braginski A I 2004 *Squid Theory* (Wiley) ch 2, pp 29–87
- [63] Chen L et al 2023 *npj Quantum Inf.* **9** 26
- [64] Tholén M O et al 2022 *Rev. Sci. Instrum.* **93** 104711
- [65] Szabo F E 2015 *Stochastic Matrix The Linear Algebra Survival Guide* ed F E Szabo (Academic) p 363
- [66] Baumgarten C, Brett M, Burovski E, Carey C J, Gommers R, Haberland M, Harris C, Larson E, Mayorov N, van Mulbregt P, Nelson A, Perktold J, Polat I, Reddy T, Virtanen P, Weckesser W, Wilson J (additional contributors) 2023 `scipy.optimize.lsqr_linear` The SciPy community, SciPy 1.10.1 (available at: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.lsqr_linear.html)
- [67] Branch M A, Coleman T F and Li Y 1999 *SIAM J. Sci. Comput.* **21** 1–23