

Above 99.9% Fidelity Single-Qubit Gates, Two-Qubit Gates, and Readout in a Single Superconducting Quantum Device

Fabian Marxer^{1,*}, Jakub Mrozek^{1,*}, Joona Andersson¹, Leonid Abdurakhimov¹,
 Janos Adam¹, Ville Bergholm¹, Rohit Beriwal¹, Chun Fai Chan¹, Saga Dahl¹,
 Soumya Ranjan Das², Frank Deppe², Olexiy Fedorets¹, Zheming Gao¹,
 Alejandro Gomez Friero¹, Daria Gusenkova², Andrew Guthrie¹, Tuukka Hiltunen¹, Hao Hsu²,
 Eric Hyypää¹, Joni Ikonen¹, Sinan Inel¹, Shan W. Jolin¹, Azad Karis¹, Seung-Goo Kim¹,
 William Kindel², Anton Komlev¹, Miikka Koistinen¹, Roope Kokkonen¹, Snigdha Kumar¹,
 Hsiang-Sheng Ku², Julia Lamprich², Sami Laine¹, Alessandro Landra¹, Lan-Hsuan Lee¹,
 Nizar Lethif², Per Liebermann¹, Wei Liu¹, Kunal Mitra¹, Tuomas Mylläri¹,
 Caspar Ockeloen-Korppi¹, Tuure Orell¹, Alexander Plyshch¹, Jukka Rabinä¹, Arthur Rebello¹,
 Michael Renger², Outi Reentilä¹, Jussi Ritvas¹, Sampo Saarinen¹, Otto Salmenkivi¹,
 Matthew Sarsby¹, Mykhailo Savitskyi¹, Ville Selinmaa¹, Matthew Steggle¹, Eelis Takala¹,
 Ivan Takmakov¹, Brian Tarasinski¹, Jani Tuorila¹, Alpo Välimaa¹, Jeroen Verjauw²,
 Jaap Westdorp¹, Nicola Wurz², Wei Qiu¹, Lihuang Zhu¹, Juha Hassel¹, Johannes Heinsoo¹,
 Attila Geresdi² and Antti Vepsäläinen^{1,§}
¹*IQM Quantum Computers, Espoo 02150, Finland*
²*IQM Quantum Computers, Munich 80992, Germany*



(Received 4 September 2025; revised 3 February 2026; accepted 26 March 2026; published 20 May 2026)

Achieving high-fidelity single-qubit gates, two-qubit gates, and qubit readout is critical for building scalable, error-corrected quantum computers. However, device parameters that enhance one operation often degrade the others, making simultaneous optimization challenging. Here, we demonstrate that careful tuning of qubit-coupler coupling strengths in a superconducting circuit with two transmon qubits coupled via a tunable coupler enables high-fidelity single- and two-qubit gates, without compromising readout performance. Furthermore, we introduce a new calibration protocol for diabatic controlled-Z gates, phased-averaged leakage error amplification (PALEA), which enables efficient suppression of coherent gate errors and leakage to noncomputational states. Using PALEA in combination with careful tuning of qubit-coupler coupling strengths, we achieve a 40 h-averaged CZ gate fidelity of 99.93 %, simultaneous single-qubit gate fidelities of 99.98 %, and readout fidelities over 99.94 % in a single device. Our results demonstrate a viable path toward scaling up superconducting quantum processors while maintaining consistently high gate and readout fidelities.

DOI: [10.1103/n86s-2b88](https://doi.org/10.1103/n86s-2b88)

I. INTRODUCTION

Scalable, fault-tolerant quantum computing requires low error rates across all core operations: single-qubit gates, two-qubit gates, qubit readout, qubit reset, and

leakage removal. To enable resource-efficient implementation of logical operations, these error rates must fall well below the thresholds set by quantum error correction protocols [1–3]. Among various quantum hardware platforms, superconducting qubits—particularly transmons—have emerged as a scalable and practical platform for quantum computing, enabled by fast gate and readout operations, flexible control, and mature fabrication techniques [4–6].

In such systems, single-qubit gate fidelities now routinely exceed 99.9 % [3,7–10], and comparable fidelities have been achieved for two-qubit gates [11–14]. This progress has been driven by recent advances in tunable coupler designs, such as grounded transmon couplers [11,15–19], floating transmon couplers [20–22],

*These two authors contributed equally.

†Contact author: fabian@meetiqm.com

‡Contact author: jakub@meetiqm.com

§Contact author: avepsalainen@meetiqm.com

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.

inductively mediated coupling schemes [13,14], and double-transmon couplers [12], all of which enable dynamically controlled qubit-qubit interactions while suppressing residual coupling [23].

At the same time, high-fidelity readout [24–29] has become increasingly critical, not only for final measurements but also for mid-circuit measurements required in quantum error correction [3,30,31] and feedback-based protocols [32–34]. Recent progress has been enabled by improved measurement protocols and readout hardware, including shelving-based approaches [26,35–37] as well as optimized parametric amplification chains [38–41]. In particular, the latest experiments have reported readout fidelities exceeding 99.9 % using optimized dispersive readout with high-performance amplifiers [28,29].

Despite these advances, simultaneously achieving high-fidelity single-qubit gates, two-qubit gates, and readout in a single device remains a major challenge. Device parameters optimized for one operation often degrade another. A key trade-off lies in the qubit-coupler coupling strength: strong coupling facilitates fast entangling gates, but also increases qubit-qubit hybridization, leading to simultaneous single-qubit gate errors. Similarly, incorporating high-fidelity readout without degrading qubit coherence or gate performance requires careful architectural integration.

In this work, we demonstrate an optimization strategy that addresses the trade-offs between single-qubit gates and two-qubit gates, while maintaining high-fidelity readout. Our device consists of two transmon qubits coupled via a tunable coupler. To guide the design, we simulate both single-qubit errors—dominated by hybridization-induced crosstalk and incoherent relaxation—and two-qubit errors arising from leakage, incoherent processes, and residual swap interactions. By sweeping the qubit-coupler coupling strength, we identify a parameter regime that minimizes the combined gate error, which we then target in the device fabrication. Furthermore, to enable precise calibration of the conditional-Z (CZ) gate and to suppress coherent over- and under-rotation in the $\text{span}(|11\rangle, |02\rangle)$ subspace, we introduce a novel gate error characterization method, which we call the phase-averaged leakage error amplification (PALEA) protocol. This method is specifically designed to coherently amplify CZ gate errors, especially their associated population leakage to the second excited state of the higher frequency qubit, while remaining robust against other leakage processes. Using PALEA to improve leakage calibration accuracy, we demonstrate that, using the same number of repetitions in the experiment, the leakage can be systematically reduced by at least a factor of two compared to standard leakage amplification methods [18], significantly enhancing compatibility with quantum error correction codes. Together with the qubit-coupler coupling optimization, we achieve a CZ gate fidelity of 99.93 % and single-qubit gate fidelities above 99.98 %.

Simultaneously, our readout architecture is designed to enable high-fidelity readout within a few hundred nanoseconds. Each qubit is coupled to its own readout resonator, which in turn is connected to an individual Purcell filter, not only to protect the qubit from Purcell decay but also to suppress the off-resonant driving of untargeted readout resonators [42,43]. The design parameters are chosen to achieve an optimal ratio of dispersive shift χ to resonator linewidth κ (targeting $\chi/\kappa \approx 0.5$), enabling fast readout while maintaining high signal contrast [25]. Using a 240 ns readout pulse combined with shelving to the second excited state and a traveling-wave parametric amplifier (TWPA), we achieve simultaneous readout fidelities above 99.94 % for both qubits in our setup [26,35–37]. Using a pulse of the same duration but without the shelving gate, we additionally demonstrate a measurement with quantum non-demolition-ness (QNDness) of 99.3 % for both qubits, where QNDness quantifies how well the measurement preserves the qubit state across repeated measurements [29,44,45].

Together, these results establish that high-fidelity gate and readout operations are achievable. Notably, the layout and coupling topology of our approach is readily scalable to larger quantum processing units (QPUs), including the square-grid architecture necessary for surface code implementations [5,46], although system-level effects such as crosstalk and frequency crowding will need to be addressed in future multi-qubit implementations. Specifically, the qubit-coupler architecture demonstrated here—where the qubits are coupled to a single tunable coupler—can be naturally extended to a 2D grid by coupling each qubit to four independent couplers without significantly altering the coupling strengths or the anharmonicities of the qubits and couplers. This ensures that gate fidelity can in principle be preserved even as the system scales. Moreover, the chip layout offers sufficient space to integrate a dedicated Purcell filter for each readout resonator on the same layer as the qubits [22,43,47], enabling scalable, high-fidelity readout without requiring multilayer fabrication or complex routing.

II. OPTIMAL QUBIT-COUPLER COUPLING STRENGTH

Our device consists of two flux-tunable floating transmon qubits coupled via a flux-tunable floating transmon coupler [Fig. 1(a)], see Appendix A for the full experimental setup. The coupling structure is similar to the architecture introduced in Ref. [22]. In this qubit-coupler-qubit topology, the couplings between Qubit 1 (Q_1) and the coupler, g_{1c} , Qubit 2 (Q_2) and the coupler, g_{2c} , as well as the direct qubit-qubit interaction g_{12} between Q_1 and Q_2 play a crucial role in the performance of both single- and two-qubit gates. For simplicity, we define a symmetric coupling $g_{qc} := g_{1c} = g_{2c}$ for the numerical analysis.




FIG. 1. (a) Quasi-lumped-element circuit diagram of the device: two transmon qubits (blue and orange) are coupled via a central floating coupler qubit (red). Each qubit is equipped with its own XY drive line (green), Z flux line (violet), and a readout structure consisting of a readout resonator coupled to a Purcell filter (lime green). The readout structure is capacitively coupled to a common transmission line (yellow). The coupler qubit is flux-controlled but not driven with microwave pulses. Gray arrows indicate the effective coupling strengths between Q_1 and the coupler (g_{1c}), Q_2 and the coupler (g_{2c}), and the direct coupling between the two qubits (g_{12}). (b) Simulated single-qubit gate, CZ gate, and two-qubit Clifford gate errors as a function of the qubit-coupler coupling strength $g_{qc}/2\pi$, assuming an identical qubit-coupler coupling for both qubits, $g_{qc} := g_{1c} = g_{2c}$, and a fixed qubit relaxation time of $T_1 = 100 \mu\text{s}$. The vertical dotted line marks the coupling strength used in this work ($g_{qc}/2\pi = 66 \text{ MHz}$). (c) Qubit-coupler coupling strengths, g_{qc} , versus qubit relaxation times reported in selected previous works [17,18,20–22,48]. The black line represents the simulated optimal coupling that minimizes the two-qubit Clifford error, and the shaded green region denotes the $\pm 5\%$ band around this optimum.

While tunable couplers are usually employed to cancel the static ZZ interaction between the computational states, the bare qubit states typically remain hybridized for most qubit-qubit detunings in this type of circuit [18,23,49], see Appendix B 1. This residual hybridization gives rise to hybridization crosstalk—a coherent crosstalk mechanism in which driving one qubit unintentionally excites or induces phase shifts in other unwanted eigenstates, even in the absence of direct coupling between the driveline and those qubits. Consequently, minimizing the transverse coupling strength—governed by g_{qc} and g_{12} —helps reduce such unwanted crosstalk, improving simultaneous single-qubit gate fidelities.

However, two-qubit gates benefit from a stronger g_{qc} . First, higher couplings allow faster CZ gate operation, reducing incoherent error from relaxation and dephasing. Second, with stronger coupling, the same effective qubit-qubit interaction strength can be achieved with a smaller frequency excursion of the coupler. This arrangement reduces the hybridization of the computational states with the coupler mode during the gate, thereby mitigating the impact of the coupler’s typically lower coherence on the gate fidelity [17,18,48].

To explore this trade-off quantitatively, we simulate gate errors as a function of qubit-coupler coupling strength g_{qc} , see Fig. 1(b). The simulations assume qubit relaxation times of $T_1^{Q1} = T_1^{Q2} = 100 \mu\text{s}$, together with qubit and coupler frequencies and anharmonicities being similar to those of our fabricated device (see Appendix A). In the model, both qubits are assumed to have identical T_1 and pure dephasing times T_ϕ . The coupler is assigned a relaxation time that is four times shorter than that of the qubits, reflecting its higher energy participation ratio (EPR) in lossy interfaces [50], which—based on our present design—is estimated to be four times larger. All components are modeled as coupled three-level Duffing oscillators [4]. Additionally, to isolate the impact of g_{qc} , the direct qubit-qubit coupling g_{12} is adjusted for each value of g_{qc} such that the coupler idling frequency (where $\zeta = 0$) remains fixed throughout the sweep.

For single-qubit gates, we use master equation simulations to evaluate simultaneous $\sqrt{X}$ operations on both qubits to capture both incoherent and coherent (hybridization-induced) errors. We also simulate the CZ gate to extract leakage and incoherence contributions, see Appendix B for further details. The resulting trends show that single-qubit errors increase with g_{qc} , while CZ gate errors decrease. To evaluate overall gate performance, we calculate the estimated two-qubit Clifford gate error by weighting the single-qubit gate and two-qubit gate errors by their average prevalence in a random Clifford gate: $4.65 \sqrt{X}$ gates and 1.5 CZ gates per Clifford [2,51]. These weighted error contributions yield the Clifford error curve shown in Fig. 1(b).

From this analysis, we identify an optimal coupling strength near $g_{qc}/(2\pi) = 75 \text{ MHz}$ that minimizes the total Clifford error. The fabricated device yields $g_{1c}/(2\pi) = 69 \text{ MHz}$ and $g_{2c}/(2\pi) = 63 \text{ MHz}$, corresponding to a geometric average $g_{qc}/(2\pi) = 66 \text{ MHz}$, in close agreement with the design target.

To investigate the role of coherence, we repeat the simulation with T_1 values ranging from $10 \mu\text{s}$ to 1 ms . As shown in Fig. 1(c), we observe that the optimal g_{qc} decreases with increasing T_1 due to the hybridization errors of the single-qubit gates becoming the dominating error mechanism, see Appendix B for further details.

We also compare our results to selected prior works that employ similar transmon-based qubit-coupler-qubit

topology and report both T_1 and g_{qc} [17,18,20–22,48]. Most devices operate within a $\pm 5\%$ range of the simulated optimal coupling. Notably, the device in Ref. [17] stands out as an exception, employing a significantly higher g_{qc} (by a factor of 3 to 4), possibly due to focusing solely on the CZ gate fidelity. We emphasize, however, that our simulation assumptions do not exactly match those of the compared works, as the qubit and coupler frequencies, anharmonicities, and coherence times differ from our model. Hence, this comparison should be regarded as qualitative.

Our findings provide guidance for future designs: as qubit coherence times increase beyond $100\ \mu\text{s}$ [52–57], further reducing g_{qc} becomes beneficial. This not only lowers single-qubit and residual errors but also helps mitigate long-range hybridization effects such as next-nearest-neighbor crosstalk—an increasingly important factor in larger QPUs [22].

III. CZ GATE CALIBRATION

In this work, the CZ gate is implemented using simultaneous flux pulses on both the coupler and Q_1 , with virtual Z rotations applied to both qubits at the end of the gate to correct for dynamical phases. The flux pulse applied to the qubit consists of a constant plateau between cosine-shaped rising and falling edges to bring the $|11\rangle$ and $|02\rangle$ states rapidly close to resonance. The coupler flux pulse follows a Slepian-shaped envelope to minimize population leakage from Q_1 to the coupler during the gate [58]. To account for distortions in the flux lines, we precompensate the pulses using exponential filters [59].

Assuming a fixed gate duration, the calibration of this CZ gate implementation involves the minimization of three main imperfections: single-qubit Z rotations, which we offset using virtual Z gates [51]; the ZZ over-rotation angle, which is most sensitive to the frequency of Q_1 during the gate [22]; and the population leakage to $|02\rangle$, sensitive to the frequency of the coupler during the gate. As design and coherence time improvements enable higher fidelity gates, the calibration techniques for all the necessary gate parameters need to improve in both precision and speed to keep the gate fidelity limited by incoherent errors. This can be achieved through methods based on phase estimation, employing error amplification with n gates to approach the Heisenberg limit of error $\mathcal{O}(1/n)$, such as robust phase estimation [60,61], Floquet calibration [62,63], or matrix-element amplification using dynamical decoupling (MEADD) [64].

Here, we use the MEADD protocol to minimize the error in two of the three CZ gate parameters: the ZZ rotation angle and the single-qubit virtual Z rotations. Precise optimization of the third imperfection—leakage to the $|02\rangle$, caused by the required diabatic transition between the $|11\rangle$ and $|02\rangle$ states—is more challenging. Most existing error

amplification protocols which could be used to minimize the $|02\rangle$ state leakage through extending them to the two-excitation subspace are either prone to interference by other error sources, slow, or incompatible with our hardware, see discussion in Appendix C. To address this, we introduce a modified version of the MEADD protocol, called PALEA, that employs higher excited state dynamics and phase averaging to isolate the leakage parameter, eliminating the need to control or calibrate the qubit drive phase at each repetition and thereby simplifying and accelerating the implementation in our architecture.

A. Leakage error amplification

The diabatic CZ gate between transmon qubits with tunable couplers is frequently described in terms of tuning the states $|11\rangle$ and $|02\rangle$ to resonance, and then activating their coupling for a fixed time [65]. In practice, however, these two levels are not exactly on resonance, and thus this model is often insufficient. Here, we model an imperfect CZ gate as an arbitrary two-level unitary $\hat{U}_g \in \text{SU}(2)$ acting within the $\text{span}(|11\rangle, |02\rangle)$ subspace in the laboratory reference frame (see discussion in Appendix C 1). Ignoring the common phase acquired by both the $|11\rangle$ and $|02\rangle$ states, this unitary can be decomposed as $\hat{R}_z(\alpha)\hat{R}_x(\theta)\hat{R}_z(\beta)$, where $\hat{R}_j(\phi) = \exp(-i\frac{\phi}{2}\hat{\sigma}_j)$ denotes a rotation by angle ϕ around axis j , and σ_j is the corresponding Pauli operator in the $\{|11\rangle, |02\rangle\}$ basis [66]. The resulting action on the $|11\rangle$ and $|02\rangle$ states is thus given by

$$\begin{aligned}\hat{U}_g |11\rangle &= e^{i\frac{\alpha+\beta}{2}} \cos(\theta/2) |11\rangle + ie^{i\frac{\alpha-\beta}{2}} \sin(\theta/2) |02\rangle, \\ \hat{U}_g |02\rangle &= ie^{-i\frac{\alpha-\beta}{2}} \sin(\theta/2) |11\rangle - e^{-i\frac{\alpha+\beta}{2}} \cos(\theta/2) |02\rangle.\end{aligned}\quad (1)$$

The angle of interest, θ , governs the population transfer between the $|11\rangle$ and $|02\rangle$ states and is primarily controlled by the coupler frequency—and thus by the coupler flux pulse amplitude. Any deviation from $\theta = 0$ indicates an error. Hence, the aim of a leakage amplification experiment is to precisely measure the over-rotation angle θ . In contrast, the angles α and β , describing the difference of phases acquired by states $|11\rangle$ and $|02\rangle$, are irrelevant, as the $|02\rangle$ state lies outside the computational basis.

The basic leakage amplification experiment [18] repeatedly applies CZ gate candidates to an initial $|11\rangle$ state, see Fig. 2(a). As shown in Fig. 2(b), sweeping both the number of CZ gates and the coupler flux pulse amplitude while measuring the population of the $|02\rangle$ state [noted as unwanted population in Fig. 2(b)] reveals low-contrast oscillations when the coupler flux pulse amplitude significantly deviates from the optimal value of 0.447. These oscillation frequencies cannot easily be used to characterize the leakage, as they depend on the phases of the eigenvalues of $\hat{U}_g$, not just θ , see discussion in Appendix C.




FIG. 2. (a) Circuit diagram of the standard error amplification experiment. An X_{π}^{01} pulse on each qubit excites the system to the $|11\rangle$ state, followed by repeated applications of the CZ gate. (b) Measured unwanted population in the $|02\rangle$ state as a function of the number of cycles and the coupler flux pulse amplitude. (c) Circuit diagram of the phase-averaged leakage error amplification experiment. In each cycle, a CZ gate is followed by a dynamical decoupling layer D , consisting of an X_{π}^{01} gate on the lower-frequency qubit (Q_1) and an X_{π}^{12} gate on the higher frequency qubit (Q_2). (d) Measured unwanted population as a function of the number of cycles and the coupler flux pulse amplitude. Here, population in the $|02\rangle$ ($|11\rangle$) state is considered unwanted (wanted) for even n , and population in $|11\rangle$ ($|02\rangle$) state is unwanted (wanted) for odd n . (e) Vertical slices from panels (b) and (d) at the flux amplitudes indicated by dashed lines. The solid curve shows a fit to the PALEA model [Eq. (5)], yielding an over-rotation angle $|\theta_{\text{fit}}| = 0.383 \pm 0.002$. In contrast, the standard error amplification experiment signal shows low-contrast coherent oscillations at a frequency 2μ such that $\theta < 2\mu < \pi$. (f) Unwanted population, averaged over the number of cycles, as a function of the amplitude of the flux pulse. Solid lines represent Lorentzian fits used to extract the center offset x_0 , indicated by the vertical dashed line.

In addition, this experiment is susceptible to an accidental resonance with a leakage mode, for example to the coupler, appearing as vertical lines of elevated unwanted population [18]. Using Slepian-shaped coupler flux pulses reduces the leakage to the coupler and thus the presence of the associated spurious lines [18,58], though residual effects may still interfere with the calibration of the leakage to $|02\rangle$, see Appendix E.

B. PALEA protocol

To avoid the aforementioned issues in amplifying the over-rotation angle θ , we introduce the PALEA protocol.

A dynamical decoupling (DD) layer D is added after each CZ gate, consisting of an $X_{\pi}^{12}(\phi_1)$ gate with phase ϕ_1 , acting on the span($\{|1\rangle, |2\rangle\}$) subspace of the higher frequency transmon and a phased- $X_{\pi}^{01}(\phi_0)$ gate with phase ϕ_0 on the lower-frequency transmon, see Fig. 2(c). Hence, D acts on the span($\{|11\rangle, |02\rangle\}$) subspace as a π rotation,

$$D = \hat{R}_z(\phi_0 - \phi_1) \hat{R}_x(\pi) \hat{R}_z(\phi_1 - \phi_0). \quad (2)$$

This is essentially an adaptation of the SWAP error protocol from Eqs. (55)–(60) of Ref. [64], adjusting the DD layer to act on the span($\{|11\rangle, |02\rangle\}$) subspace instead of the span($\{|10\rangle, |01\rangle\}$) subspace. Importantly, to correctly implement the D operation in the laboratory frame, special care needs to be taken to ensure that the phases of the subsequent D layers within the same gate sequence are properly accounted for, see Appendix D 2.

After n cycles of CZ gates followed by the DD layers the probability of measuring $|11\rangle$ is

$$V_n(\theta, \Delta\phi) = \left| \langle 11 | \left[\hat{R}_z(\Delta\phi) \hat{R}_x(\pi - \theta) \right]^n | 11 \rangle \right|^2, \quad (3)$$

where $\Delta\phi = 2\phi_0 - 2\phi_1 - \alpha + \beta$ depends both on the difference of the phases of the applied X_{π}^{01} and X_{π}^{12} gates and the intrinsic Z rotations induced by each CZ gate between the $|11\rangle$ and $|02\rangle$ states, see discussion in Appendix D. Due to the D layer acting as a π rotation between the states $|11\rangle$ and $|02\rangle$, the cycle unitary $\hat{R}_z(\Delta\phi) \hat{R}_x(\pi - \theta)$ is now dominated by the X rotation rather than the Z rotation for small θ values, greatly increasing visibility in the final Z measurement. In contrast to the MEADD protocol, which evaluates the dynamics for only two specific values of $\Delta\phi$, PALEA averages over a wide range of random angles. This averaging is introduced either passively—through the inability to reset the phase difference between the X_{π}^{01} and X_{π}^{12} gates between repetitions, as is common for systems using CZ-based gate sets—or it can be implemented actively by manually varying the phase difference in each repetition. When sampling a broad range of $\Delta\phi$ values the expected value of the measured $|11\rangle$ state population is given by

$$I_n(\pi - \theta) = \frac{1}{2\pi} \int_0^{2\pi} d\phi V_n(\theta, \phi). \quad (4)$$

The purpose of the averaging is to decouple the unimportant angles α and β from the target over-rotation angle θ and perform only one averaged measurement.

The integral in Eq. (4) can be solved analytically (see Appendix D 3), yielding a closed-form expression for the

signal in the PALEA experiment,

$$I_n(\pi - \theta) = 1 - \cos^2(\theta/2) \sum_{m=0}^{n-1} (-1)^m P_m(\cos \theta), \quad (5)$$

where P_m is the m th Legendre polynomial. The model predicts a high-contrast signal without requiring additional calibration—adaptive or otherwise—of any other parameter, allowing for an accurate estimation of θ with fewer measurement repetitions. Equation (5) can be quickly evaluated using the Clenshaw recursion algorithm for Legendre series, see Appendix D 4.

C. Calibration using PALEA

By fitting the model to experimental data, we can accurately fit the over-rotation angle θ , and hence, the leakage to the $|02\rangle$ state. For example, the signal I_n at a coupler flux pulse amplitude of 0.436 [Fig. 2(d)] can be fitted with the model in Eq. (5) to extract $|\theta_{\text{fit}}| = 0.383 \pm 0.002$, corresponding to a population transfer to the $|02\rangle$ state of $\sin^2(\theta_{\text{fit}}/2) \approx 3.6\%$ per CZ gate, see Fig. 2(e). Conversely, the standard error amplification experiment yields a low-contrast signal for the coupler flux pulse amplitude of 0.436 [Fig. 2(b)], oscillating at a faster frequency $2\mu = 2 \cos^{-1}(\cos(\theta/2) \cos[(\alpha + \beta)/2])$ (Appendix C).

As shown in Fig. 2(f), averaging the PALEA signal over the number of cycles yields a sharp dip which can be accurately fitted with a simple Lorentzian model to determine the optimal coupler flux pulse amplitude x_0 . This enables fast and efficient calibration of a high-fidelity CZ gate. In contrast, averaging the signal of the standard error amplification experiment produces a broader dip, reducing the precision of identifying the optimal coupler flux pulse amplitude and potentially increasing the leakage to the $|02\rangle$ state—unless a significantly higher number of repetitions is used, see Appendix D 4.

The decoupling of the over-rotation angle θ from the Z rotation angles allows using PALEA as part of an optimal quantum control [67] optimization protocol varying a large number of parameters of exotic pulse shapes, from Nelder-Mead to Bayesian optimization with Gaussian processes [68,69], without the need to find additional Z rotation angles, greatly increasing its utility. Additionally, including the dynamical decoupling D to target a specific pair of states and averaging over the phases of the involved X rotations allows one to isolate the amplified exchange process between these states—here $|11\rangle$ and $|02\rangle$. The other processes which may involve the state $|11\rangle$, but are not explicitly targeted by D —such as leakage to the coupler—are effectively converted to incoherent processes and thus not amplified, see discussion in Appendix E 1. This removes a possible systematic bias introduced to the calibration by coupler leakage processes being randomly

amplified at certain coupler pulse amplitudes, see details in Appendix E 2.

Finally, although the PALEA protocol is demonstrated here in the context of the diabatic transmon-transmon CZ gates, its principles are broadly applicable. The protocol could be used to precisely measure the exchange errors in two-qubit gates across other qubit modalities, such as leakage errors in CZ gates between fluxonium qubits coupled with transmon tunable couplers [70], or SWAP errors in diabatic CZ gates between quantum dot spin qubits [71,72]. See Appendix D 6 for details.

IV. OPTIMIZING AND BENCHMARKING THE CZ GATE

To accurately characterize gate fidelity, we use iterative interleaved randomized benchmarking (IRB) [13,22,73]. In this protocol, n CZ gates are interleaved between Clifford gates instead of just one. Figure 3(a) illustrates the protocol for $n = 1, 3, 5$. The extracted total error after applying n interleaved CZ gates, $\epsilon(n)$, is shown in Fig. 3(b). We fit this dependence to a quadratic model

$$\epsilon_{\text{quad}}(n) = an^2 + bn + c$$

to disentangle coherent and incoherent error sources: incoherent errors increase linearly with n , while coherent errors and errors with long temporal correlations accumulate quadratically [73]. A further advantage of iterative IRB is the explicit extraction of the constant offset term c , which may arise from residual non-Markovian effects—such as flux or drive pulse distortions from the surrounding Clifford gates [9] or coherent flux noise [74]—unrelated to the interleaved CZ gates. In standard IRB, this offset is indistinguishable from the CZ gate error, leading to systematic over- or underestimation, whereas iterative IRB separates it from the intrinsic gate error.

We define the CZ gate error as the slope of the curve at $n = 1$, i.e., $\epsilon_{\text{CZ}} := \epsilon'_{\text{quad}}(1) = 2a + b$ [75,76]. Unlike $\epsilon_{\text{quad}}(1)$, which resembles the gate error definition of standard IRB, $\epsilon'_{\text{quad}}(1)$ is independent of c while still capturing the coherent error of a single CZ gate. Although coherent error contribution grows with n , in many applications such residual errors can be converted into incoherent errors using protocols such as randomized compiling [77], ensuring that the coherent errors do not add up constructively.

To minimize coherent CZ gate errors, we calibrate both the coupler and qubit flux pulse amplitudes using PALEA, discussed in Sec. III. In addition, we optimize the gate durations by sweeping the coupler and qubit flux pulse widths separately. Full recalibration and iterative IRB with three-state discrimination is performed for each duration setting to benchmark both the CZ gate error and the leakage to the $|2\rangle$ states (see Appendix G). We find the optimal




FIG. 3. (a) Iterative interleaved randomized benchmarking results for one, three, and five iteratively interleaved CZ gates, where the error bars indicate the standard deviation. The extracted errors per n -CZ gates, ϵ_{nCZ} , are shown with fit uncertainties. Inset: Iterative interleaved randomized benchmarking circuit, consisting of a Clifford gate C followed by n CZ gates, repeated m times, and finalized with an inverting Clifford gate C_r . (b) Extracted error vs number of iteratively interleaved CZ gates. The black line shows a quadratic fit $\epsilon_{\text{quad}}(n)$; the slope at $n = 1$ gives the CZ gate error $\epsilon_{CZ} = (4.9 \pm 0.7) \times 10^{-4}$ (dashed line). (c) Time series of ϵ_{CZ} over 40 h. Error bars indicate the standard error from the quadratic fit. The average CZ gate error, $\bar{\epsilon}_{CZ} = (6.5 \pm 0.2) \times 10^{-4}$, is shown as a dashed line. (d) Histogram of ϵ_{CZ} values from (c), with the same average indicated.

CZ gate performance at a total gate duration of 33 ns, consisting of a 22 ns coupler pulse overlapping with a 27 ns qubit pulse and 3 ns buffers on each side.

After calibration and optimization of the CZ gate, we benchmark it using iterative IRB over a continuous period of 40 h, recalibrating only the single-qubit gates between measurements [Fig. 3(c)]. For each run of the iterative IRB experiment, we use a fixed seed to generate the set of Clifford sequences and find an average CZ error of

$$\bar{\epsilon}_{CZ} = (6.5 \pm 0.2) \times 10^{-4},$$

where the uncertainty is the standard error of the mean. The lowest observed CZ gate error during the measurement series was $\epsilon_{CZ}^{\min} = (4.9 \pm 0.7) \times 10^{-4}$. The histogram

of data points in Fig. 3(d) confirms a nearly Gaussian distribution with standard deviation $\sigma(\bar{\epsilon}_{CZ}) = 1.1 \times 10^{-4}$.

The corresponding average Clifford gate error is measured to be

$$\bar{\epsilon}_{\text{Clifford}} = (2.79 \pm 0.05) \times 10^{-3},$$

in good agreement with the estimate based on the Clifford gate decomposition

$$\epsilon_{\text{Clifford}}^{\text{estimate}} = 1 - ((1 - \epsilon_{CZ})^{1.5} \cdot (1 - \epsilon_{\sqrt{X}})^{4.65}) = 2.4 \times 10^{-3},$$

using $\epsilon_{\sqrt{X}} \approx 3 \times 10^{-4}$. Note that in these experiments, the single-qubit gate duration was set to 40 ns, which had not yet been fully optimized. The improved fidelities presented later in the paper use shorter, optimized gate durations.

We also compare the iterative IRB result to the CZ error estimated using standard IRB analysis, which yields

$$\bar{\epsilon}_{CZ}^{\text{IRB}} = (7.3 \pm 0.2) \times 10^{-4}.$$

This value is significantly higher than the gate error extracted from iterative IRB. In our data, the quadratic fit reveals an average offset of $c = 1.1 \times 10^{-4}$, likely from residual non-Markovian effects unrelated to the CZ gate. Subtracting this offset from the standard IRB result yields $\bar{\epsilon}_{CZ}^{\text{IRB}} - c = 6.2 \times 10^{-4}$, in close agreement with the iterative IRB value, indicating that standard IRB mistakenly attributes the offset c to the CZ gate error.

V. ERROR BUDGET

To evaluate the overall performance of our quantum device, we construct a comprehensive error budget encompassing single-qubit gates, two-qubit (CZ) gates, and readout. While the CZ gate fidelity is a key benchmark for entangling operations, achieving consistently low errors across all layers of the control stack is critical for scalability. This section details how each component is individually optimized, calibrated, and characterized. We emphasize that the error budget reported here does not include state-preparation errors of the qubits and the coupler, which have to be treated separately from the quoted gate fidelities.

A. Single-qubit gates

Single-qubit gates are implemented using cosine-shaped envelopes combined with the derivative removal by adiabatic gate (DRAG) technique [78,79] to suppress leakage to the second excited state. The quadrature envelope is optimized for leakage minimization, and the resulting phase errors are corrected using virtual Z gates [51]. Our native gate set consists of $\sqrt{X}$ rotations.

To quantify the single-qubit gate fidelity, we use purity randomized benchmarking (RB) with three-state discrimination to capture incoherent, coherent, and leakage errors

[9,80–82]. We sweep the gate duration to identify the optimal operating point, finding a minimum error at 26 ns. This relatively long duration, compared to recent literature [9,83], is primarily constrained by compression effects in the up-conversion mixer, which limit the usable drive power.

Figures 4(a) and 4(b) show the purity RB results for Q_1 , categorizing single-qubit gate errors into incoherent, coherent, and leakage errors. We extract a single-qubit gate error of $\epsilon_{\sqrt{X}}^{Q_1} = (1.59 \pm 0.02) \times 10^{-4}$, with the dominant contribution from incoherent error, $\epsilon_{\sqrt{X}}^{Q_1, \text{incoh}} = (1.26 \pm 0.01) \times 10^{-4}$. The residual error consists of coherent calibration errors such as control pulse distortion, and a negligible amount of leakage errors to the $|2\rangle$, remaining below the measurable threshold of 1×10^{-6} . Figures 4(c) and 4(d) show analogous results for Q_2 , where we find $\epsilon_{\sqrt{X}}^{Q_2} = (1.68 \pm 0.02) \times 10^{-4}$, composed of $\epsilon_{\sqrt{X}}^{Q_2, \text{incoh}} = (1.62 \pm 0.01) \times 10^{-4}$, $\epsilon_{\sqrt{X}}^{Q_2, \text{coh}} = (4 \pm 2) \times 10^{-6}$, and leakage $\epsilon_{\sqrt{X}}^{Q_2, \text{leak}} = (3.7 \pm 0.5) \times 10^{-6}$.

The low leakage errors are expected given the relatively long gate durations used and the additional suppression from the DRAG pulse. Simulations for a $\sqrt{X}$ gate with a duration of 26 ns, qubit-qubit detuning of 110 MHz, and qubit relaxation times of $T_1 \approx 100 \mu\text{s}$ predict a gate error of $\epsilon_{\sqrt{X}}^{\text{sim}} = 1.3 \times 10^{-4}$ for both qubits, fully limited by incoherent error, agreeing well with the measured values. At these gate durations, detuning, and optimized qubit-coupler couplings, the expected hybridization error is below 1×10^{-5} , and therefore does not limit the performance. This is experimentally confirmed by comparing gate fidelities of the two qubits individually, where we observe nearly identical performance indicating negligible hybridization error. These results also imply that the microwave crosstalk, measured to be approximately -20 dB, is not a limiting factor in our setup at the given qubit-qubit detuning.

We note that the simulated $\sqrt{X}$ gate errors in Fig. 1(b) are based on optimally short gate durations (on the order of 15 ns), while our experiments use longer pulses due to limitations mentioned above. Hence, the simulated gate error




FIG. 4. (a) and (b) Purity randomized benchmarking results for Q_1 . Shown are the average $\langle \sigma_z \rangle$, the sequence purity $\langle \sqrt{P} \rangle$, and the leakage population p_f as a function of the number of Clifford gates. The extracted $\sqrt{X}$ gate error is $\epsilon_{\sqrt{X}}^{Q_1} = (1.59 \pm 0.02) \times 10^{-4}$. (c) and (d) Same as (a) and (b), but for Q_2 yielding $\epsilon_{\sqrt{X}}^{Q_2} = (1.68 \pm 0.02) \times 10^{-4}$. (e) Error budget for the CZ gate extracted from leakage interleaved randomized benchmarking. Errors for CZ (purple) and identity gates (blue) are shown as a function of the number of gate repetitions. CZ data are fitted with a quadratic model, identity with a linear model. The blue shaded region indicates the CZ error limit $\epsilon_{\text{CZ}}^{\text{lim}}$ in the absence of coupler-induced effects. Contributions from incoherent errors ($\epsilon_{\text{CZ}}^{\text{incoh}}$), leakage to the second excited states ($\epsilon_{\text{CZ}}^{\text{leak}}$), and coherent errors ($\epsilon_{\text{CZ}}^{\text{coh}}$) are shown in shades of purple. (f) and (g) Single-shot readout clouds for Q_1 and Q_2 , based on two million repetitions. Using a shelving technique via an X_{π}^{12} pulse, $|e\rangle$ state populations are moved mostly to $|f\rangle$ state prior to readout. Black lines indicate thresholds between $|0\rangle$ and $|\bar{0}\rangle$. (h) and (i) Probability density vs in-phase readout signals used to extract readout fidelities: $\epsilon_{\text{RO}}^{Q_1} = (5.7 \pm 0.2) \times 10^{-4}$ and $\epsilon_{\text{RO}}^{Q_2} = (5.4 \pm 0.2) \times 10^{-4}$. (j) Summary of combined error contributions for each qubit, separated into single-qubit gate errors (incoherent, coherent, leakage), two-qubit gate errors (incoherent, leakage), and readout errors (decay, overlap, mixing).

reaches 1×10^{-4} , even when averaging over nonideal qubit-qubit detunings.

B. CZ gate error decomposition

To assess the two-qubit gate performance, we use the iterative IRB method described in Sec. IV. In addition to measuring the CZ gate error directly, we seek to disentangle its error contributions.

We first determine the coherence-limited error floor of the qubits by applying iterative IRB with an interleaved identity (ID) gate of the same duration as the CZ gate (33 ns). The extracted error per ID gate is $\epsilon_{\text{ID}} = (4.1 \pm 0.2) \times 10^{-4}$ [Fig. 4(e)], which serves as an estimate for the lower bound of CZ gate error in the absence of coupler-induced decoherence, $\epsilon_{\text{CZ}}^{\text{lim}}$.

Next, we estimate the incoherent component of the CZ gate error, $\epsilon_{\text{CZ}}^{\text{incoh}}$, by measuring T_1 and T_2 of the hybridized qubit modes as a function of the coupler flux bias, and integrating over the flux pulse waveform to compute effective coherence times (see Appendix H). This yields $\epsilon_{\text{CZ}}^{\text{incoh}} = 6.2 \times 10^{-4}$. The increased value compared to ϵ_{ID} confirms the impact of the coupler's shorter relaxation time ($T_1^{\text{TC}} = 33 \mu\text{s}$) versus the qubit lifetimes ($T_1^Q \approx 100 \mu\text{s}$).

To characterize leakage errors to higher qubit excited states, we implement leakage randomized benchmarking [82,84] using three-state discrimination of the iterative IRB data. We extract a total leakage error to the second excited states of both qubits, $L_{\text{CZ}}^{\text{tot}} = 1 \times 10^{-4}$. Most leakage originates from Q_2 ($L_{\text{CZ}}^{Q_2} = 9 \times 10^{-5}$), which transiently populates the $|2\rangle$ state during the CZ gate. We attribute the leakage remaining after the PALEA optimization to dephasing in the $\text{span}(|11\rangle, |02\rangle)$ subspace during the gate.

Summing the incoherent and leakage contributions gives $\epsilon_{\text{CZ}}^{\text{incoh}} + L_{\text{CZ}}^{\text{tot}} = 7.2 \times 10^{-4}$, which closely matches the measured average CZ error of $\bar{\epsilon}_{\text{CZ}} = 6.5 \times 10^{-4}$, suggesting that most error contributions are captured by this decomposition. The residual error is attributed to coherent error $\epsilon_{\text{CZ}}^{\text{coh}}$, which manifests as a quadratic contribution in iterative IRB and is likely caused by flux pulse distortions that induce inter-gate memory effects.

We also characterize the SWAP and coupler leakage errors and demonstrate that neither limits the performance of our CZ gate. The SWAP error arises from coherent population exchange between the $|10\rangle$ and $|01\rangle$ states and is described by a two-level unitary operator $\hat{U}_s$ acting in the $\text{span}(|10\rangle, |01\rangle)$ subspace. Analogous to the leakage analysis in Sec. III, $\hat{U}_s$ is decomposed as $\hat{R}_z(\alpha_s)\hat{R}_x(\theta_s)\hat{R}_z(\beta_s)$. Here, θ_s quantifies the SWAP error, while α_s and β_s correspond to conditional and single-qubit phases that can be calibrated separately. Using PALEA, we measure $\theta_s = (0.0198 \pm 0.0006) \text{ rad}$, which corresponds to a SWAP error of $\epsilon_{\text{CZ}}^{\text{SWAP}} = (3.9 \pm 0.2) \times$

10^{-5} , confirming that SWAP errors are not a limiting factor for the CZ gate fidelity (see Appendix F1 for further discussion).

Leakage to the coupler is more challenging to quantify, as the coupler lacks direct XY control and readout. The absence of a dedicated drive and measurement line precludes the use of PALEA or randomized benchmarking techniques to assess population transfer to the coupler excited state. Instead, we adopt a modified Floquet calibration method [62], in which a variable delay τ is inserted between repeated CZ gates. During this delay, the state $|101\rangle$ and the leaked state $|011\rangle$ acquire a relative phase under the free Hamiltonian evolution of the system. For specific τ values, this phase evolution constructively amplifies any residual coupler population, enabling its detection. We measure a leakage-induced rotation angle of $\theta_{\text{TC}} = (0.041 \pm 0.001) \text{ rad}$ per CZ gate, corresponding to a coupler leakage error of $\epsilon_{\text{CZ}}^{\text{TC}} = \frac{4}{5} - \frac{4}{5} \cos^4(\theta_{\text{TC}}/4) = (1.7 \pm 0.1) \times 10^{-4}$, where we have modeled the coupler leakage as an incoherent decay of Q_1 . This decay is already partially accounted for by the difference between ϵ_{ID} and ϵ_{CZ} (see Appendixes E1 and F2). We attribute the small coupler leakage primarily to the use of a Slepian-shaped envelope for the coupler flux pulse, which effectively minimizes leakage during the CZ gate.

C. Shelved readout and fidelity analysis

For high-fidelity measurement, we use a shelving-based readout technique: an X_π^{12} gate is applied just prior to a rectangular readout pulse to transfer the $|e\rangle$ state to $|f\rangle$, thereby slightly increasing its effective readout lifetime [26,35–37]. We begin by optimizing the readout frequency, followed by a joint optimization of the readout amplitude and duration for maximum readout fidelity. Finally, after optimizing the integration weights, the configuration consists of a 40 ns shelving gate followed by a 240 ns readout pulse, resulting in a total duration of 280 ns. In practice, subsequent operations are constrained by the time required to allow resonator photons to decay as well as control-system timing constraints. Due to a limitation in the electronics firmware, the readout pulse is followed by a fixed idle time of 347 ns before the next readout can be triggered. As a result, the time between the start times of two consecutive readout operations is 627 ns. This interval is sufficient for complete photon depletion under our operating conditions, although the intrinsic resonator ring-down time may be shorter than this firmware-imposed delay.

Figures 4(f) and 4(g) show single-shot in-phase and quadrature distributions for Q_1 and Q_2 based on two million samples. $|g\rangle$ state preparations are mostly classified as $|0\rangle$, with rare $|\bar{0}\rangle$ misclassifications mostly due to measurement-induced excitations. $|e\rangle$ state preparations—shelved to $|f\rangle$ with occasional $f \rightarrow e$ relaxation—are classified as $|\bar{0}\rangle$; rare $e \rightarrow g$ relaxation yields

TABLE I. Readout performance metrics for shelved and quantum non-demolition-optimized readout.

	Q ₁		Q ₂	
	Shelved	QND-opt.	Shelved	QND-opt.
Readout error, ϵ_{RO}	$(5.7 \pm 0.2) \times 10^{-4}$	$(1.92 \pm 0.03) \times 10^{-3}$	$(5.4 \pm 0.2) \times 10^{-4}$	$(1.59 \pm 0.03) \times 10^{-3}$
QND error, $1 - \mathcal{Q}$	$(1.06 \pm 0.02) \times 10^{-2}$	$(6.7 \pm 0.2) \times 10^{-3}$	$(1.11 \pm 0.02) \times 10^{-2}$	$(5.9 \pm 0.2) \times 10^{-3}$
Readout-induced leakage, L	$(1.005 \pm 0.006) \times 10^{-2}$	$(2.1 \pm 0.2) \times 10^{-5}$	$(1.05 \pm 0.005) \times 10^{-2}$	$(6.6 \pm 0.3) \times 10^{-4}$

$|0\rangle$ misclassifications. Counting the assignment errors, we extract readout errors of $\epsilon_{\text{shelved}}^{\text{Q1}} = (5.7 \pm 0.2) \times 10^{-4}$ and $\epsilon_{\text{shelved}}^{\text{Q2}} = (5.4 \pm 0.2) \times 10^{-4}$, with uncertainties obtained via bootstrapping [85].

From the histogram fits shown in Figs. 4(h) and 4(i), we identify the main readout error sources for Q₁ as decay error $\epsilon_{\text{shelved}}^{\text{Q1,decay}} = (3.8 \pm 0.2) \times 10^{-4}$, mixing error $\epsilon_{\text{shelved}}^{\text{Q1,mix}} = (1.48 \pm 0.01) \times 10^{-4}$, and negligible overlap error $\epsilon_{\text{shelved}}^{\text{Q1,overlap}} = (3.6 \pm 0.1) \times 10^{-5}$. For Q₂, the error profile follows the same pattern, with decay error $\epsilon_{\text{shelved}}^{\text{Q2,decay}} = (3.6 \pm 0.2) \times 10^{-4}$ as the leading contribution, while mixing and overlap errors remain at levels comparable to those of Q₁.

Although shelving increases the effective relaxation time during measurement, qubit decay remains the dominant limitation to the readout fidelity. Potential improvements include shelving to the third excited state, which further extends the effective lifetime of the readout manifold.

While shelving provides high-fidelity readout suitable for final measurements, it is less practical for mid-circuit measurements required in quantum error correction [3] due to the potentially increased risk of leakage to the $|f\rangle$ state. Compatibility with such measurements can be restored if high-fidelity reset from the $|f\rangle$ state is available [86–90], enabling the qubit to be returned to its premeasurement state when needed.

Ideally, mid-circuit measurements are implemented as quantum non-demolition (QND) readouts, which preserve the qubit state regardless of the measurement outcome [44]. In practice, no readout is perfectly QND, and the goal is to maximize its QNDness $\mathcal{Q}$, defined as [45]

$$\mathcal{Q} = (\mathcal{Q}_g + \mathcal{Q}_e) / 2, \quad (6)$$

with

$$\begin{aligned} \mathcal{Q}_g &= P(g, 0|g) + P(g, 1|g), \\ \mathcal{Q}_e &= P(e, 0|e) + P(e, 1|e), \end{aligned} \quad (7)$$

where $P(a, b | c)$ denotes the probability that the qubit ends up in state a and is assigned outcome b , given the qubit is initially in state c . We adapt the readout-induced leakage benchmarking (RILB) protocol [45] to extract $\mathcal{Q}$; see Appendix I for details on the experiment and models.

Using the fidelity-optimized shelved readout parameters and appending an additional X_{π}^{12} gate after the readout pulse to return population from $|f\rangle$ to $|e\rangle$, we observe that the QNDness is limited by elevated leakage rates (Table I). This leakage is attributed to the decay from $|f\rangle$ to $|e\rangle$ during the readout pulse, followed by re-excitation to $|f\rangle$ via the second X_{π}^{12} gate.

As an alternative, we optimize a standard (unshelved) readout pulse for QNDness, which typically involves lowering the readout amplitude to reduce the number of photons in the resonator and minimize excitation to higher states [45,91,92]. The QND-optimized readout reduces leakage by up to two orders of magnitude relative to the shelved readout and improves the QNDness up to $\mathcal{Q} = 99.41\%$, albeit with reduced fidelity; see Table I. The readout implementation optimized for QNDness may be readily utilized in quantum error correction experiments with a no-reset scheme [93].

These results place our QND performance alongside other high-fidelity QND measurements reported for QPU-compatible architectures [29,87,94–97].

VI. CONCLUSIONS

We have demonstrated simultaneous high-fidelity single-qubit gates, CZ gates, and readout in a superconducting quantum device composed of two transmon qubits coupled via a tunable floating coupler. Through targeted device-level simulations, we identified the optimal qubit-coupler coupling strength that minimizes the total Clifford error by balancing hybridization-induced crosstalk and incoherent gate errors. This design strategy was validated experimentally, yielding single-qubit gate fidelities above 99.98%, a 40-hour-averaged CZ gate fidelity of 99.93%, and readout fidelities exceeding 99.94%.

A key enabler of this performance is our newly developed PALEA protocol, which allows precise calibration of coherent over- and under-rotation errors in the $\text{span}(|11\rangle, |02\rangle)$ subspace. Combined with iterative interleaved randomized benchmarking, PALEA enables both the suppression and quantification of coherent errors, leading to a factor of two reduction in leakage and an accurate estimation of the residual CZ gate error.

We also presented a detailed error budget that quantifies the incoherent, coherent, and leakage components of both single- and two-qubit gates, as well as readout. The readout

fidelity was achieved via shelving to the second excited state and was broken down into decay, mixing, and overlap contributions. In addition, we optimized nonshelved readout for QNDness, obtaining values exceeding 99.3% for both qubits.

Importantly, the device architecture and control methodology studied here are compatible with scalable layouts. Our design allows each qubit to be coupled to multiple tunable couplers while preserving gate parameters, and the chip layout supports dedicated Purcell filters for each qubit on a single fabrication layer [43,47]. These features make our approach well-suited for extension to 2D square-lattice architectures required, for example, for surface code implementations.

Together, our results provide a blueprint for achieving uniformly high gate and readout fidelities in a superconducting platform designed with scalability in mind. While fast reset and leakage-removal protocols remain to be integrated, these advances represent an important step toward fault-tolerant quantum computing.

ACKNOWLEDGMENTS

We acknowledge Ali Yurtalan, Lucas Ortega, and Pavel Titov for supporting the construction and maintenance of the experimental setup, and Balint Csatari, Rakhim Davletkaliyev, Guillermo Alonso, Jyrgen Luus, and Umut Deniz Özugurel for additional software support. We would additionally like to thank the rest of the IQM team for creating the entire infrastructure, laying the foundation of this work. We would like to thank Sumeru Hazra and Wei Dai for discussions on the RILB protocol.

A.V.E. conceptualized the project, F.M. and J.M. planned and executed the experiments and analyzed the experimental data, A.G., A.L., E.T., T.M., J.R., conducted and analyzed the microwave simulations, J.A.N., S.L., T.O., M.ST., H.H., O.S., and J.T. provided theoretical modeling and simulations, A.G., C.O.-K., A.L., and S.R.D. designed and simulated the sample, F.M., J.M., J.A.D., M.ST., M.SAV., R.B., C.F.C., O.F., Z.G., A.G.F., D.G., T.H., E.H., J.I., S.I., S.J., M.K., J.L., N.L., P.L., J.R., B.T., J.W., A.V.A., V.B., and A.V.E. developed the experiment and analysis software, A.P., S.-G.K., W.L., W.Q., A.K.O., L.Z., H.-H.L., and A.K.A. fabricated, coordinated, and selected the device, K.M., V.S., and S.D. bonded and packaged the device, R.K., M.SAR., and M.SAV. maintained and improved the experimental setup, L.A., S.K., A.R., S.S., and M.ST. worked on coherence improvements for this sample, W.K., H.-S.K., F.D., J.H.A., M.R., and J.V. provided background support, F.M., J.M., J.A.N., J.A.D., J.H.E., and A.V.E. drafted or revised the manuscript, and J.H.E., A.G., and A.V.E. supervised the work.

This work was supported by Business Finland through Projects QuTi (No. 40755/31/2020) and CfoQ (No.

787/31/2025). Parts of this work are included in patent applications filed by IQM Finland Oy.

DATA AVAILABILITY

The data that support the findings of this article are openly available [98].

APPENDIX A: EXPERIMENTAL SETUP

The experiments are performed in a cryogenic dilution refrigerator with a base temperature of approximately 10 mK. A schematic of the full wiring configuration is shown in Fig. 5. The sample consists of two frequency-tunable floating transmon qubits coupled via a frequency-tunable floating transmon coupler. A summary of the key device parameters is provided in Table II.




FIG. 5. Schematic of the experimental setup. The background colors depict the different temperature stages of the dilution refrigerator. The coaxial microwave cables (straight black lines) are used to control and read out the qubit and coupler states. The three-wave-mixing traveling-wave parametric amplifier is biased via a twisted-pair cable and pumped through a coaxial microwave line. The schematic of the sample in the white box consists of two flux-tunable qubits (blue and orange), a readout structure for each qubit (lime green), and a flux-tunable coupler (red).

TABLE II. Summary of device parameters.

Quantity, symbol (unit)	Qubit 1	Qubit 2	Coupler
Readout resonator frequency, $\omega_R/2\pi$ (GHz)	6.190	6.350	—
Readout Purcell filter frequency, $\omega_{\text{Rpf}}/2\pi$ (GHz)	6.187	6.329	—
Effective readout resonator bandwidth, $\kappa_{\text{eff}}/2\pi$ (MHz)	6.1	3.4	—
Readout circuit dispersive shift, $\chi_R/2\pi$ (MHz)	2.5	2.6	—
Qubit frequency at sweet spot, $\omega_{\text{ge}}/2\pi$ (GHz)	4.799	4.910	4.850
Qubit frequency at idling, $\omega_{\text{ge}}/2\pi$ (GHz)	4.799	4.910	3.672
Transmon anharmonicity, α (MHz)	−191	−191	−216
Energy relaxation time at idling, T_1 (μs)	86	102	33
Transverse relaxation time at idling, T_2 (μs)	74	63	1.4
Echoed transverse relaxation time at idling, T_2^{cho} (μs)	140	104	3.1
Pure dephasing time at idling, T_ϕ (μs)	130	91	1.4
Qubit-coupler coupling, $g_{\text{qc}}/2\pi$ (MHz)	69.6	62.2	—

Each qubit has an individual microwave drive line, equipped with 39 dB of attenuation distributed across the various temperature stages to minimize thermal noise. Additional low-pass and Eccosorb filters are installed at the mixing chamber stage to suppress spurious high-frequency signals. The Eccosorb filters add approximately another 15 dB attenuation at the qubit frequencies.

The qubits are driven using a Zurich Instruments HDAWG arbitrary waveform generator (AWG), which produces intermediate-frequency pulses that are up-converted to the qubit frequencies using IQ mixers. Dedicated HDAWG channels are also used to control the voltage on the flux line of Q_1 and the coupler. For Q_2 , which requires only static flux control, we use a Qblox SPI D5a voltage source module to deliver DC biases.

The readout bus consists of a transmission line capacitively coupled to the readout structure of each qubit. Each readout structure comprises a dedicated resonator-Purcell filter pair, designed to suppress spontaneous emission into the environment while maintaining the coupling strength required for fast, high-fidelity measurement. Readout pulses are generated using a Zurich Instruments UHFQA, with frequency multiplexing applied to address both qubits simultaneously. The UHFQA output is up-converted and sent through the dilution refrigerator input chain.

The readout signal is first amplified by Arctic Instruments' three-wave-mixing TWPA at the 10 mK stage and subsequently by a 4 K high-electron-mobility transistor amplifier. The amplified signal is then downconverted at room temperature and digitized by the UHFQA. The TWPA is pumped by a Rhode & Schwarz microwave generator at 11 GHz and biased using two DC output channels of a Qblox SPI D5a voltage source. In the same measurement setup, we independently estimated the total readout-chain quantum efficiency to be approximately $\eta \approx 0.13$, measured on a similar chip in the same cryostat and wiring configuration [99].

For reference, Fig. 5 includes the location and labeling of all attenuators, filters, amplifiers, and circulators, as well as the thermal stage at which each component is mounted. All flux and drive lines are thermally anchored at each temperature stage to minimize heat load and thermal noise coupling to the qubits.

APPENDIX B: QUBIT-COUPLER COUPLING SIMULATION

1. ZZ interaction and transverse coupling

For a system comprising two transmon qubits coupled via a coupler, the effective two-qubit Hamiltonian can be written as

$$\hat{H}/\hbar = \sum_{i=1,2} \frac{1}{2} \tilde{\omega}_i \hat{\sigma}_i^z + g_{XY} (\hat{\sigma}_1^+ \hat{\sigma}_2^- + \hat{\sigma}_2^- \hat{\sigma}_1^+), \quad (\text{B1})$$

where $\tilde{\omega}_i$ is the angular frequency of the eigenmode corresponding to Q_i , $\hat{\sigma}_i^z$ is the Pauli Z operator, $\hat{\sigma}_1^+$ and $\hat{\sigma}_2^-$ are the raising and lowering operators, and g_{XY} is the effective transverse coupling strength between the qubits. In our device, the floating coupler frequency ω_c is below the bare qubit frequencies ω_i for $i \in \{1, 2\}$, resulting in a transverse coupling strength that is approximately given by [22,23]

$$g_{XY} \approx g_{12} - \frac{g_{1c}g_{2c}}{2} \left(\frac{1}{\Delta_{1c}} + \frac{1}{\Delta_{2c}} - \frac{1}{\Sigma_{1c}} - \frac{1}{\Sigma_{2c}} \right), \quad (\text{B2})$$

where $\Delta_{ic} = \omega_i - \omega_c$ and $\Sigma_{ic} = \omega_i + \omega_c$.

The ZZ interaction is defined as

$$\zeta = (\omega_{|11\rangle} - \omega_{|01\rangle}) - (\omega_{|10\rangle} - \omega_{|00\rangle}), \quad (\text{B3})$$

where $\omega_{|ij\rangle}$ represents the angular frequency of the corresponding eigenstate. An analytical expression for ζ can be obtained using a Schrieffer-Wolff transformation that retains terms up to at least fourth order in the coupling

strength [18,100]. This analysis reveals that, for most qubit-qubit detunings, the points where ζ is canceled and where the transverse coupling g_{XY} vanishes occur at different coupler frequencies.

2. Assumptions in the simulation

To simulate single- and two-qubit gate errors, we use a custom simulation framework built on the open-source Python library QuTiP [101]. As described in the main text, our simulations rely on several simplifying assumptions, which we elaborate on here.

First, both qubits are assumed to have identical relaxation and dephasing times: $T_1^{Q1} = T_1^{Q2} = T_\phi^{Q1} = T_\phi^{Q2}$. While this is not exact, it closely reflects the behavior observed in our device (see Fig. 14). Second, the coupler is assigned a relaxation and dephasing time that is four times shorter than that of the qubits. The reduced relaxation time of the coupler is motivated by its larger EPR in lossy interfaces—which is approximately four times higher compared to the EPR of the qubits (0.8×10^{-3} vs 0.2×10^{-3})—and by coherence measurements of our fabricated device (see Table II). Third, all elements—the two qubits and the coupler—are modeled as three-level Duffing oscillators. Adding a fourth level did not significantly alter the simulation results. The anharmonicities of the qubits and coupler are set to values matching those of the fabricated device (see Table II). Finally, to isolate the impact of the qubit-coupler coupling strength g_{qc} , we vary the direct qubit-qubit coupling g_{12} in such a way that the coupler idling frequency remains constant across the sweep. This leads to g_{12} values in the range [0.7, 17.1] MHz.

3. Single-qubit gate error mechanisms

Single-qubit gate errors are primarily influenced by incoherent processes, hybridization-induced crosstalk, and, for short pulses, leakage to higher energy levels. Hybridization crosstalk arises from the fact that the qubits are slightly hybridized with each other due to the choice of the idling configuration. Although the coupler is biased to cancel the ZZ interaction ($\zeta = 0$), the eigenstates of the coupled system— $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$ —are not perfectly aligned with the bare states $|gg\rangle$, $|ge\rangle$, $|eg\rangle$, $|ee\rangle$ [18]. As a result, driving one qubit can unintentionally excite the hybridized mode of the other, even in the absence of classical crosstalk.

This hybridization crosstalk depends on (i) the qubit-coupler coupling strength g_{qc} (larger g_{qc} leads to stronger hybridization), (ii) the qubit-qubit detuning $\Delta\omega$, and (iii) the gate duration (shorter pulses have broader frequency content and are more likely to excite neighboring eigenstates). To capture these effects, we perform a four-dimensional sweep over $\sqrt{X}$ gate duration, g_{qc} , relaxation time T_1 (which sets both T_1 and T_ϕ), and qubit-qubit

detuning $\Delta\omega$, simulating both qubits simultaneously. The drive waveform is a cosine-shaped envelope, matching our experimental implementation.

Figures 6(a) and 6(d) show the simulated $\sqrt{X}$ gate error versus gate duration for various g_{qc} values, for Q_1 and Q_2 , respectively. For small g_{qc} , the gate error is coherence-limited for durations longer than approximately 16 ns. For shorter gates, leakage to the $|2\rangle$ state becomes the dominant source of error. While such leakage could in principle be mitigated using advanced pulse shaping techniques such as FAST DRAG or higher-order DRAG [9], in practice we observe no benefit due to mixer compression at high drive powers.

For larger g_{qc} , hybridization crosstalk becomes significant. In this regime, the gate error decreases with increasing duration, as longer gates have narrower frequency spectra and interact more adiabatically with unwanted transitions. Oscillations observed in the Q_2 traces originate from transient population of the $|2\rangle$ state due to the proximity of the $|11\rangle$ and $|02\rangle$ levels. These oscillations are absent in Q_1 , which is the lower-frequency qubit and does not couple strongly to $|20\rangle$.

From each duration sweep, we extract the gate duration that minimizes the error and plot the corresponding minimum $\sqrt{X}$ gate error as a function of g_{qc} for several T_1 values [Figs. 6(b) and 6(e)]. The overall trend shows increasing gate error with increasing g_{qc} . At high g_{qc} , the system is hybridization-limited, whereas at low g_{qc} , it becomes coherence-limited—yielding two distinct slopes.

To evaluate how the $\sqrt{X}$ gate error depends on qubit-qubit detuning, we simulate five values of $\Delta\omega/(2\pi) \in \{60, 85, 110, 135, 160\}$ MHz, shown in Figs. 6(c) and 6(f). These detunings are symmetrically spaced around $\Delta\omega/(2\pi) = 110$ MHz, where the $\sqrt{X}$ gate error reaches its minimum. This configuration corresponds to both the $|01\rangle$ and $|10\rangle$, as well as the $|11\rangle$ and $|02\rangle$ states being far detuned from each other.

The simulated detuning range of 100 MHz is chosen to reflect realistic constraints in large-scale processors, where nearest-neighbor and next-nearest-neighbor detunings must be engineered to avoid spectral crowding.

The trend of increasing gate error with g_{qc} is consistent across all $\Delta\omega$ values. As expected, hybridization error is minimized near $\Delta\omega/(2\pi) = 110$ MHz and increases for both lower and higher detunings. Interestingly, we observe that the gate errors at smaller detunings (60 and 85 MHz) are higher than their larger-detuning counterparts (135 and 160 MHz), despite their symmetric placement around the optimal point. This asymmetry arises because the effective transverse coupling g_{XY} —which governs the strength of hybridization between the qubits—is stronger at smaller detunings and weaker at larger ones, leading to enhanced crosstalk for low $\Delta\omega$ in our design.




FIG. 6. (a) and (d) Simulated $\sqrt{X}$ gate error vs gate duration for Q_1 and Q_2 , respectively, at $T_1 = 100 \mu\text{s}$, $\Delta\omega = 150 \text{ MHz}$, and qubit-coupler couplings from 30 to 150 MHz. (b) and (e) Simulated $\sqrt{X}$ gate error vs qubit-coupler coupling for a range of T_1 times ($10 \mu\text{s}$ – 1 ms) at fixed $\Delta\omega/(2\pi) = 150 \text{ MHz}$, and an optimized gate duration. (c) and (f) Simulated $\sqrt{X}$ gate error vs qubit-coupler coupling at $T_1 = 100 \mu\text{s}$ for five different qubit-qubit detunings $\Delta\omega/(2\pi) \in \{60, 85, 110, 135, 160\} \text{ MHz}$. (g) Simulated CZ gate error vs gate duration for different qubit-coupler couplings at $T_1 = 100 \mu\text{s}$. (h) Simulated CZ gate error vs qubit-coupler coupling for various T_1 values. (i) Simulated CZ gate duration vs qubit-coupler coupling for the same range of T_1 times.

4. Two-qubit gate error mechanisms

To estimate the CZ gate errors, ϵ_{CZ} , we perform a three-dimensional sweep over gate duration (i.e., both qubit and coupler flux pulse widths), qubit-coupler coupling g_{qc} , and T_1 . The qubit flux pulse includes 1 ns cosine-shaped rising and falling edges. Slices of the results are shown in Figs. 6(g)–6(i).

Figure 6(g) shows ϵ_{CZ} versus gate duration for various g_{qc} at fixed $T_1 = 100 \mu\text{s}$. At short durations and low g_{qc} , coherent under-rotation in the $\text{span}(|11\rangle, |02\rangle)$ subspace dominates. At longer durations or higher couplings, incoherent error becomes the limiting factor. For $g_{\text{qc}} \gtrsim 60 \text{ MHz}$, oscillatory behavior appears due to SWAP-type coherent errors, see definition in Appendix F 1. These can be minimized through careful gate timing, although leakage to $|20\rangle$ persists and ultimately limits performance.

From each duration sweep, we extract the minimum ϵ_{CZ} per g_{qc} and plot it versus coupling strength for several T_1 values in Fig. 6(h). For low g_{qc} , gate errors drop rapidly due to decreasing gate duration and reduced incoherent error. At higher couplings, the trend saturates as leakage to $|20\rangle$ becomes the dominant contribution. The apparent noise in Fig. 6(h) is not numerical; rather, it reflects the sensitivity of leakage to subtle changes in the CZ gate duration. This is also evident in Fig. 6(i), where optimal

CZ durations exhibit irregular switching between local minima.

5. Connection to main text

The single-qubit error data shown in Fig. 1(b) are generated by fixing $T_1 = 100 \mu\text{s}$, determining the optimal $\sqrt{X}$ gate duration per g_{qc} , averaging over the five $\Delta\omega$ values, and summing errors from both qubits. The CZ gate error data corresponds directly to Fig. 6(h) at $T_1 = 100 \mu\text{s}$.

To compute the Clifford error $\epsilon_{\text{Clifford}}$, we combine the two contributions using a weighted sum based on Clifford decomposition: 4.65 $\sqrt{X}$ gates and 1.5 CZ gates per Clifford, following the prescription in Ref. [2], with additionally replacing X and Y rotations by $\pi/2$ pulses and virtual-Z gates [51].

For Fig. 1(c), we apply a bivariate spline interpolation to the $\epsilon_{\text{Clifford}}(g_{\text{qc}}, T_1)$ data, extract the minimum along the g_{qc} axis for each T_1 , and plot the result as the solid black line indicating the optimal coupling strength.

APPENDIX C: DYNAMICS OF THE REPEATED DIABATIC CZ GATE

In this section, we elaborate on the theory behind the standard error amplification experiment [Figs. 2(a)

and 2(b)], as well as its extensions for characterizing leakage to the $|02\rangle$ state and the coupler.

1. Choice of the reference frame

In this Appendix and Appendixes D and E, we primarily work in the laboratory reference frame, unless specified otherwise. This choice is motivated by the nature of the diabatic CZ gate, which is implemented by dynamically tuning the frequencies of the coupler and qubit Q_1 [22]. As a result, in the laboratory frame, each of the repeated CZ gates in the same gate sequence can be described by the same unitary $\hat{U}_g$ as in Eq. (1) of the main text. Moving to a standard computational frame, defined by the frequencies of the two qubits, would cause the underlying Hamiltonian $\hat{H}$ to acquire additional time dependence as a function of the difference of the qubit frequencies. Consequently, each consecutive CZ in the same sequence would be described by a different unitary in the $\text{span}(|11\rangle, |02\rangle)$ subspace. Thus, choosing to instead work in the laboratory frame (or any interaction frame associated with the same frequency ω for both qubits) greatly simplifies the derivations.

This choice has minor implementation consequences which will be described later in Appendix D 2.

2. Standard error amplification

The goal of this section is to analyze the dynamics of the standard error amplification experiment shown in Figs. 2(a) and 2(b) of the main text. In particular, we seek to understand why the signal-to-noise ratio (SNR) of the population oscillations becomes small near the optimal coupler flux pulse amplitude—rendering this method inefficient for minimizing leakage.

To this end, we calculate the expectation value V_n after applying n repeated CZ gates to the initial state $|11\rangle$. We start from the unitary $\hat{U}_g$ defined in Eq. (1) in the main text, acting on the $\text{span}(|11\rangle, |02\rangle)$ subspace and decomposed as $\hat{U}_g = \hat{R}_z(\alpha)\hat{R}_x(\theta)\hat{R}_z(\beta)$, where θ governs the population transfer between $|11\rangle$ and $|02\rangle$. The expectation value for the population remaining in $|11\rangle$ after n applications of $\hat{U}_g$ is then

$$\begin{aligned} V_n(\theta, \alpha, \beta) &= \left| \langle 11 | \left[\hat{U}_g(\theta, \alpha, \beta) \right]^n | 11 \rangle \right|^2 \\ &= \left| \langle 11 | \left[\hat{R}_z(\alpha + \beta)\hat{R}_y(\theta) \right]^n | 11 \rangle \right|^2. \end{aligned} \quad (\text{C1})$$

In the second line, we use the identity

$$\begin{aligned} & \left(\hat{R}_z(\alpha)\hat{R}_x(\theta)\hat{R}_z(\beta) \right)^n \\ &= \hat{R}_z\left(\alpha - \frac{\pi}{2}\right) \left(\hat{R}_z(\alpha + \beta)\hat{R}_y(\theta) \right)^n \hat{R}_z\left(\beta + \frac{\pi}{2}\right), \end{aligned} \quad (\text{C2})$$

and note that the leading and trailing Z rotations do not affect the expectation value in a Z basis measurement. We

define $\phi = \alpha + \beta$ as the effective Z rotation applied in each cycle. Since α and β are not relevant for the operation of a CZ gate, ϕ remains uncalibrated and effectively random, but constant throughout the experiment.

The single cycle operator $\hat{R}(\phi, \theta) = \hat{R}_z(\phi)\hat{R}_y(\theta)$ can be diagonalized as

$$\hat{R}(\phi, \theta) = S(\phi, \theta)D(\phi, \theta)S^{-1}(\phi, \theta), \quad (\text{C3})$$

where

$$D(\phi, \theta) = \begin{pmatrix} e^{i\mu} & 0 \\ 0 & e^{-i\mu} \end{pmatrix}, \quad (\text{C4})$$

$$S(\phi, \theta) = \begin{pmatrix} \sin(\theta/2) & \sin(\theta/2) \\ \cos(\theta/2) - e^{i(\frac{\phi}{2} + \mu)} & \cos(\theta/2) - e^{i(\frac{\phi}{2} - \mu)} \end{pmatrix}. \quad (\text{C5})$$

The matrices $D(\phi, \theta)$ and $S(\phi, \theta)$ are defined in the $\{|11\rangle, |02\rangle\}$ basis, and the angle μ satisfies

$$\cos \mu = \cos(\phi/2) \cos(\theta/2). \quad (\text{C6})$$

Substituting into Eq. (C1), the expectation value becomes

$$\begin{aligned} & \left| \langle 11 | \hat{R}^n(\phi, \theta) | 11 \rangle \right|^2 \\ &= \left| \begin{pmatrix} 1 & 0 \end{pmatrix} \cdot S(\phi, \theta) D^n(\phi, \theta) S^{-1}(\phi, \theta) \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 \\ &= \frac{\cos^2(\theta/2) - \cos(2\mu) + \sin^2(\theta/2) \cos(2n\mu)}{2 \sin^2 \mu}. \end{aligned} \quad (\text{C7})$$

This expression can be written as $A + B \cos(2n\mu)$, where A and B are independent of n . The signal therefore consists of an oscillation with frequency 2μ and amplitude B , which determines the SNR of the coherent error amplification experiment.

We can express the amplitude B as

$$2B = \frac{1}{1 + \left(\frac{\sin \frac{\phi}{2}}{\tan \frac{\theta}{2}} \right)^2}. \quad (\text{C8})$$

This is a Lorentzian function centered around $\phi = 0$, with a full width at half maximum of $2 \tan(\theta/2)$. For arbitrary ϕ and small θ , the amplitude B becomes small, leading to poor SNR and ineffective error amplification. This explains the low contrast observed near the optimal coupler flux pulse amplitude (where θ is small) in Fig. 2(b) of the main text.

Finally, we note that the signal oscillates at the frequency 2μ per gate, which satisfies $\theta \leq 2\mu \leq \pi$. Equation (C6) shows that $2\mu = \theta$ is satisfied for $\phi = \alpha + \beta = 0$, but in general μ depends on the uncontrolled

angle ϕ . This means that the frequency of the measured signal is systematically larger than θ , making it difficult to accurately extract the leakage angle based on the oscillation frequency alone, even if the lowered SNR is overcome by increasing the number of repetitions in the experiment.

3. Extensions to the standard error amplification

For completeness, in this section we present two known extensions to the standard error amplification experiment that enable accurate measurement of leakage to the $|02\rangle$ state. A key requirement for these methods is that the CZ gate itself remains unmodified during error amplification. Both extensions aim to compensate for a nonzero angle $\phi = \alpha + \beta$, which arises due to the levels $|11\rangle$ and $|02\rangle$ being nonresonant during the final calibrated gate and the additional free evolution during the added buffer time. If ϕ is effectively canceled—such that $\phi \ll \theta$ —then the signal amplitude in Eq. (C8) becomes large, allowing reliable amplification of leakage error.

The first approach, shown in Fig. 7(a), introduces a fixed delay τ after each CZ gate. The delay is chosen to satisfy $\tau(\omega_{|11\rangle} - \omega_{|02\rangle}) \approx \phi$, where $\omega_{|11\rangle}$ and $\omega_{|02\rangle}$ are the angular frequencies of the corresponding energy levels. This compensates for the accumulated Z rotation in the laboratory frame. Figure 7(b) shows experimental data for a fixed delay of 15.667 ns. Near the optimal coupler flux amplitude, the signal amplitude is significantly larger than in the standard error amplification experiment [Fig. 2(b) of the main text].

A more robust alternative is Floquet calibration [62], where a Z rotation operation with angle $-\phi$ is interleaved between CZ gates to counteract ϕ , see Fig. 7(c). The Z rotation is implemented on Q_1 via two $R_y^{01}(\pi) = \hat{R}_z(\gamma)\hat{R}_x(\pi)\hat{R}_z(-\gamma)$ gates with a relative phase difference $\gamma_0 - \gamma_1 = -\phi/2$, effectively realizing an $\hat{R}_z(-\phi)$ gate.

We note that in practice the microwave pulses implementing the gates are not zero duration. As a result, there is an additional constant component to the implemented Z rotation angles originating from the free Hamiltonian evolution and possibly other effects, for example a microwave-induced AC Stark shift [102]. Since the true laboratory frame ϕ is unknown, and these additional components only affect the reference, it is always possible to experimentally determine the optimal nominal Z rotation angle. Figure 7(d) shows data using Floquet calibration with an interleaved nominal Z rotation of -1.122 rad.

Both methods perform significantly better than the standard experiment. As illustrated in Fig. 7(e), the oscillation amplitudes at near-optimal coupler flux amplitudes are markedly increased. This results in sharper signals, enabling precise determination of the optimal coupler flux amplitude, as shown in Fig. 7(f), where a Lorentzian fit is used to extract the optimal amplitude x_0 .




FIG. 7. (a) Circuit diagram for error amplification with an additional delay τ inserted after each CZ gate. The system is initialized to $|11\rangle$ using an X_{π}^{01} gate on each qubit. During the delay, the system evolves under the free Hamiltonian. (b) Measured unwanted population in the $|02\rangle$ state as a function of the number of cycles n , and the coupler flux pulse amplitude. (c) Circuit diagram of the Floquet calibration method. A Z_{ϕ}^{01} rotation is interleaved between the CZ gates using two $X_{\pi}^{01}(\gamma)$ pulses with a relative phase difference $\gamma_0 - \gamma_1 = -\phi/2$ on Q_1 . (d) Measured unwanted population in the $|02\rangle$ state as a function of CZ gate number and coupler flux amplitude using Floquet calibration. Both (b) and (d) use three-state single-shot readout corrected for assignment error. (e) Vertical slices from (b) and (d), and the standard error amplification experiment in Fig. 2(b) (main text) at selected flux pulse amplitudes (dashed lines). (f) Unwanted population averaged over the number of cycles, as a function of the coupler flux pulse amplitude. The vertical dashed line indicates the center offset x_0 of Lorentzian fits (not shown).

However, these extensions are not practical for fine calibration tasks. The protocols require precise tuning of the delay τ or Z rotation angle ϕ , because the range over which the SNR is enhanced—determined by Eq. (C8)—is narrow. Such tuning is time-consuming and sensitive to fluctuations, even if performed adaptively. Moreover, varying gate parameters in an optimization routine would necessitate adaptation of these compensating operations at each point, limiting their usefulness.

APPENDIX D: PALEA MODEL

In this section, we present the theoretical framework underlying PALEA and demonstrate its specific application to measuring the leakage from $|11\rangle$ to $|02\rangle$. Finally, we compare the accuracy and precision of PALEA to the other

error amplification protocols discussed in the previous sections. In Appendix E 2 we show that PALEA is insensitive to other leakage processes, enabling accurate calibration of a single exchange process at a time.

Similarly to the MEADD protocol [64], PALEA can be generally applied to any excitation-preserving entangling gate as long as small off-resonant exchange rate between the two levels can be assumed. The only two required components are the ability to implement the dynamical decoupling step and to average the relevant phase. The protocol is also robust to miscalibration in single-qubit rotations used.

1. Defining the components of the PALEA model

Following the introduction of the PALEA protocol in Fig. 2(c), we aim to justify Eqs. (2)–(4). First, we define the DD step D using the single-qubit operators describing phased π rotations in appropriate subspaces,

$$D = \hat{R}_z^{12}(\phi_1)\hat{R}_x^{12}(\pi)\hat{R}_z^{12}(-\phi_1) \otimes \hat{R}_z^{01}(\phi_0)\hat{R}_x^{01}(\pi)\hat{R}_z^{01}(-\phi_0), \quad (\text{D1})$$

where $\hat{R}_z^{12}$ and $\hat{R}_x^{12}$ denote rotations around the Z and X axis in the $\text{span}(|1\rangle, |2\rangle)$ subspace, and $\hat{R}_z^{01}$ and $\hat{R}_x^{01}$ are the corresponding rotations in the $\text{span}(|0\rangle, |1\rangle)$ subspace. The dynamical decoupling operator D is designed to keep the state in the subspace $\text{span}(|11\rangle, |02\rangle)$,

$$\begin{aligned} D|11\rangle &= -e^{i(\phi_1-\phi_0)}|02\rangle, \\ D|02\rangle &= -e^{i(\phi_0-\phi_1)}|11\rangle, \end{aligned} \quad (\text{D2})$$

and its action can be written as

$$D = \hat{R}_z(\phi_0 - \phi_1)\hat{R}_x(\pi)\hat{R}_z(\phi_1 - \phi_0),$$

recovering Eq. (2).

Next, for rotations of angle π we have $\hat{R}_z(\lambda)\hat{R}_x(\pi) = \hat{R}_x(\pi)\hat{R}_z(-\lambda)$ for any angle λ , which allows us to combine the action of D and $\hat{U}_g$ —defined in Eq. (1)—into a single cycle unitary,

$$\begin{aligned} D\hat{U}_g &= \hat{R}_z(\phi_0 - \phi_1)\hat{R}_x(\pi)\hat{R}_z(\phi_1 - \phi_0)\hat{R}_z(\alpha)\hat{R}_x(\theta)\hat{R}_z(\beta) \\ &= \hat{R}_z(2\phi_0 - 2\phi_1 - \alpha - \pi)\hat{R}_x(\pi - \theta)\hat{R}_z(\beta + \pi). \end{aligned} \quad (\text{D3})$$

Finally, using Eq. (C2) we recover Eq. (3),

$$V_n(\theta, \Delta\phi) = \left| \langle 11 | \left[\hat{R}_z(\Delta\phi)\hat{R}_x(\pi - \theta) \right]^n | 11 \rangle \right|^2,$$

where $\Delta\phi = 2\phi_0 - 2\phi_1 - \alpha + \beta$.

The importance of the added π rotation may be intuitively understood by careful inspection. Equation (3) is

a special case of Eq. (C1), and the solution proceeds identically. According to Eq. (C6), the cycle eigenvalues satisfy

$$\begin{aligned} \cos \mu &= \cos(\Delta\phi/2) \cos(\pi/2 - \theta/2) \\ &= \cos(\Delta\phi/2) \sin(\theta/2). \end{aligned} \quad (\text{D4})$$

Since the measured population depends only on $|2\mu|$, we may, without loss of generality, choose the positive branch $0 \leq \mu \leq \pi/2$. Taking into account the possible values of $\cos(\Delta\phi/2)$, this leads to the tighter bound $\pi - \theta \leq 2\mu \leq \pi$, which spans a narrow range if $\theta \ll \pi$.

Thus, Eq. (3) describes an iterated two-level rotation in the $\text{span}(|11\rangle, |02\rangle)$ subspace, where the rotation axis is nearly orthogonal to the Z axis when the over-rotation angle θ is small.

Following a derivation analogous to Eq. (C8), we obtain the signal amplitude:

$$2B = \frac{1}{1 + \left(\frac{\sin(\Delta\phi/2)}{\tan(\pi/2 - \theta/2)} \right)^2} \approx 1. \quad (\text{D5})$$

This indicates that the iterated rotation yields near-maximal contrast between the populations of $|11\rangle$ and $|02\rangle$ as a function of the number of cycles n for most values of $\Delta\phi$.

Comparing the PALEA protocol to the Floquet calibration discussed in Appendix C 3, we find that the cycle unitary in the absence of a DD layer is dominated by the Z component—unless the combined phase $\phi = \alpha + \beta \ll \pi$. Consequently, the signal amplitude in such protocols is large only near $\phi \approx 0$ for most values of θ , as illustrated in the simulation in Fig. 8(a) for a fixed over-rotation angle $\theta = -0.380$. In contrast, when a DD layer is applied—as in the MEADD protocol—the signal amplitude remains large over a wide range of ϕ values and over-rotation angles, as shown in Fig. 8(b).

Averaging over the angle ϕ corresponds to the “phase-averaged” aspect of PALEA. Without a DD layer, this averaging suppresses the signal significantly, resulting in a low-contrast response [Fig. 8(c)]. With a DD layer, however, the same averaging preserves high signal amplitude—this is the core advantage of the PALEA protocol.

Finally, Fig. 8(d) shows the signal averaged over both ϕ and the number of cycles while sweeping the over-rotation angle θ . In the presence of a DD layer, this produces a sharp and well-defined dip, enabling precise calibration. Without a DD layer, the dip becomes shallow and less informative.

2. Phase tracking in the PALEA protocol

The PALEA model is defined in the laboratory frame, for reasons outlined earlier in Appendix C 1. While the




FIG. 8. (a) and (b) Unwanted population as a function of the number of cycles and the Z rotation angle ϕ , for an over-rotation angle $\theta = -0.380$. Panel (a) shows the case without dynamical decoupling, corresponding to a Floquet-type calibration, while panel (b) shows the case with a dynamical decoupling layer, corresponding to the matrix-element amplification using dynamical decoupling-type calibration. (c) Unwanted population averaged over ϕ , plotted as a function of the number of cycles. (d) Unwanted population averaged over both ϕ and the number of cycles, plotted as a function of the over-rotation angle θ .

phase difference $\Delta\phi$ may vary between repetitions of the experiment, it must remain constant across the repeated cycles within a single gate sequence. Specifically, if the computational frame of each qubit is defined by the frequency of the microwave drive used to implement the required rotation, then in the laboratory frame the phase difference $\Delta\phi$ evolves at a frequency given by $\frac{1}{\hbar}(E_{|11\rangle} - E_{|02\rangle})$, where $E_{|11\rangle}$ and $E_{|02\rangle}$ are the energies of the respective states. In most experimental setups, the phase of the microwave reference used to implement the single-qubit gates is allowed to progress freely at its frequency. Without correction, this causes $\Delta\phi$ to drift between cycles. To prevent this undesired evolution, one must track the accumulated phase and actively compensate for it. In practice, this can be achieved by regressing the microwave phase by ϕ_i after each cycle, with $\phi_1 = -\tau\tilde{\omega}_1$ for qubit Q₁ and $\phi_2 = -\tau(\tilde{\omega}_2 + \alpha_2)$ for qubit Q₂, where τ is the duration of one cycle. In most experimental frameworks it only requires setting the virtual Z rotations attached to the CZ gate to the above values.

3. Solving for the expectation value in the PALEA model

Since the eigenvalues and eigenvectors of the cycle with the added π rotation is less sensitive to $\Delta\phi$, averaging over $\Delta\phi$ results in a high-contrast signal from which θ can be

extracted. To justify this intuition and solve the integral in Eq. (4), we prove the following lemma:

Lemma D1. For $\hat{U} \in \text{SU}(2)$, $n \in \mathbb{N}$, the average of the matrix element

$$\left| \langle 0 | \left[\hat{R}_z(\phi) \hat{U} \right]^n | 0 \rangle \right|^2 \quad (\text{D6})$$

over ϕ is given by

$$I_n(\theta) = 1 - \sin^2(\theta/2) \sum_{m=0}^{n-1} P_m(\cos \theta), \quad (\text{D7})$$

where θ is such that $\hat{U} = \hat{R}_z(\alpha)\hat{R}_y(\theta)\hat{R}_z(\beta)$ and P_m is the m th Legendre polynomial.

Proof. The average can be expressed as an integral using the ZYZ Euler angle decomposition for $\hat{U}$,

$$I_n(\theta) = \frac{1}{2\pi} \int_0^{2\pi} d\phi \left| \langle 0 | \left[\hat{R}_z(\phi) \hat{R}_z(\alpha) \hat{R}_y(\theta) \hat{R}_z(\beta) \right]^n | 0 \rangle \right|^2. \quad (\text{D8})$$

Note

$$\begin{aligned} & \left| \langle 0 | \left[\hat{R}_z(\phi) \hat{R}_z(\alpha) \hat{R}_y(\theta) \hat{R}_z(\beta) \right]^n | 0 \rangle \right|^2 \\ &= \left| \langle 0 | \left[\hat{R}_z(\phi + \alpha + \beta) \hat{R}_y(\theta) \right]^n | 0 \rangle \right|^2 \end{aligned} \quad (\text{D9})$$

utilizing Eq. (C2), and we can define

$$\hat{R}(\phi, \theta) = \hat{R}_z(\phi) \hat{R}_y(\theta). \quad (\text{D10})$$

■

a. Expanding the integral

Following the derivation in Eqs. (C1)–(C7), the integrand of Eq. (D8) can be expanded as

$$\begin{aligned} & \left| \langle 0 | \hat{R}^n(\phi, \theta) | 0 \rangle \right|^2 \\ &= \left| \begin{pmatrix} 1 & 0 \end{pmatrix} \cdot S(\phi, \theta) D^n(\phi, \theta) S^{-1}(\phi, \theta) \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 \\ &= \frac{\cos^2(\theta/2) - \cos(2\mu) + \sin^2(\theta/2) \cos(2n\mu)}{2 \sin^2 \mu}, \end{aligned} \quad (\text{D11})$$

leading to

$$I_n(\theta) = \frac{1}{2\pi} \int_0^{2\pi} d\phi \frac{\cos^2(\theta/2) - \cos[2\mu(\phi)] + \sin^2(\theta/2) \cos[2n\mu(\phi)]}{2 \sin^2[\mu(\phi)]}. \quad (\text{D12})$$

b. Generating function

We define the generating function $I(z; \theta)$ as a formal power series

$$I(z; \theta) = \sum_{n=0}^{\infty} I_n(\theta) z^n. \quad (\text{D13})$$

Substituting here $I_n(\theta)$ from Eq. (D12), we obtain

$$I(z; \theta) = \frac{1}{4\pi} \int_0^{2\pi} d\phi \frac{\cos^2(\theta/2) - \cos[2\mu(\phi)]}{\sin^2[\mu(\phi)]} \sum_{n=0}^{\infty} z^n + \frac{\sin^2(\theta/2)}{4\pi} \int_0^{2\pi} d\phi \frac{\sum_{n=0}^{\infty} \cos[2n\mu(\phi)] z^n}{\sin^2[\mu(\phi)]}. \quad (\text{D14})$$

The first term contains the geometric series

$$\sum_{n=0}^{\infty} z^n = \frac{1}{1-z}, \quad (\text{D15})$$

while the second contains a related series

$$\begin{aligned} \sum_{n=0}^{\infty} \cos(2n\mu) z^n &= \frac{1}{2} \sum_{n=0}^{\infty} \left[(e^{2i\mu} z)^n + (e^{-2i\mu} z)^n \right] \\ &= \frac{1}{2} \left(\frac{1}{1 - e^{2i\mu} z} + \frac{1}{1 - e^{-2i\mu} z} \right) \\ &= \frac{1 - \cos(2\mu)z}{1 - 2 \cos(2\mu)z + z^2} \\ &= \frac{1 + z - 2z \cos^2 \mu}{(1 + z)^2 - 4z \cos^2 \mu}. \end{aligned} \quad (\text{D16})$$

Using these, the generating function becomes

$$I(z; \theta) = \frac{1}{4\pi} \int_0^{2\pi} d\phi \frac{\cos^2(\theta/2) - \cos[2\mu(\phi)]}{\sin^2[\mu(\phi)]} \frac{1}{1-z} + \frac{\sin^2(\theta/2)}{4\pi} \int_0^{2\pi} d\phi \frac{1}{1 - \cos^2[\mu(\phi)]} \frac{1 + z - 2z \cos^2[\mu(\phi)]}{(1 + z)^2 - 4z \cos^2[\mu(\phi)]}. \quad (\text{D17})$$

Performing partial fraction decomposition to the integrand of the second term by treating $\cos^2[\mu(\phi)]$ as the independent variable, we can further simplify this as

$$I(z; \theta) = \frac{1}{1-z} - \frac{z(1+z)}{1-z} \frac{\sin^2(\theta/2)}{2\pi} \int_0^{2\pi} \frac{d\phi}{(1+z)^2 - 4z \cos^2[\mu(\phi)]}. \quad (\text{D18})$$

Based on Eq. (C6), we have

$$\cos^2[\mu(\phi)] = \cos^2(\phi/2) \cos^2(\theta/2) = \frac{\cos^2(\theta/2)}{2} [1 + \cos(\phi)]. \quad (\text{D19})$$

Substituting then leads to

$$I(z; \theta) = \frac{1}{1-z} - \frac{z(1+z)}{1-z} \frac{\sin^2(\theta/2)}{2\pi} \int_0^{2\pi} \frac{d\phi}{(1+z)^2 - 2z \cos^2(\theta/2) + 2z \cos^2(\theta/2) \cos \phi}. \quad (\text{D20})$$

The remaining integral can be converted to a contour integral along the unit circle in the complex plane by defining $w = e^{i\phi}$, giving us

$$I(z; \theta) = \frac{1}{1-z} - \frac{z(1+z)}{1-z} \frac{\sin^2(\theta/2)}{2\pi i} \oint_{|w|=1} \frac{dw}{-z \cos^2(\theta/2)(w^2 + 1) + [(1+z)^2 - 2z \cos^2(\theta/2)]w}. \quad (\text{D21})$$

The integrand has two first-order poles at

$$w_{\pm} = -1 + \frac{1+z}{2z} \times \frac{1+z \pm \sqrt{1-2z \cos \theta + z^2}}{\cos^2(\theta/2)}, \quad (\text{D22})$$

but only w_- is inside the unit circle. Using the residue theorem, we finally obtain

$$I(z; \theta) = \frac{1}{1-z} - \frac{z}{1-z} \frac{\sin^2(\theta/2)}{\sqrt{1-2z \cos \theta + z^2}}. \quad (\text{D23})$$

c. Extracting the integrals

Utilizing again the geometric series (D15) as well as the generating function

$$\frac{1}{\sqrt{1-2xt+t^2}} = \sum_{n=0}^{\infty} P_n(x) t^n \quad (\text{D24})$$

of the Legendre polynomials, we can expand the generating function (D23) as

$$I(z; \theta) = \sum_{n=0}^{\infty} z^n - \sin^2(\theta/2) \sum_{n,m=0}^{\infty} P_m(\cos \theta) z^{n+m+1}. \quad (\text{D25})$$

Changing the summation variables from (n, m) to $(n+m, m)$, this can be further written as

$$I(z; \theta) = 1 + \sum_{n=1}^{\infty} \left[1 - \sin^2(\theta/2) \sum_{m=0}^{n-1} P_m(\cos \theta) \right] z^n. \quad (\text{D26})$$

From here we can immediately read that

$$I_0(\theta) = 1 \quad (\text{D27})$$

and

$$I_n(\theta) = 1 - \sin^2(\theta/2) \sum_{m=0}^{n-1} P_m(\cos \theta). \quad (\text{D28})$$

According to the above lemma, the signal from the phase-averaged experiment is given by

$$\begin{aligned} I_n(\pi - \theta) &= 1 - \cos^2(\theta/2) \sum_{m=0}^{n-1} P_m(-\cos \theta) \\ &= 1 - \cos^2(\theta/2) \sum_{m=0}^{n-1} (-1)^m P_m(\cos \theta) \end{aligned}$$

using the parity condition of Legendre polynomials, and we finally recover Eq. (5).

4. Analyzing the robustness of PALEA

For small θ , Eq. (5) describes a signal flipping between 0 and 1. To regularize it, we can consider

$$\tilde{I}_n(\theta) = (-1)^n \left[I_n(\pi - \theta) - \frac{1}{2} \right] + \frac{1}{2} \quad (\text{D29})$$

instead. It can be shown to have a limiting form when $n \gg 1$ and $\theta \ll 1$,

$$\tilde{I}_n(\theta) \approx \frac{1}{2} - \frac{1}{2} J_0(n\theta), \quad (\text{D30})$$

where J_0 is a Bessel function of the first kind. This expression is easier to understand and analyze than the full model. Its maximum is approximately 0.7, which is another way of confirming that higher SNR can be achieved for rotations which are nearly π , rather than nearly 0, seen in Fig. 8.

If a measured quantity depends on the phase as a sine, $V \propto \sin(n\theta)$, the phase can only be estimated up to the Nyquist frequency π/n . However, the signal following Eq. (D30) is not entirely periodic, and only assumes values between 0 and 0.35 once, for values of $n\theta < 1.8$. Thus if a high population of the wanted outcome state is measured after a large number of gates, one might be certain the angle θ must be small and its absolute value may be calculated from a single phase-averaged measurement. This greatly simplifies calibration, as one might fix n at a value guaranteeing required precision, and vary the gate parameters to unequivocally identify a region where $|\theta|$ is below a certain threshold. This is why a simple sweep of a coupler pulse amplitude, shown in Fig. 2(d) works well as a quick calibration experiment.

Instead of fitting a Lorentzian, one might fit the full model [Eq. (5)] directly to the two-dimensional experimental data. The data seen in Figs. 2(d) and 9(a) were collected using single-shot readout calibrated for three-state discrimination and corrected for assignment error. Since the ideal π rotation would flip the state between $|11\rangle$ and $|02\rangle$, we define the correct outcome to be $|11\rangle$ for even n and $|02\rangle$ for odd n . Then, the unwanted population is defined as

$$UP_n = 1 - \frac{\text{Number of correct outcomes}}{\text{Number of the outcomes } 11 \text{ or } 02}, \quad (\text{D31})$$

which is an unbiased estimator in presence of incoherent processes causing the decay to the single- and zero-excitation subspaces [62]. The variance increases as the total number of in-subspace repetitions decreases. Figure 9(b) shows the fit of the full model [Eq. (5)] to the data, assuming that θ depends on the coupler pulse amplitude polynomially. Figure 9(c) shows the data and fit averaged over the number of cycles, and Fig. 9(d) the θ values extracted as fit parameters.




FIG. 9. (a) Experimental data of the unwanted population as a function of the number of cycles and the coupler flux pulse amplitude. Here, population in the $|02\rangle$ ($|11\rangle$) state is considered unwanted (wanted) for even n , and population in $|11\rangle$ ($|02\rangle$) state is unwanted (wanted) for odd n . (b) Two-dimensional fit of the unwanted population as a function of the number of cycles and the coupler flux pulse amplitude to the model in Eq. (D29). (c) Unwanted population averaged over the number of cycles as a function of the coupler flux pulse amplitude. The solid line indicates the fit to the phase-averaged leakage error amplification model averaged over the number of cycles. (d) The over-rotation angles, θ , as a function of the coupler flux pulse amplitude, extracted from experimental data as parameters of the fit shown in (b). The dashed line indicates the optimal coupler flux pulse amplitude, where $\theta = 0$.

The Legendre series in Eq. (5) may appear computationally expensive to evaluate and fit. However, the evaluation using an appropriate version of the Clenshaw algorithm using the Legendre polynomial recurrence relation [103–105] has time complexity $\mathcal{O}(n)$ for a single n and thus $\mathcal{O}(n^2)$ for all numbers of cycles until n and is implemented in the `numpy` Python library as `polynomial.legendre.legval`. Alternatively, one can use the Bessel approximation from Eq. (D30) for fitting.

The standard error in the estimate of θ is inversely proportional to the derivative,

$$\frac{\partial \tilde{I}_n}{\partial \theta} \approx -\frac{n}{2} J_1(n\theta). \quad (\text{D32})$$

The Bessel function $J_1(x)$ has an asymptotic form $J_1(x) \sim (x/2)$ for $x < 1$ [106], so for $n\theta < 1$, the error in the estimate of θ follows $\mathcal{O}(1/n^2)$, which is the same scaling one could reach with a sine function.

For large x , $J_1(x)$ has an envelope given by the absolute value of the Hankel function of the first kind $H_1(x) = J_1(x) + iY_1(x)$, namely $J_1(x) < |H_1(x)|$, and $|H_1(x)|$ has an asymptotic form $|H_1(x)| \sim \sqrt{2/\pi x}$ [106].

In turn, for large n we obtain $\frac{\partial \tilde{I}_n}{\partial \theta} < \sqrt{n/2\pi\theta}$, recovering the $\mathcal{O}(1/\sqrt{n})$ scaling. Thus, PALEA is at its most effective with regard to precision in characterizing θ when $n\theta$ is small and calibration of the set of pulse parameters for which $\theta = 0$. This covers the relevant use cases, since error amplification is unnecessary for large θ .

To concretely demonstrate the sensitivity of PALEA, and verify these relationships for the full model defined in Eq. (D29), further numerical simulations of a realistic calibration procedure may be performed, as demonstrated in the following section.

5. Simulating the estimation power of PALEA

We aim to illustrate the ability of PALEA to measure the angle θ in the context of a realistic automatic calibration procedure. Such a procedure must be fast, that is, use a limited number of repetitions, each corresponding to one quantum measurement, and not require human intervention. We compare PALEA to the nonadaptive Floquet calibration using Z rotations [Fig. 7(c)] and to the MEADD protocol. We assume a fixed repetition budget of N total repetitions for each procedure. For each repetition, we calculate an ideal Born-rule probability p of measuring the target state— $|11\rangle$ for the procedures described in this work. We then generate a random bit from a Bernoulli distribution with probability p , which is flipped with probability ϵ , corresponding to a readout error. The generated bits are then averaged to obtain a data point. If a procedure requires collecting k data points to fit them with a model, the total amount of repetitions N is split equally so that each point receives N/k repetitions.

To perform the Floquet calibration, we need to first find the appropriate angle ϕ for the Z rotation. We simulate an experiment using ten total cycles, sweeping the angle ϕ and fit the generated points with the model given in Eq. (C7). Using the recovered estimate for ϕ , we then simulate another experiment sweeping the number of cycles from 0 to 49, and fit the data with a sinusoidal model, extracting an estimate for θ . We divide the repetition budget equally between the two experiments.

For MEADD, no additional calibrations are needed. We simulate two experiments using the PALEA circuit, but fixing the angle $\phi_0 - \phi_1$ to 0 and $\pi/2$ for the first and second experiment, respectively, following the `iSWAP` MEADD procedure from Ref. [64]. Because we are not interested in the sign of θ , we skip the third circuit. The number of cycles is swept from 0 to 48 with step 2. The repetition budget is divided equally between the two experiments and each realization of the procedure assumes a different, random value of $\alpha - \beta$ defined in Eq. (1). We fit the two data sets with sinusoidal models, extracting two angles μ_0 and μ_1 , and calculate $\theta = \sqrt{\mu_0^2 + \mu_1^2}$ following a small angle approximation of Eq. (D4).




FIG. 10. Simulations of phase-averaged leakage error amplification, nonadaptive Floquet calibration, and hypothetical matrix-element amplification using dynamical decoupling protocol assuming a finite number of repetitions and 0.05 readout error. (a) Simulated frequencies of measuring the target state in phase-averaged leakage error amplification, using total of 4000 repetitions with target angle $\theta = 0.15$. The fit to the noisy simulated points is the source of the estimated angle. (b) Histograms of estimated θ values extracted from 1000 runs of Floquet calibration, matrix-element amplification using dynamical decoupling and phase-averaged leakage error amplification. Matrix-element amplification using dynamical decoupling and phase-averaged leakage error amplification both estimate the angle correctly on average, Floquet calibration overestimates it, because the Z angle rotation is not correct due to a finite number of repetitions, introducing a systematic error. (c) Mean squared error produced by each estimator with 4000 total repetitions as a function of the target angle. For small angles, phase-averaged leakage error amplification and matrix-element amplification using dynamical decoupling are equally good, but then phase-averaged leakage error amplification signal decreases because it is proportional to a Bessel function J_0 and not a cosine. Floquet calibration underperforms because of the systematic error. (d) Mean squared error produced by each estimator with the target angle $\theta = 0.15$ as a function of the total repetition budget. Mean squared error of phase-averaged leakage error amplification and matrix-element amplification using dynamical decoupling is proportional to $1/N$. The Floquet calibration struggles for small number of repetitions because of the compounding error from incorrect estimation of the Z rotation angle, but it reaches the same sensitivity for large number of repetitions.

For PALEA, we simulate the signal averaged over 40 values of $\phi_0 - \phi_1$ and fit the data with the model from Eq. (5). The number of cycles is swept identically as in the MEADD simulation. An example of a generated signal and fit is shown in Fig. 10(a) using $N = 4000$ and target $\theta = 0.15$.

We perform each procedure 1000 times, generating estimates of θ and present their histograms in Fig. 10(b). PALEA and MEADD both estimate θ correctly, but the nonadaptive Floquet calibration seems to systematically overestimate it. This is because the fitted angle μ follows Eq. (C6), which is larger than θ if $\phi \neq 0$ —the estimation of ϕ is itself erroneous if not enough repetitions are used for it. This demonstrates the fundamental reason why this procedure is too resource-intensive to be used in automatic calibration—the two estimation errors compound.

The target angle θ is then varied, shown in Fig. 10(c), revealing that PALEA and MEADD have equal sensitivity for small angles. However, for larger θ values, the lower signal in PALEA—following a Bessel function—causes lowered SNR compared to a pure sinusoidal oscillation, as discussed following Eq. (D32). The Floquet calibration yields an order of magnitude higher mean squared error (MSE), because of the systematic error inherent to the procedure.

Finally, we investigate the MSE when the total repetition budget N is varied. As shown in Fig. 10(d), the Floquet calibration eventually reaches the same MSE as the two other methods, but only after nearly a million total repetitions.

As discussed in Appendix D4, PALEA is at its best when estimating smaller angles θ . This simulation does not include any decoherence errors, which would normally limit the total number of cycles performed, so the relevant quantity here is the product of the angle θ and the maximum number of cycles used. PALEA performs best when this product is approximately 5. In a realistic calibration procedure, a previous coarse experiment using just one CZ gate would have been performed, identifying a region where the angle θ is small. Thus the increased MSE for larger angles is not relevant for such a procedure.

While PALEA suffers from slightly lowered SNR in ideal conditions, in realistic settings the $|11\rangle - |02\rangle$ dynamics are not the only dynamics present. The CZ gate suffers from out-of-subspace effects which may pollute the calibration, and averaging the phases might be helpful to reduce experimental sensitivity to this pollution.

6. Broader applicability of PALEA

The PALEA protocol is discussed in this work in the context of characterizing and minimizing the rate of coherent exchange in the $\text{span}(|11\rangle, |02\rangle)$ subspace. The underlying mechanisms behind the protocol—phase averaging and dynamical decoupling—could however be applied to measurement and optimization of small unwanted exchange processes across quantum computing modalities. The only requirement is an implementation of a phased π rotation between the participating states, replacing the dynamical decoupling layer used in this work.

One example is a diabatic CZ gate in quantum dot spin qubits [71,72]. It is implemented by tuning an exchange coupling J between the two quantum dots for a short duration to perform a 2π rotation in the span($\{|\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle\}$) subspace while acquiring geometric conditional phase. The remaining nonzero SWAP rate could be amplified and measured using a variation of the PALEA protocol, utilizing single-qubit X rotations, similarly to the experiments performed in this work (see Appendix F 1) to measure SWAP error in CZ gates.

A general framework for geometric two-qubit gates using stimulated Raman adiabatic passage (STIRAP) was introduced in Ref. [70]. The population of the state $|11\rangle$ is rotated through two auxiliary states to acquire a geometric phase. Whenever an auxiliary state is such that it is possible to design a phased π rotation connecting it to $|11\rangle$ —as is the case with the fluxonium-transmon system discussed there—it may be possible to use PALEA to amplify the leakage process to that state and measure the associated exchange rate precisely.

APPENDIX E: OUT-OF-SUBSPACE DYNAMICS

In addition to the main exchange dynamics between the $|11\rangle$ and $|02\rangle$ states, CZ gates might induce other unwanted transitions, such as leakage to the coupler or other neighboring modes, if any are present. First, we will review the tools to measure the effect of those additional dynamics and then discuss how to prevent them from introducing systematic errors in the CZ gate calibration.

1. Measuring coupler leakage

While the extensions of the standard error amplification experiments discussed in Appendix C are less practical for leakage amplification calibration, they are effective tools for probing out-of-subspace dynamics, including coherent leakage to the coupler [18].

The goal is to observe coherent population exchange between the $|101\rangle$ and $|011\rangle$ states (with notation $|Q_1, TC, Q_2\rangle$)—i.e., excitation transfer from Q_1 to the coupler. This two-level system, with interaction strength g_{1c} , can be described with a single Hamiltonian

$$\hat{H}/\hbar = \begin{pmatrix} 0 & g_{1c} \\ g_{1c} & \omega_c(t) - \omega_1 \end{pmatrix}, \quad (\text{E1})$$

where $\omega_c(t)$ is the time dependent coupler frequency and ω_1 is the static qubit frequency. Although the coupling g_{1c} depends on $\omega_c(t)$ as $g_{1c}(t) = \beta_{1c}\sqrt{\omega_1\omega_c(t)}$ [22], we approximate it as constant. Transforming to an interaction frame with $\hat{H}_0(t) = (\omega_c(t) - \omega_1)|011\rangle\langle 011|$ and

corresponding unitary

$$\hat{U}_0 = e^{-i\int^t dt' \hat{H}_0/\hbar} = \exp\left(-\frac{i}{2}\hat{I}\left[-\omega_1 t + \int^t dt' \omega_c(t')\right]\right), \quad (\text{E2})$$

where $\hat{I}$ is the identity operator, yields a Landau-Zener Hamiltonian

$$\hat{H}/\hbar = \frac{1}{2}[(\omega_c(t) - \omega_1)\hat{\sigma}_z + 2g_{1c}\hat{\sigma}_x]. \quad (\text{E3})$$

Analogously to Eq. (C1), the unitary $U_c = e^{-i\int^t dt' \hat{H}/\hbar}$ can be decomposed into $\hat{R}_z$ and $\hat{R}_y$ rotations as

$$\hat{U}_c = \hat{R}_z(\alpha_c)\hat{R}_x(\theta_c)\hat{R}_z(\beta_c), \quad (\text{E4})$$

where angle θ_c quantifies the leakage process between the $|101\rangle$ and $|011\rangle$ states, and the angles α_c and β_c are the Z axis rotation angles. Again, the goal is to amplify the rotation $\hat{R}_x(\theta_c)$ for accurate measurement of the coupler leakage.

Although Slepian-shaped flux pulses reduce the leakage to the coupler [18,58], some leakage will persist. To coherently amplify the angle θ_c , the total Z rotation $\alpha_c + \beta_c$ must vanish, as is evident from Eq. (E4). This can be accomplished by inserting delays or phase corrections between the CZ gates. Here, we choose the delay approach, shown in Fig. 7(a).

We perform the experiment using 30 cycles of CZ gates interleaved with delays. The measurement of $|02\rangle$ [Fig. 11(a)] reveals the coherent exchange between the $|11\rangle$ and $|02\rangle$ states described earlier [vertical slices of this figure correspond to horizontal slices of Fig. 2(b) and Fig. 7(b)]. In contrast, the population of $|01\rangle$, seen in Fig. 11(b), shows diagonal stripes. We identify these stripes as leakage to the $|011\rangle$ state, i.e., leakage to the coupler. A Fourier transform of this signal along the delay axis reveals a frequency of 1.129 GHz, matching the qubit-coupler detuning ($4.799 - 3.672$ GHz = 1.127 GHz; see Table II), confirming that the peaks originate from the coupler mode. The precise delays at which the peaks occur depend on the amplitude of the coupler flux pulse approximately linearly. Note that the experiment is effectively a discrete pulsed spectroscopy, revealing a plethora of possible processes, which interfere with the coupler leakage process.

Next, we focus on characterizing the leakage error using the amplitude corresponding to the final calibrated CZ gate, which minimizes the $|02\rangle$ leakage. We select the delay of 7.26 ns, at which the coupler leakage process is amplified [Fig. 11(c)], by modeling the data using Eqs. (C6) and (C7), and setting $\phi = \tau(\omega_1 - \omega_c)$.

Finally, we measure the population of $|01\rangle$ versus the number of CZ gates, shown in Fig. 11(d). The measured population consists of the sum of the $|011\rangle$ state




FIG. 11. (a) Measured population of $|02\rangle$ and (b) $|01\rangle$ vs flux pulse amplitude of the tunable coupler and delay between CZ gates, τ . The number of cycles is fixed to 30. (c) Zoom in on one of the peaks in population of $|01\rangle$ for the optimal CZ gate amplitude, indicated by the orange segment in (b). The solid line shows a fit to the model in Eq. (C7) with 30 cycles. (d) Evolution of the population of $|01\rangle$ as a function of the number of cycles, using the delay $\tau = 7.26$ ns. The solid line indicates a fit to a model including the coherent coupler leakage to $|011\rangle$ and the incoherent damping involving $|001\rangle$.

population, involved in the coherent process, and $|001\rangle$ state population, where the higher frequency qubit may have incoherently decayed. We model the $|01\rangle$ population assuming the $|101\rangle - |011\rangle$ is the only coherent exchange, characterized by Y and Z rotation angles, accompanied by the incoherent amplitude and phase damping dynamics. We obtain $\theta_c = 0.041 \pm 0.001$, $\alpha_c + \beta_c = 0.025 \pm 0.001$. The signal in Fig. 11(d) decays mostly due to dephasing in the $\text{span}(|101\rangle, |011\rangle)$ subspace. Additionally, due to limitation in the temporal resolution of the used AWG, the Z rotation component is not fully canceled, resulting in a slightly reduced population oscillation amplitude.

We note that since only the states of Q_1 and Q_2 are measured, distinguishing coupler leakage from leakage to other modes becomes more difficult in larger QPUs, where unmeasured spectator modes may interfere with calibration [107]. In our case, the device is a small test chip optimized for high-fidelity operation, with the coupler as the only relevant leakage mode.

Unfortunately, methods requiring the use of X rotations—such as PALEA—could not be used to investigate

coupler leakage, as the device is missing a dedicated driveline for the coupler. However, Floquet calibration methods can reveal and quantify leakage to modes outside the computational subspace—such as to the coupler—with enough measurement time available. They are however less effective when the goal is to quickly and efficiently isolate and quantify leakage to the $|02\rangle$ state as part of the automated calibration procedure. In such cases, leakage to other spectator modes can interfere with the measurement, potentially distorting the inferred $|02\rangle$ leakage and leading to misleading calibration outcomes. Hence, there is a need for a more robust method that selectively amplifies the target leakage process while remaining insensitive to unwanted dynamics.

2. Coupler leakage effects in PALEA

The out-of-subspace dynamics, such as coupler leakage, may bleed into the measurement of the $|11\rangle - |02\rangle$ dynamics, interfering with the data. The goal of this section is to investigate the effect of these dynamics on the results of the PALEA protocol and to confirm that it solves this problem.

Figure 12(a) shows the $|11\rangle$ state population from the same run as Fig. 2(b). The coupler leakage signatures also shown in Fig. 11(b), bleed into this measurement. In contrast, no such traces appear in the $|11\rangle$ state population when measured with PALEA. An intuitive understanding of this may rely on the fact that the amplification of the coupler population occurs only for certain repeated Z rotations in the $\text{span}(|101\rangle, |011\rangle)$ subspace, and PALEA averages over all of them, converting the coupler leakage to incoherent background. This is not entirely trivial however, as the dynamical decoupling step D would excite any potential coupler leakage state $|011\rangle$ to a 4-excitation state $|112\rangle$. Assuming the energy levels $|101\rangle, |002\rangle, |011\rangle$ are all approximately on resonance during the gate, $|112\rangle$ would be the highest energy 4-excitation state. During idling however, this state is lower energy than all 4-excitation states involving the coupler in the ground state. Therefore all these states are crossed by $|112\rangle$ during the coupler flux pulse. The exact dynamical analysis is thus too complicated to be tackled analytically. We observe no pulse-amplitude-dependent coupler leakage in experiments, so we perform numerical simulations. We compare the PALEA dynamics to those of phase-averaged Floquet calibration circuits discussed in Appendix E 1. Figure 12(c) shows the total population of the states $|011\rangle$ and $|112\rangle$ in a Floquet calibration experiment, where the Z rotation angle has been chosen to show the worst-case scenario leakage. Simulation of the same gates, but using the PALEA circuit, is shown in Fig. 12(d). PALEA converts the coherent coupler leakage traces to a two orders of magnitude lower incoherent decay with the number of cycles.




FIG. 12. (a) Measured population of the state $|11\rangle$ in the standard error amplification experiment, same measurement as in Fig. 2(b). Visible vertical line at the amplitude 0.437 is a signature of coupler leakage. (b) Measured population of the state $|11\rangle$ in phase-averaged leakage error amplification, same measurement as in Fig. 2(d). No coupler leakage signatures visible. (c) and (d) Simulated total population of the coupler leakage states $|011\rangle$ and $|112\rangle$ as a function of number of cycles and the coupler frequency shift, using (c) the Floquet calibration sequence and (d) phase-averaged leakage error amplification. Simulations include the lowest four states of both qubits and the coupler, as well as realistic device parameters. The optimal Slepian pulse CZ gate uses coupler frequency shift 0.807 GHz.

To conclude the analysis of PALEA, we note that by converting various interfering processes into noise, PALEA can be considered robust to coherent interference. Thus by introducing averaging to MEADD, we trade a slightly worse SNR for additional robustness to interfering processes.

APPENDIX F: OTHER COHERENT ERRORS IN CZ GATE

1. SWAP error

A CZ gate is an example of a general excitation-preserving gate. In addition to two-excitation dynamics discussed above, single-excitation dynamics might additionally contribute to the overall coherent error. These dynamics may contribute to the ZZ rotation angle and the

single-qubit Z rotations, but those are calibrated separately. For this reason we assume the SWAP error mechanism is a unitary rotation around the $XX + YY$ operator,

$$\mathcal{E}_s = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_s/2) & -i \sin(\theta_s/2) & 0 \\ 0 & -i \sin(\theta_s/2) & \cos(\theta_s/2) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (\text{F1})$$

where we define SWAP angle to be θ_s . The simplest way to calculate the gate fidelity contribution is to calculate the entanglement fidelity of the above operator with identity [108], noting that it commutes with the ideal CZ gate. Defining $|\phi\rangle$ as a Bell state in the two copies of the system

$$|\phi\rangle = \frac{1}{2}(|00\rangle + |11\rangle + |22\rangle + |33\rangle), \quad (\text{F2})$$

we have

$$F_e(\mathcal{E}_s) = |\langle\phi| I \otimes \mathcal{E}_s |\phi\rangle|^2 = \left| \frac{1}{4} \left(2 + 2 \cos\left(\frac{\theta_s}{2}\right) \right) \right|^2 = \cos^4\left(\frac{\theta_s}{4}\right). \quad (\text{F3})$$

Thus the SWAP contribution to average gate infidelity is

$$1 - F(\mathcal{E}_s) = 1 - \frac{4F_e(\mathcal{E}_s) + 1}{4 + 1} = \frac{4}{5} - \frac{4}{5} \cos^4\left(\frac{\theta_s}{4}\right). \quad (\text{F4})$$

We measure the SWAP angle using PALEA, preparing the $|10\rangle$ state and using two phased- X_{π}^{01} gates as the dynamical decoupling step. Thus the experimental circuit is now identical to one of the MEADD [64] circuits, and only differs by the applied phase averaging. The SWAP angle is measured after setting the width of the coupler pulse and calibrating the coupler and qubit pulses for the ZZ rotation angles and the $|11\rangle - |02\rangle$ dynamics. The data analysis is identical to PALEA analysis, replacing the counts of 11 and 02 with counts of 10 and 01.

We find the SWAP angle below 0.1 rad for all measured widths, and below 0.04 rad for widths longer than 21 ns, which does not limit the gate fidelity, see Fig. 13(c). We therefore validate the design choices described in Appendix B.

2. Contribution to average gate fidelity from coupler leakage

In Appendix E 1 we demonstrated a measurement of a coupler leakage angle $\theta_c = 0.041 \pm 0.001$, but we did not calculate the expected contribution to fidelity from this process. The exchange between the basis states $|101\rangle$ and




FIG. 13. (a) Circuit and pulse diagram of the CZ gate. A Slepian-shaped flux pulse with duration $\tau_{Z,TC}$ is applied to the tunable coupler, while a flat-top cosine pulse with a rising edge time of 1 ns and a duration $\tau_{Z,Q1}$ is applied to Q1. Q2 remains idle during the gate. Virtual Z rotations (not shown) are applied afterward on each qubit to correct for dynamic phases. (b) Measured Clifford error, CZ gate error, and leakage to the $|2\rangle$ state for Q1 and Q2 as a function of the qubit flux pulse duration $\tau_{Z,Q1}$. (c) Measured Clifford error, CZ gate error, $|2\rangle$ state leakage for both qubits, and SWAP error, plotted as a function of the coupler flux pulse duration $\tau_{Z,TC}$.

$|011\rangle$ takes the system state from a computational to a non-computational state, making the entire process nontrace preserving. However, if the coupler is traced out, it could be understood as an incoherent process in which the state $|11\rangle$ simply decays to $|01\rangle$. Making that approximation, the average gate infidelity contribution may be calculated by assuming the coupler leakage is an error channel $\mathcal{E}_c$ such that $\mathcal{E}_c(|11\rangle) = \cos(\theta_c/2)|11\rangle + \sin(\theta_c/2)|01\rangle$, and $\mathcal{E}_c(|10\rangle) = \cos(\theta_c/2)|10\rangle + \sin(\theta_c/2)|00\rangle$, but it does not affect the other two states. Following the derivation above, it results in the formula identical to Eq. (F4), which for the reported value of θ_c yields $\frac{4}{5} - \frac{4}{5}\cos^4(\theta_c/4) = (1.7 \pm 0.1) \times 10^{-4}$.

This value could be understood as part of the difference between the reported error per ID gate of $\epsilon_{ID} = (4.1 \pm 0.2) \times 10^{-4}$ and the incoherent contribution to CZ error $\epsilon_{CZ}^{\text{incoh}} = 6.2 \times 10^{-4}$ reported in Sec. VB, and represents another way of looking at the same process. Either the qubits hybridize with the coupler, and the excitations

in that basis decay with increased rates, or the qubits leak to the coupler's excited state and then the coupler decays.

3. Leakage to $|20\rangle$ state

For strong couplings and short gate durations, leakage to $|20\rangle$, where the lower qubit's second excited state is populated, may limit the gate fidelity according to simulations. Note that this state's detuning from $|11\rangle$ is approximately double that of the detuning between $|01\rangle$ and $|10\rangle$, while the effective coupling increases by a factor of $\sqrt{2}$.

We could use either PALEA or the Floquet calibration with delays between the CZ gates, analogous to the one used to measure the coupler leakage shown in Fig. 11, to measure the rate of exchange between $|11\rangle$ and $|20\rangle$. Note however the assumptions of the model are only satisfied if the stronger exchange in the $\text{span}(|11\rangle, |02\rangle)$ subspace can be eliminated; otherwise the experiments need to be adjusted to take into account three-level dynamics. One such modification could be to utilize a qutrit dynamical decoupling step [109] on both transmons, with a cyclic $X_3 = X_\pi^{01} X_\pi^{12}$ gate on one and $X_3^\dagger$ on the other, cycling through states $|11\rangle$, $|02\rangle$ and $|20\rangle$ —all of this is however beyond the scope of this work.

We find, using both PALEA and the delay-based amplification, that for gates shorter than 21 ns, the coherent leakage to $|20\rangle$ can be measured to be below 0.03 rad per gate, not limiting the fidelity, while for longer gates it is too small to be measurable. At those durations, the leakage to $|02\rangle$ cannot however be entirely eliminated and is larger than the leakage to $|20\rangle$, violating the model. We conclude that a more appropriate method is to measure the total average leakage rate using the interleaved leakage randomized benchmarking [82,84], see Sec. VB of the main text. We note that the leakage to the lower qubit's $|2\rangle$ in an RB circuit may originate from the direct leakage to $|20\rangle$, but it can also come from more complicated processes such as leakage to $|02\rangle$, excitation to $|12\rangle$ and population transfer to $|21\rangle$.

APPENDIX G: CZ GATE OPTIMIZATION

While calibrating the qubit and coupler flux pulse amplitudes already yields high CZ gate fidelities, further improvement requires careful optimization of the pulse durations. Specifically, both the qubit and coupler flux pulse durations affect leakage to the $|2\rangle$ states and SWAP-type errors and thus must be tuned to minimize total gate error. The CZ gate circuit and pulse timing are illustrated in Fig. 13(a).

1. Optimization of the qubit flux pulse duration

We begin by fixing the coupler flux pulse duration to 21 ns, a value suggested by simulations to yield high-fidelity gates [see Fig. 6(g)]. We then sweep the qubit flux

pulse duration $\tau_{Z,Q1}$, recalibrate the CZ gate at each step using the PALEA protocol and related calibrations, and finally benchmark the gate using iterative IRB. The results are shown in Fig. 13(b).

We extract the $|2\rangle$ state leakage for both qubits using three-state discrimination. For Q_1 , this is defined as the total population in states $|20\rangle$, $|21\rangle$, and $|22\rangle$; analogously for Q_2 . The leakage per CZ gate is estimated by fitting a linear model to the leakage extracted via iterative IRB, similar to how the CZ fidelity per gate is determined in Sec. IV of the main text.

At long qubit pulse durations (28–30 ns), both the Clifford and CZ gate errors rise significantly, despite low leakage. This increase is attributed to coherent errors caused by uncompensated flux pulse distortions in the nanosecond time scales. Since the total gate duration is fixed to 33 ns, longer qubit pulses reduce the buffer time between consecutive gates, allowing pulse distortions to bleed into subsequent operations. These distortions manifest as a large quadratic term in the iterative IRB fit.

For short qubit durations (23–25 ns), the dominant error source becomes leakage to the $|2\rangle$ state of Q_2 . This behavior also possibly stems from nanosecond-scale flux distortions, which affect the optimal shape of the coupler pulse. As the Slepian-shaped coupler pulse assumes a static qubit frequency, distortions in the qubit pulse introduce nonideal Slepian pulse parameters.

The optimal trade-off between leakage and coherent errors is found at a qubit flux pulse duration of approximately 27 ns, where both leakage and overall CZ gate error are minimized. At this point, $|2\rangle$ state leakage in both qubits remains below 10^{-4} .

2. Optimization of the coupler flux pulse duration

In the second sweep, we fix the qubit flux pulse duration to 27 ns and vary the coupler flux pulse duration $\tau_{Z,TC}$. The results of this sweep, shown in Fig. 13(c), exhibit a mostly flat response in Clifford error, CZ gate error, and $|2\rangle$ state leakage.

Because the CZ gate is largely coherence-limited in this regime, the purpose of this sweep is to avoid regions of enhanced leakage—possibly caused by interactions with weakly coupled parasitic modes—which appear at seemingly random durations. One such region is seen near 22.6 ns, where leakage to the $|2\rangle$ state increases sharply for both qubits, reducing overall gate fidelity. The lowest observed CZ error occurs at a coupler flux duration of approximately 22 ns, which we adopt for the final gate parameters reported in the main text (Sec. IV).

SWAP-type errors, as described in Appendix F 1, show an oscillatory behavior as a function of the coupler flux pulse duration [Fig. 13(c)], similar to the oscillations seen in the simulations in Fig. 6(g). Since the errors due to the SWAP-type errors are mostly at levels below 1×10^{-4} , they are

negligible with our setup. For coupler flux pulse durations shorter than 20 ns (not shown in the figure), the SWAP-type errors become the limiting error source for the CZ gate.

APPENDIX H: INCOHERENT ERROR OF THE CZ GATE

To estimate the coherence-limited CZ gate fidelity, we measure the coherence times of the hybridized qubit modes as a function of the coupler frequency and compute an effective average over the flux pulse trajectory, following the approach in Ref. [22].

Figure 14(a) shows the hybridized qubit frequencies and the bare coupler frequency versus the coupler frequency. The vertical dashed lines indicate the operating points corresponding to the idling configuration (ZZ off) and the CZ gate (ZZ on). During the CZ gate, the coupler is flux-pulsed close to the qubit frequencies to enable strong coupling. To mitigate leakage to the coupler mode, a Slepian-shaped flux pulse is used.




FIG. 14. (a) Qubit-like mode frequencies and bare coupler frequency as a function of the coupler frequency. Vertical dashed lines indicate the ZZ off and ZZ on configurations. (b) Energy relaxation time T_1 of both qubits as a function of coupler frequency. (c) Exponential pure dephasing time $T_{\phi,1}$ extracted from Ramsey experiments. Due to Q_1 being dominated by Gaussian dephasing noise, $T_{\phi,1}$ is assumed to be constant at its idling value for Q_1 . (d) Gaussian pure dephasing time $T_{\phi,2}$ extracted from the same experiment. Dashed lines in (b)–(d) represent smoothed interpolations.

We measure the energy relaxation time T_1 by preparing each qubit in its excited state with an X_π pulse, then flux pulsing it (along with the coupler) to the same frequencies used during the CZ gate. The qubit flux pulse has the same envelope as during the CZ gate, while the coupler is pulsed with a flat-top cosine envelope, with a rise time of 20 ns to ensure adiabatic transitions. This measurement is repeated for each qubit individually, and T_1 is extracted via exponential fits to the time-resolved population decay. The results are shown in Fig. 14(b).

To extract pure dephasing times, we perform a similar Ramsey experiment with an initial $\sqrt{X}$ pulse. The resulting signal is fitted with the model

$$p(t) = a(t) \exp\left(-\frac{t}{2T_1} - \frac{t}{T_{\phi,1}} - \left(\frac{t}{T_{\phi,2}}\right)^2\right), \quad (\text{H1})$$

where $T_{\phi,1}$ and $T_{\phi,2}$ are the Gaussian and exponential components of the dephasing time, and T_1 is fixed from prior measurements. The exponential and Gaussian dephasing components are shown in Figs. 14(c) and 14(d).

For Q_1 , which is flux-pulsed away from its sweet spot by approximately 80 MHz, the Gaussian component dominates, making $T_{\phi,1}$ extraction unreliable. Therefore, we assume $T_{\phi,1}^{Q1}$ to be constant and equal to its value of 158 μs at the idling point.

We apply a smoothing interpolation (rather than direct interpolation) to the coherence data to filter out short-term fluctuations, such as those from parasitic two-level systems, which are known to cause temporary dips in coherence.

Using the smoothed data, we compute effective coherence times during the CZ gate by integrating the coherence rate along the flux pulse trajectory. The extracted effective values are

$$\begin{aligned} T_{1,\text{eff}}^{Q1} &= 63 \mu\text{s}, & T_{\phi,1,\text{eff}}^{Q1} &= 157 \mu\text{s}, & T_{\phi,2,\text{eff}}^{Q1} &= 12 \mu\text{s}, \\ T_{1,\text{eff}}^{Q2} &= 83 \mu\text{s}, & T_{\phi,1,\text{eff}}^{Q2} &= 145 \mu\text{s}, & T_{\phi,2,\text{eff}}^{Q2} &= 83 \mu\text{s}. \end{aligned}$$

The total incoherent error contribution to the CZ gate is then calculated using the model in Ref. [65],

$$\begin{aligned} \epsilon_{\text{CZ}}^{\text{incoh}} &= \frac{3}{10} \frac{\tau}{T_{1,\text{eff}}^{Q1}} + \frac{3}{8} \frac{\tau}{T_{\phi,1,\text{eff}}^{Q1}} + \frac{3}{8} \left(\frac{\tau}{T_{\phi,2,\text{eff}}^{Q1}}\right)^2 \\ &+ \frac{1}{2} \frac{\tau}{T_{1,\text{eff}}^{Q2}} + \frac{31}{40} \frac{\tau}{T_{\phi,1,\text{eff}}^{Q2}} + \frac{31}{40} \left(\frac{\tau}{T_{\phi,2,\text{eff}}^{Q2}}\right)^2, \quad (\text{H2}) \end{aligned}$$

where $\tau = 33$ ns is the total CZ gate duration.

Substituting the effective coherence times into Eq. (H2) yields an estimated incoherent error of $\epsilon_{\text{CZ}}^{\text{incoh}} = 6.2 \times 10^{-4}$, in close agreement with the experimentally measured value of $\epsilon_{\text{CZ}} = 6.5 \times 10^{-4}$.

APPENDIX I: MEASURING QNDness

QND measurement refers to a measurement process that preserves the quantum state of the system, meaning the qubit remains in the eigenstate to which it was projected by the readout operation [44]. This property is especially desirable for mid-circuit measurements in quantum error correction, as it removes the need to reset and reinitialize the qubit afterward.

To quantify the QND nature of the readout, we define a ‘‘QNDness’’ metric $\mathcal{Q}$, which reflects the probability that the qubit remains in its original state after measurement, independent of the measurement outcome [44,45],

$$\mathcal{Q} = (\mathcal{Q}_g + \mathcal{Q}_e) / 2, \quad (\text{I1})$$

with

$$\begin{aligned} \mathcal{Q}_g &= P(g, 0 | g) + P(g, 1 | g), \\ \mathcal{Q}_e &= P(e, 0 | e) + P(e, 1 | e), \end{aligned} \quad (\text{I2})$$

and $P(a, b | c)$ denotes the probability that the qubit ends up in state a and is assigned outcome b , given the qubit is initially in state c . For instance, $P(g, 1 | g)$ is the probability of the qubit remaining in $|g\rangle$ but being misassigned to $|1\rangle$.

The aim of this section is to experimentally evaluate the QND character of our readout operations. To this end, we adopt the readout-induced leakage benchmarking (RILB) protocol introduced in Ref. [45], with modifications in the analysis, which allow us to extract not only the leakage rate but also the QNDness $\mathcal{Q}$.

Directly measuring $\mathcal{Q}$ is challenging due to the interplay of various error mechanisms. These include assignment errors such as $P(g, 1 | g)$ and $P(e, 0 | e)$, transition errors such as $P(e | g)$ and $P(g | e)$, and leakage to higher excited states such as $P(f | g)$ and $P(f | e)$. The RILB protocol is designed to disentangle these effects by analyzing the decay of different correlation types between measurements. These decay rates are sensitive to leakage and transition errors but largely independent of assignment errors, making it possible to extract $\mathcal{Q}$ and the leakage.

1. RILB experiment

The circuit diagram for the RILB experiment is shown in Fig. 15(a). The protocol consists of n repeated readout operations, interleaved with randomized bit-flips. The outcome of the n th measurement is denoted $r_n \in \{0, 1\}$, and the applied premeasurement operation is $i_n \in \{0, 1\}$, where 1 indicates a bit-flip (X gate) and 0 indicates the identity operation (I gate).

To extract the local correlations between consecutive measurements, we compute the binary ‘‘flipped-or-not’’




FIG. 15. (a) Circuit diagram of the readout-induced leakage benchmarking experiment. Readout operations are interleaved with randomized bit-flips. Global correlations indicate the correlations between the initial state and expected n th readout outcome, whereas local correlations indicate the expected correlations between two consecutive readout outcomes. (b) and (c) Readout-induced leakage benchmarking experiment performed on Q_1 and Q_2 with (a) a readout operation consisting of only a readout pulse optimized for quantum non-demolition-ness and (b) a readout operation using the shelving technique via an X_{π}^{12} pulse before and after the readout pulse. The experiments are performed and averaged over 500 randomized bit-flip sequences. The global and local correlations are fitted with exponential decay models.

quantity,

$$o_n = r_n \oplus r_{n-1},$$

and compare it to the applied operation to get the local correlation,

$$C_n^{\text{local}} = 1 - o_n \oplus i_n,$$

where $\oplus$ denotes the XOR operation. This definition matches that of Ref. [45] up to a scaling factor.

The original RILB protocol does not allow one to extract the QNDness directly. To solve this issue, we extend the protocol to calculate another quantity from the same experimental data. We define a global correlation, which compares the state after n rounds to the initial state $|0\rangle$. Let

$$h_n = i_0 \oplus i_1 \oplus \dots \oplus i_n,$$

then the global correlation is defined as

$$C_n^{\text{global}} = 1 - r_n \oplus h_n.$$

Following the computation of the correlations, we average them with respect to all possible bit-flip sequences, obtaining the expectation values $\langle C_n^{\text{local}} \rangle$ and $\langle C_n^{\text{global}} \rangle$. While the correlations are straightforward to compute, modeling their decay as a function of n is nontrivial. In the next sections, we derive analytical models for the expectation values of both global and local correlation decay, allowing us to extract Q , leakage, and seepage.

2. Derivation of the global correlation model

We use a simple classical hidden Markov chain model to describe the global correlation, assuming the qubit can be in one of $k + 2$ states: $|g\rangle$, $|e\rangle$, or some leakage state $|l_i\rangle$, where $i \in [0, k]$ and k is the number of leakage states communicating with $|g\rangle$ or $|e\rangle$. This generalization accounts for the fact that more than one leakage state may be populated during the measurement [45]. We note this assumption of Markovianity may be broken for example if the readout resonator becomes increasingly populated with photons during the experiment. We further assume the single-qubit gates used are perfect. We do allow a state-preparation error, but assume the qubit is never prepared in a leakage state.

We start the derivation by defining the bare transition probabilities which do not depend on the assigned outcome,

$$\begin{aligned} p_g &= P(e, 0 | g) + P(e, 1 | g), \\ p_e &= P(g, 0 | e) + P(g, 1 | e), \\ L_{g,i} &= P(l_i, 0 | g) + P(l_i, 1 | g), \\ L_{e,i} &= P(l_i, 0 | e) + P(l_i, 1 | e), \\ S_{g,i} &= P(g, 0 | l_i) + P(g, 1 | l_i), \\ S_{e,i} &= P(e, 0 | l_i) + P(e, 1 | l_i), \end{aligned} \quad (13)$$

as well as assignment probabilities, which we assume only depend on the state of the qubit after the readout,

$$\begin{aligned} e_0 &= P(0 | e), \\ e_1 &= P(1 | g), \\ w_i &= P(0 | l_i). \end{aligned} \quad (14)$$

This is a somewhat risky assumption, but in practice we will see that the w_i values do not enter the final formula. The e_i values can be interpreted as readout SNR errors or overlap errors.

A hidden Markov chain model [110,111] requires an initial state column vector $\mathbf{v}_0$, a transition matrix T , and a row emission probability vector $\mathbf{b}$ used to extract the required observation from the model. Thus the expected value of

the global and local correlations are given by

$$\begin{aligned}\langle C_n^{\text{global}} \rangle &= \mathbf{b}_g T^n \mathbf{v}_0, \\ \langle C_n^{\text{local}} \rangle &= \mathbf{b}_l T^n \mathbf{v}_0.\end{aligned}\quad (15)$$

The aim is to find T and the two emission probability vectors $\mathbf{b}_g$ and $\mathbf{b}_l$ corresponding to global and local correlations. Assuming Markovianity, the global and local correlation expectations must only depend on the expected state after n steps. The emission probability vectors can be understood as linear functions from the set of states to $\mathbb{R}$, which are used to calculate those expectations. Without specifying the vectors $\mathbf{b}_g$ and $\mathbf{b}_l$, the above equation is simply a statement that those functions are in fact linear, which is a simple consequence of the definition of the expected value.

We define a stochastic population vector $\mathbf{v}$ containing the total probabilities of the system being in a wanted, unwanted or any of the leakage states across all possible realizations of the experiment. The wanted state for a specific realization of the experiment is defined with reference to the number of X gates performed thus far, which is $|g\rangle$ for an even number of X gates preceding the last measurement or $|e\rangle$ for an odd number of X gates. Then, v^0 is the total probability that the system is in the wanted computational state and v^1 is the total probability that the system is in the unwanted computational state, that is, whatever v^0 is not. Finally, v^{i+2} is the probability the system is in the i th leakage state, so that $\mathbf{v}_0$ has length $k+2$. The initial population vector is

$$\mathbf{v}_0 = \begin{bmatrix} 1-h \\ h \\ 0 \\ \vdots \end{bmatrix}, \quad (16)$$

where h is the state-preparation error.

Due to a nonstandard definition of the basis for the vector $\mathbf{v}$, we need to take additional care to show that the vector of probabilities after n readouts is indeed equal to $T^n \mathbf{v}_0$ for some square stochastic matrix T . Formally, we will demonstrate it using the principle of the mathematical induction.

The state before any readouts is the initial state $T^0 \mathbf{v}_0 = \mathbf{v}_0$. We need to show that, assuming the state vector after n readouts is $\mathbf{v}_n = T^n \mathbf{v}_0$, the vector describing the state probabilities after $n+1$ readouts is $\mathbf{v}_{n+1} = T^{n+1} \mathbf{v}_0 = T \mathbf{v}_n$.

Suppose that after the n th readout, the wanted state is $|g\rangle$. Then, if there is no X gate between the n th and $(n+1)$ st readout, the vector evolves according to the transition

matrix

$$T_0 = \begin{bmatrix} 1-p_g - \sum_i^k L_{g,i} & p_e & S_{g,0} & \cdots \\ p_g & 1-p_e - \sum_i^k L_{e,i} & S_{e,0} & \cdots \\ L_{g,0} & L_{e,0} & \ddots & \\ \vdots & \vdots & & \ddots \end{bmatrix}, \quad (17)$$

where the unspecified matrix elements are transition probabilities from one leakage state to another, specifically $T_0^{i+2,j+2}$ is the probability of transition from leakage state i to leakage state j . They are positive real numbers such that each column sums to 1.

If there is a flip, the basis of the vector $\mathbf{v}$ is redefined such that its components stay the same after the flip. The basis of T_0 is also transformed, which amounts to permuting the first two rows and the first two columns, resulting in the transition matrix

$$T_1 = \begin{bmatrix} 1-p_e - \sum_i^k L_{e,i} & p_g & S_{e,0} & \cdots \\ p_e & 1-p_g - \sum_i^k L_{g,i} & S_{g,0} & \cdots \\ L_{e,0} & L_{g,0} & \ddots & \\ \vdots & \vdots & & \ddots \end{bmatrix}. \quad (18)$$

The transition matrix reflects probabilities of certain transitions occurring, which depends on the presence of the X gate used to flip the qubit state. Accounting for the X gate being present before the readout with probability of $\frac{1}{2}$, and using the principle of total probability we may write $T^{i,j} = \frac{1}{2} T_0^{i,j} + \frac{1}{2} T_1^{i,j}$ for each matrix element, leading to the effective transition matrix

$$T = \frac{1}{2}(T_0 + T_1) = \begin{bmatrix} 1-p-L & p & S_0 & \cdots \\ p & 1-p-L & S_0 & \cdots \\ L_0 & L_0 & \ddots & \\ \vdots & \vdots & & \ddots \end{bmatrix}, \quad (19)$$

where $p = \frac{1}{2}(p_g + p_e)$, $L = \frac{1}{2}(\sum_i^k L_{g,i} + \sum_i^k L_{e,i})$ are average state switch and leakage probabilities, $L_i = \frac{1}{2}(L_{g,i} + L_{e,i})$ are average leakage rates from computational subspace to leakage state i , and $S_i = \frac{1}{2}(S_{g,i} + S_{e,i})$ are average seepage rates from leakage state i to the computational subspace. The first two columns are identical after the second

row, so are the first two rows after the second column. The QNDness defined earlier can be expressed in terms of these quantities: $Q = 1 - p - L$. The derivation for the case when the wanted state is $|e\rangle$, not $|g\rangle$, follows analogously. Thus, again following the total probability, we can write $\mathbf{v}_{n+1} = T\mathbf{v}_n$, which by mathematical induction implies that $\mathbf{v}_n = T^n\mathbf{v}_0$ for every $n \in \mathbb{N}$.

In addition to a left eigenvector $\mathbf{u}_0 = [1, 1, 1, \dots]$ with eigenvalue 1, shared by all stochastic matrices, the matrix T always has a left eigenvector $\mathbf{u}_1 = [1, -1, 0, \dots]$ with eigenvalue $1 - 2p - L$. This can be verified by direct calculation and is a consequence of the first two rows being identical after the second column. Additionally, $\mathbf{v}_- = \mathbf{u}_1^T$ is a right eigenvector with the same eigenvalue, because the first two columns are identical after the second row.

Consider now the emission probability vector $\mathbf{b}_g$ corresponding to the global correlation. Assuming the wanted state after n readouts is $|g\rangle$, the first component of the vector $\mathbf{v}_n$ contributes $1 - e_1$, second contributes e_0 , and each i th component afterward contributes w_i . If the wanted state is $|e\rangle$, the contributions change to $1 - e_0$, e_1 and $1 - w_i$. Which state is wanted is random with probability $\frac{1}{2}$, so the total emission probability vector is

$$\begin{aligned} \mathbf{b}_g &= \left[1 - \frac{1}{2}(e_0 + e_1), \frac{1}{2}(e_0 + e_1), \frac{1}{2}(w_0 + 1 - w_0), \dots \right] \\ &= \left[1 - e, e, \frac{1}{2}, \dots \right] \\ &= \left(\frac{1}{2} - e \right) \mathbf{u}_1 + \left[\frac{1}{2}, \frac{1}{2}, \dots \right] = \left(\frac{1}{2} - e \right) \mathbf{u}_1 + \frac{1}{2} \mathbf{u}_0, \end{aligned} \quad (\text{I10})$$

where $e = \frac{1}{2}(e_0 + e_1)$. The expected value of the global correlation is then given by

$$\begin{aligned} \langle C_n^{\text{global}} \rangle &= \mathbf{b}_g T^n \mathbf{v}_0 \\ &= \left(\frac{1}{2} - e \right) \mathbf{u}_1 T^n \mathbf{v}_0 + \frac{1}{2} \mathbf{u}_0 T^n \mathbf{v}_0 \\ &= \left(\frac{1}{2} - e \right) (1 - 2p - L)^n \mathbf{u}_1 \cdot \mathbf{v}_0 + \frac{1}{2} \mathbf{u}_0 \cdot \mathbf{v}_0 \\ &= \left(\frac{1}{2} - e \right) (1 - 2h)(1 - 2p - L)^n + \frac{1}{2}, \end{aligned} \quad (\text{I11})$$

where the third equality follows because $\mathbf{u}_0$ and $\mathbf{u}_1$ are left eigenvectors of T . The global correlation is therefore expected to follow a single exponential decay with offset $\frac{1}{2}$, regardless of the number of leakage states.

Note that this model heavily relies on proper randomization of the sequences of flips, and it is crucial to ensure the set of those sequences satisfies the required properties.

3. Derivation of the local correlation model

The local correlation model uses the same transition matrix T , but a different emission probability vector— $\mathbf{b}_l$. To derive it, we will consider a sum of quantities

$$P(r_{n-1} \oplus r_n \oplus i_{n+1} = 0, s_{n+1} | s_n) P(s_n) \quad (\text{I12})$$

spanning all possible cases, where s_n is the real qubit state after the n th measurement. Note that only $P(s_n)$ depends on n , according to the assumption of Markovianity.

Specifically, $P(s_n = g) = P(s_n = e) = \frac{1}{2}(v_n^0 + v_n^1)$, because the states $|g\rangle$ or $|e\rangle$ may be wanted or unwanted with probability $\frac{1}{2}$, while $P(s_n = l_i) = v_n^{i+2}$. Consequently, the two first components of the emission probability vector $\mathbf{b}_l$ have to be equal, because we must have

$$b_l^0 v_n^0 + b_l^1 v_n^1 = A_0 \frac{1}{2}(v_n^0 + v_n^1) + A_1 \frac{1}{2}(v_n^0 + v_n^1), \quad (\text{I13})$$

where A_0 and A_1 are independent of n . Unfortunately, their values depend in general on all the introduced transition and assignment probabilities, because the presence of the flip does affect the transition probability and we may no longer rely on the average leakage rates—we need to use the matrix T_0 directly. We may however write

$$\begin{aligned} A &= \frac{1}{2}(A_1 + A_0) \\ &= \frac{1}{2} \sum_{s_n \in \{g, e\}} \sum_{s_{n+1}} P(r_{n-1} \oplus r_n \oplus i_{n+1} = 0, s_{n+1} | s_n), \end{aligned} \quad (\text{I14})$$

where A in general depends on all other model parameters.

We will now show that each leakage state contributes half its population to the local correlation. Consider an i th leakage state l_i . The transition probability to another state is given by T_0^{i+2j} . The presence of the flip does not change the state, so it does not affect the transition probability. Without the flip, the bit assignments for (r_{n-1}, r_n) as 00 and 11 contribute to the local correlation, but with the flip, the assignments 10 and 01 contribute to it. These are all possible assignments, and the probability of obtaining them is independent of the presence of the flip, thus for leakage states

$$\begin{aligned} P(r_{n-1} \oplus r_n \oplus i_{n+1} = 0, s_{n+1} = j | s_n = l_i) \\ &= \frac{1}{2} P(s_{n+1} = j | s_n = l_i) \\ &= \frac{1}{2} T_0^{i+2j}, \end{aligned} \quad (\text{I15})$$

where the factor of $\frac{1}{2}$ is the probability of the flip. The total expected contribution factor to the local correlation

if $s_n = l_i$ is

$$\sum_j^{k+2} \frac{1}{2} T_0^{i+2,j} = \frac{1}{2}, \quad (\text{I16})$$

because columns of T_0 sum to 1.

We may finally define the emission probability vector for local correlation

$$\mathbf{b}_l = [A, A, \frac{1}{2}, \dots] = \left(A - \frac{1}{2}\right) [1, 1, 0, \dots] + \frac{1}{2} \mathbf{u}_0. \quad (\text{I17})$$

As shown before, $\mathbf{v}_- = [1, -1, 0, \dots]^T$ is a right eigenvector of T with eigenvalue $(1 - 2p - L)$. Define additionally $\mathbf{v}_+ = [1, 1, 0, \dots]^T$, which allows us to write $\mathbf{v}_0 = \frac{1}{2} \mathbf{v}_+ + (\frac{1}{2} - h) \mathbf{v}_-$. Thus we obtain for the expectation value of the local correlation

$$\begin{aligned} \langle \mathcal{C}_n^{\text{local}} \rangle &= \mathbf{b}_l T^n \mathbf{v}_0 \\ &= \frac{1}{2} \mathbf{b}_l T^n \mathbf{v}_+ + \left(\frac{1}{2} - h\right) \mathbf{b}_l T^n \mathbf{v}_- \\ &= \frac{1}{2} \mathbf{b}_l T^n \mathbf{v}_+ + \left(\frac{1}{2} - h\right) (1 - 2p - L)^n \mathbf{b}_l \cdot \mathbf{v}_- \\ &= \frac{1}{2} \mathbf{b}_l T^n \mathbf{v}_+, \end{aligned} \quad (\text{I18})$$

because the vector $\mathbf{b}_l$ is orthogonal to $\mathbf{v}_- = \mathbf{u}_1^T$ associated with eigenvalue $(1 - 2p - L)$. Thus we see that the state preparation error does not affect the local correlation. However, $\mathbf{b}_l$ is not guaranteed to be orthogonal to any other eigenvector. In almost all cases, the matrix T will be diagonalizable and have k more eigenvectors and eigenvalues, each of which may be dependent on any remaining leakage and seepage rates, including the transition probabilities from one leakage state to another. Thus, we expect the local correlation to be a decaying function of n , with k exponential contributions.

However, typically only one leakage state is relevant. In such a case, the matrix T becomes

$$T = \begin{bmatrix} 1 - p - L & p & \frac{S}{2} \\ p & 1 - p - L & \frac{S}{2} \\ L & L & 1 - S \end{bmatrix}, \quad (\text{I19})$$

where S is now total seepage rate and gains a left eigenvector $\mathbf{u}_2 = [-(L/S), -(L/S), 1]$ with eigenvalue $1 - L - S$. Following Eq. (I18) we obtain

$$\langle \mathcal{C}_m^{\text{local}} \rangle = \frac{L(A - \frac{1}{2})(1 - L - S)^m + AS + \frac{1}{2}L}{L + S}, \quad (\text{I20})$$

thus reconstructing the model from Ref. [45]. We may then estimate the leakage rate L and the seepage rate S by fitting

the amplitude, offset and decay rate of the exponentially decaying data.

4. Estimating the QNDness

In case of one relevant leakage state, the local correlation allows measurement of the leakage rate. This rate can then be used to calculate the QNDness $\mathcal{Q}$ by combining the result with the global correlation decay constant estimated using Eq. (I11).

In case of more than one relevant leakage state, the total leakage rate from the computational subspace states may no longer be estimated as accurately without additional assumptions specifying certain transition rates to be 0, because the local correlation is a sum of multiple exponential decays, whose rates are eigenvalues of T with generally complicated analytical expressions. In other words, the model is no longer fully identifiable [112]—a model with k leakage states can have up to $k^2 + k + 1$ parameters, including A and all the independent rates included in the matrix T . It then generates data following a sum of k exponential decays, which can be fitted with $2k + 1$ parameters: a global offset, and one scaling factor and decay rate per included exponential contribution. Thus, the model becomes nonidentifiable once $k^2 + k + 1 > 2k + 1$ or already for $k \geq 2$. It is unknown to us whether the total leakage rate from the computational subspace is an identifiable parameter without additional assumptions.

In such a case, one may bound the QNDness $\mathcal{Q}$ using just the global correlation rate $2p + L$,

$$1 - 2p - L < \mathcal{Q} < 1 - p - \frac{L}{2}, \quad (\text{I21})$$

where the lower bound may be taken if leakage is the limiting factor.

5. Imperfect single-qubit gates

For the purpose of the above derivation, we have assumed perfect single-qubit gates. Given low errors demonstrated in the Sec. VA of the main text, this is a reasonable assumption, and here we demonstrate how to further test its validity. Suppose our single-qubit X rotation is not perfect and the entirety of its infidelity is due to a depolarizing error λ . In such a case we would have $\lambda = \epsilon_X \approx 2\epsilon_{\sqrt{X}}$ because our X gates are implemented using two $\sqrt{X}$ gates.

Since during the RILB experiment any X rotation is immediately followed by a Z basis measurement, we can assume its action is effectively given by a classical

stochastic matrix E_λ ,

$$E_\lambda = \begin{bmatrix} 1-\lambda & \lambda & 0 & \cdots \\ \lambda & 1-\lambda & 0 & \cdots \\ 0 & 0 & 1 & \\ \vdots & \vdots & & \ddots \end{bmatrix}. \quad (I22)$$

This matrix is applied to the matrix T when the X rotation is present, which occurs with probability $\frac{1}{2}$. The wanted state is also $|g\rangle$ or $|e\rangle$ with probability $\frac{1}{2}$, and thus the new effective transition matrix becomes

$$T' = \frac{1}{4}(T_0 + T_1 E_\lambda + T_0 E_\lambda + T_1) = T \left(\frac{1}{2}I + \frac{1}{2}E_\lambda \right), \quad (I23)$$

which is a matrix structurally identical to T , with the substitution

$$p \rightarrow p + \frac{\lambda}{2} - \frac{L\lambda}{2} - p\lambda. \quad (I24)$$

Thus, by assuming perfect X rotations we may overestimate the real p —and so underestimate the QNDness—by an additive factor up to $\lambda/2 \approx \epsilon_{\sqrt{X}}$, while the measured estimate for the leakage rate is unaffected.

Note that the models for global and local correlations rely on symmetry introduced by random flips. If the single-qubit gate's imperfections are of a different nature, it might break that symmetry, violating the model.

We include this correction in the reported uncertainty in the QNDness.

6. QND-optimized readout versus shelved readout

Figure 15(b) shows experimental data from the RILB experiment using a QND-optimized readout configuration. This readout consists of a single rectangular pulse with a duration of 240 ns and a wait time of 350 ns after the end of the readout pulse. This delay was not optimized and reflects a limitation of the current control software, resulting in a total readout duration of 590 ns. We extract the local and global correlations for both Q_1 and Q_2 as described in Appendix I 1, and fit the data using the models from Eqs. (I11) and (I20). Then, we calculate the QNDness Q by combining the results. The extracted parameters are as follows:

	Q_1	Q_2
$1 - Q$	$(6.7 \pm 0.2) \times 10^{-3}$	$(5.9 \pm 0.2) \times 10^{-3}$
L	$(2.1 \pm 0.2) \times 10^{-5}$	$(6.6 \times 0.3) \times 10^{-4}$

The large relative uncertainties in the extracted leakage rates are due to the flatness of the local correlation curves, indicating negligible leakage under these conditions.

In contrast, Fig. 15(c) shows RILB data obtained using a shelved readout scheme, where the qubit is excited to the second excited state before readout (parameters as used in the main text). A second X_π^{12} pulse is applied after 350 ns wait time, the resulting in a total shelved readout duration of 670 ns. Due to larger overall leakage, for local correlation we select either the single exponential model from Eq. (I20) or a phenomenological two-exponential model by minimizing the Bayesian Information Criterion [113].

In case of both qubits, the single exponential model explains the data better, implying the presence of one dominating leakage state. The extracted parameters are

	Q_1	Q_2
$1 - Q$	$(1.06 \pm 0.02) \times 10^{-2}$	$(1.11 \pm 0.02) \times 10^{-2}$
L	$(1.005 \pm 0.006) \times 10^{-2}$	$(1.05 \pm 0.005) \times 10^{-2}$

Comparing the two configurations highlights the main drawback of shelved readout: a roughly two orders of magnitude increase in the leakage rate. This degradation is attributed to the decay of the $|f\rangle$ state to $|e\rangle$ during the readout pulse, followed by the subsequent X_π^{12} pulse re-exciting $|e\rangle$ to $|f\rangle$, resulting in persistent population in the $|f\rangle$ state after measurement. On the other hand, the switching rate p does not contribute to QNDness, implying that the shelved readout is working as designed by trading the higher leakage for the substantial decrease in the switching rate p .

The comparison also yields an important insight for the RILB method: even though the global correlation exponential decays in Figs. 15(b) and 15(c) are similar, the QNDness of the shelved readout is almost entirely limited by leakage, whereas the QNDness of the QND-optimized readout is almost entirely limited by the switching rate p between the two computational states. As a result, the QND error differs by a factor of 2. Ultimately, both the switching rate and leakage contribute to QNDness and are important metrics, and we demonstrate that it is possible to suppress one in favor of the other. The final choice of which readout is better depends on the application; for example for quantum error correction, reducing leakage is typically much more important, because the leakage error cannot be corrected by the codes without the use of additional resources such as the leakage reduction unit [114].

-
- [1] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [2] Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C. White, Josh

- Mutus, Austin G. Fowler, Brooks Campbell *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature* **508**, 500 (2014).
- [3] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2025).
- [4] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [5] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G.S.L. Brandao, David A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [6] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [7] Youngseok Kim, Andrew Eddins, Sajant Anand, Ken Xuan Wei, Ewout van den Berg, Sami Rosenblatt, Hasan Nayfeh, Yantao Wu, Michael Zaletel, Kristan Temme, and Abhinav Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
- [8] Zhiyuan Li, Pei Liu, Peng Zhao, Zhenyu Mi, Huikai Xu, Xuehui Liang, Tang Su, Weijie Sun, Guangming Xue, Jing-Ning Zhang, Weiyang Liu, Yirong Jin, and Haifeng Yu, Error per single-qubit gate below 10^{-4} in a superconducting qubit, *npj Quantum Inf.* **9**, 111 (2023).
- [9] Eric Hyppä, Antti Vepsäläinen, Miha Papič, Chun Fai Chan, Sinan Inel, Alessandro Landra, Wei Liu, Jürgen Luus, Fabian Marxer, Caspar Ockeloen-Korppi, Sebastian Orbell, Brian Tarasinski, and Johannes Heinsoo, Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes, *PRX Quantum* **5**, 030353 (2024).
- [10] David A. Rower, Leon Ding, Helin Zhang, Max Hays, Junyoung An, Patrick M. Harrington, Ilan T. Rosen, Jeffrey M. Gertler, Thomas M. Hazard, Bethany M. Niedzielski, Mollie E. Schwartz, Simon Gustavsson, Kyle Serniak, Jeffrey A. Grover, and William D. Oliver, Suppressing counter-rotating errors for fast single-qubit gates with fluxonium, *PRX Quantum* **5**, 040342 (2024).
- [11] Leon Ding, Max Hays, Youngkyu Sung, Bharath Kannan, Junyoung An, Agustin Di Paolo, Amir H. Karamlou, Thomas M. Hazard, Kate Azar, David K. Kim, Bethany M. Niedzielski, Alexander Melville, Mollie E. Schwartz, Jonilyn L. Yoder, Terry P. Orlando, Simon Gustavsson, Jeffrey A. Grover, Kyle Serniak, and William D. Oliver, High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler, *Phys. Rev. X* **13**, 031035 (2023).
- [12] Rui Li, Kentaro Kubo, Yinghao Ho, Zhiguang Yan, Yasunobu Nakamura, and Hayato Goto, Realization of high-fidelity CZ gate based on a double-transmon coupler, *Phys. Rev. X* **14**, 041050 (2024).
- [13] Wei-Ju Lin, Hyunheung Cho, Yinqi Chen, Maxim G. Vavilov, Chen Wang, and Vladimir E. Manucharyan, 24 days-stable CNOT gate on fluxonium qubits with over 99.9% fidelity, *PRX Quantum* **6**, 010349 (2025).
- [14] Helin Zhang, Chunyang Ding, D. K. Weiss, Ziwen Huang, Yuwei Ma, Charles Guinn, Sara Sussman, Sai Pavan Chitta, Danyang Chen, Andrew A. Houck, Jens Koch, and David I. Schuster, Tunable inductive coupler for high-fidelity gates between fluxonium qubits, *PRX Quantum* **5**, 020326 (2024).
- [15] X. Li, T. Cai, H. Yan, Z. Wang, X. Pan, Y. Ma, W. Cai, J. Han, Z. Hua, X. Han *et al.*, Tunable coupler for realizing a controlled-phase gate with dynamically decoupled regime in a superconducting circuit, *Phys. Rev. Appl.* **14**, 024070 (2020).
- [16] Yuan Xu, Ji Chu, Jiahao Yuan, Jiawei Qiu, Yuxuan Zhou, Libo Zhang, Xinsheng Tan, Yang Yu, Song Liu, Jian Li *et al.*, High-fidelity, high-scalability two-qubit gate scheme for superconducting qubits, *Phys. Rev. Lett.* **125**, 240503 (2020).
- [17] Michele C. Collodo, Johannes Herrmann, Nathan Lacroix, Christian Kraglund Andersen, Ants Remm, Stefania Lazar, Jean-Claude Besse, Theo Walter, Andreas Wallraff, and Christopher Eichler, Implementation of conditional phase gates based on tunable zz interactions, *Phys. Rev. Lett.* **125**, 240502 (2020).
- [18] Youngkyu Sung, Leon Ding, Jochen Braumüller, Antti Vepsäläinen, Bharath Kannan, Morten Kjaergaard, Ami Greene, Gabriel O. Samach, Chris McNally, David Kim *et al.*, Realization of high-fidelity CZ and ZZ-free iSWAP gates with a tunable coupler, *Phys. Rev. X* **11**, 021058 (2021).
- [19] Yangsen Ye *et al.*, Realization of high-fidelity controlled-phase gates in extensible superconducting qubits design with a tunable coupler, *Chin. Phys. Lett.* **38**, 100301 (2021).
- [20] Eyob A. Sete, Angela Q. Chen, Riccardo Manenti, Shobhan Kulshreshtha, and Stefano Poletto, Floating tunable coupler for scalable quantum computing architectures, *Phys. Rev. Appl.* **15**, 064063 (2021).
- [21] J. Stehlik, D. M. Zajac, D. L. Underwood, T. Phung, J. Blair, S. Carnevale, D. Klaus, G. A. Keefe, A. Carniol, M. Kumph *et al.*, Tunable coupling architecture for fixed-frequency transmon superconducting qubits, *Phys. Rev. Lett.* **127**, 080505 (2021).
- [22] Fabian Marxer *et al.*, Long-distance transmon coupler with CZ-gate fidelity above 99.8%, *PRX Quantum* **4**, 010314 (2023).
- [23] Fei Yan, Philip Krantz, Youngkyu Sung, Morten Kjaergaard, Daniel L. Campbell, Terry P. Orlando, Simon Gustavsson, and William D. Oliver, Tunable coupling scheme for implementing high-fidelity two-qubit gates, *Phys. Rev. Appl.* **10**, 054062 (2018).
- [24] Evan Jeffrey, Daniel Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and John M. Martinis, Fast accurate state measurement with superconducting qubits, *Phys. Rev. Lett.* **112**, 190504 (2014).
- [25] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočník, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger,

- C. Eichler, and A. Wallraff, Rapid, high-fidelity, single-shot dispersive readout of superconducting qubits, *Phys. Rev. Appl.* **7**, 054020 (2017).
- [26] Liangyu Chen, Hang-Xi Li, Yong Lu, Christopher W. Warren, Christian J. Križan, Sandoko Kosen, Marcus Rommel, Shah Nawaz Ahmed, Amr Osman, Janka Biznárová *et al.*, Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier, *npj Quantum Inf.* **9**, 26 (2023).
- [27] François Swiadek, Ross Shillito, Paul Magnard, Ants Remm, Christoph Hellings, Nathan Lacroix, Quentin Ficheux, Dante Colao Zanuz, Graham J. Norris, Alexandre Blais, Sebastian Krinner, and Andreas Wallraff, Enhancing dispersive readout of superconducting qubits through dynamic control of the dispersive shift: Experiment and theory, *PRX Quantum* **5**, 040326 (2024).
- [28] Peter A. Spring, Luka Milanovic, Yoshiki Sunada, Shiyu Wang, Arjan F. van Loo, Shuhei Tamate, and Yasunobu Nakamura, Fast multiplexed superconducting-qubit readout with intrinsic Purcell filtering using a multiconductor transmission line, *PRX Quantum* **6**, 020345 (2025).
- [29] Pavel D. Kurilovich, Thomas Connolly, Charlotte G. L. Böttcher, Daniel K. Weiss, Sumeru Hazra, Vidul R. Joshi, Andy Z. Ding, Heekun Nho, Spencer Diamond, Vladislav D. Kurilovich, Wei Dai, Valla Fatemi, Luigi Frunzio, Leonid I. Glazman, and Michel H. Devoret, High-frequency readout free from transmon multi-excitation resonances, [arXiv:2501.09161](https://arxiv.org/abs/2501.09161).
- [30] Google Quantum AI, Suppressing quantum errors by scaling a surface code logical qubit, *Nature* **614**, 676 (2023).
- [31] Youwei Zhao, Yangsen Ye, He-Liang Huang, Yiming Zhang, Dachao Wu, Huijie Guan, Qingling Zhu, Zuolin Wei, Tan He, Sirui Cao *et al.*, Realization of an error-correcting surface code with superconducting qubits, *Phys. Rev. Lett.* **129**, 030501 (2022).
- [32] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, Initialization by measurement of a superconducting quantum bit circuit, *Phys. Rev. Lett.* **109**, 050507 (2012).
- [33] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo, Feedback control of a solid-state qubit using high-fidelity projective measurement, *Phys. Rev. Lett.* **109**, 240502 (2012).
- [34] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard, Persistent control of a superconducting qubit by stroboscopic measurement feedback, *Phys. Rev. X* **3**, 021008 (2013).
- [35] Salvatore S. Elder, Christopher S. Wang, Philip Reinhold, Connor T. Hann, Kevin S. Chou, Brian J. Lester, Serge Rosenblum, Luigi Frunzio, Liang Jiang, and Robert J. Schoelkopf, High-fidelity measurement of qubits encoded in multilevel superconducting circuits, *Phys. Rev. X* **10**, 011001 (2020).
- [36] François Mallet, Florian R. Ong, Agustin Palacios-Laloy, François Nguyen, Patrice Bertet, Denis Vion, and Daniel Esteve, Single-shot qubit readout in circuit quantum electrodynamics, *Nat. Phys.* **5**, 791 (2009).
- [37] Can Wang, Feng-Ming Liu, He Chen, Yi-Fei Du, Chong Ying, Jian-Wen Wang, Yong-Heng Huo, Cheng-Zhi Peng, Xiaobo Zhu, Ming-Cheng Chen, Chao-Yang Lu, and Jian-Wei Pan, Longitudinal and nonlinear coupling for high-fidelity readout of a superconducting qubit, *Phys. Rev. Lett.* **135**, 060803 (2025).
- [38] B. Yurke, L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. A. Whittaker, Observation of parametric amplification and deamplification in a Josephson parametric amplifier, *Phys. Rev. A* **39**, 2519 (1989).
- [39] B. Yurke, M. L. Roukes, R. Movshovich, and A. N. Pargellis, A low-noise series-array Josephson junction parametric amplifier, *Appl. Phys. Lett.* **69**, 3078 (1996).
- [40] Josh Y. Mutus, Theodore C. White, Rami Barends, Yu Chen, Zijun Chen, Ben Chiaro, Andrew Dunsworth, Evan Jeffrey, Julian Kelly, Anthony Megrant, Charles Neill, Peter J. J. O'Malley, Pedram Roushan, Daniel Sank, Amit Vainsencher, James Wenner, Karl M. Sundqvist, Andrew N. Cleland, and John M. Martinis, Strong environmental coupling in a Josephson parametric amplifier, *Appl. Phys. Lett.* **104**, 263513 (2014).
- [41] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret, Optimizing the nonlinearity and dissipation of a SNAIL parametric amplifier for dynamic range, *Phys. Rev. Appl.* **10**, 054020 (2018).
- [42] Johannes Heinsoo, Christian Kraglund Andersen, Ants Remm, Sebastian Krinner, Theodore Walter, Yves Salathé, Simone Gasparinetti, Jean-Claude Besse, Anton Potočník, Andreas Wallraff, and Christopher Eichler, Rapid high-fidelity multiplexed readout of superconducting qubits, *Phys. Rev. Appl.* **10**, 034040 (2018).
- [43] Jami Rönkkö, Olli Ahonen, Ville Bergholm, Alessio Calzona, Attila Geresdi, Hermann Heimonen, Johannes Heinsoo, Vladimir Milchakov, Stefan Pogorzalek, Matthew Sarsby, Mykhailo Savitskyi, Stefan Seegerer, Fedor Šimkovic, P. V. Sriluckshmy, Panu T. Vesanen, and Mikio Nakahara, On-premises superconducting quantum computer for education and research, *EPJ Quantum Technol.* **11**, 32 (2024).
- [44] A. Lupaşcu, S. Saito, T. Picot, P. C. de Groot, C. J. P. M. Harmans, and J. E. Mooij, Quantum non-demolition measurement of a superconducting two-level system, *Nat. Phys.* **3**, 119 (2007).
- [45] S. Hazra, W. Dai, T. Connolly, P. D. Kurilovich, Z. Wang, L. Frunzio, and M. H. Devoret, Benchmarking the readout of a superconducting qubit for repeated measurements, *Phys. Rev. Lett.* **134**, 100601 (2025).
- [46] Sebastian Krinner, Nathan Lacroix, Ants Remm, Agustin Di Paolo, Elie Genois, Catherine Leroux, Christoph Hellings, Stefania Lazar, François Swiadek, Johannes Herrmann, Graham J. Norris, Christian Kraglund Andersen, Markus Müller, Alexandre Blais, Christopher Eichler, and Andreas Wallraff, Realizing repeated quantum error correction in a distance-three surface code, *Nature* **605**, 669 (2022).
- [47] Leonid Abdurakhimov, Janos Adam, Hasnain Ahmad, Olli Ahonen, Manuel Algaba, Guillermo Alonso, Ville Bergholm, Rohit Beriwal, Matthias Beuerle, Clinton Bockstiegel *et al.*, Technology and performance benchmarks of IQM's 20-qubit quantum computer, [arXiv:2408.12433](https://arxiv.org/abs/2408.12433).

- [48] Mark Field, Angela Q. Chen, Ben Scharmann, Eyob A. Sete, Feyza Oruc, Kim Vu, Valentin Kosenko, Joshua Y. Mutus, Stefano Poletto, and Andrew Bestwick, Modular superconducting-qubit architecture with a multichip tunable coupler, *Phys. Rev. Appl.* **21**, 054063 (2024).
- [49] S. Vallés-Sanclemente, T. H. F. Vroomans, T. R. van Abswoude, F. Brulleman, T. Stavenga, S. L. M. van der Meer, Y. Xin, A. Lawrence, V. Singh, M. A. Rol *et al.*, Optimizing the frequency positioning of tunable couplers in a circuit QED processor to mitigate spectator effects on quantum operations, [arXiv:2503.13225](#).
- [50] Zlatko K. Mineev, Zaki Leghtas, Shantanu O. Mundhada, Lysander Christakis, Ioan M. Pop, and Michel H. Devoret, Energy-participation quantization of Josephson circuits, *npj Quantum Inf.* **7**, 131 (2021).
- [51] David C. McKay, Christopher J. Wood, Sarah Sheldon, Jerry M. Chow, and Jay M. Gambetta, Efficient Z gates for quantum computing, *Phys. Rev. A* **96**, 022330 (2017).
- [52] Mikko Tuokkola, Yoshiki Sunada, Heidi Kivijärvi, Jonatan Albanese, Leif Grönberg, Jukka-Pekka Kaikkonen, Visa Vesterinen, Joonas Govenius, and Mikko Möttönen, Methods to achieve near-millisecond energy relaxation and dephasing times for a superconducting transmon qubit, *Nat. Commun.* **16**, 5421 (2025).
- [53] Matthew P. Bland, Faranak Bahrami, Jeronimo G. C. Martinez, Paal H. Prestegard, Basil M. Smitham, Atharv Joshi, Elizabeth Hedrick, Shashwat Kumar, Ambrose Yang, Alexander C. Pakpour-Tabrizi, Apoorv Jindal, Ray D. Chang, Guangming Cheng, Nan Yao, Robert J. Cava, Nathalie P. de Leon, and Andrew A. Houck, Millisecond lifetimes and coherence times in 2D transmon qubits, *Nature* **647**, 343 (2025).
- [54] Alexander P. M. Place *et al.*, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nat. Commun.* **12**, 1779 (2021).
- [55] Peter A. Spring, Shuxiang Cao, Takahiro Tsunoda, Giulio Campanaro, Simone Fasciati, James Wills, Mustafa Bakr, Vivek Chidambaram, Boris Shteynas, Lewis Carpenter, Paul Gow, James Gates, Brian Vlastakis, and Peter J. Leek, High coherence and low cross-talk in a tileable 3D integrated superconducting circuit architecture, *Sci. Adv.* **8**, eabl6698 (2022).
- [56] Shingo Kono, Jiahe Pan, Mahdi Chegnizadeh, Xuxin Wang, Amir Youssefi, Marco Scigliuzzo, and Tobias J. Kippenberg, Mechanically induced correlated errors on superconducting qubits with relaxation times exceeding 0.4 ms, *Nat. Commun.* **15**, 3950 (2024).
- [57] J. Van Damme, S. Massar, R. Acharya, T. Ivanov, D. Perez Lozano, Y. Canvel, M. Demarets, D. Vangoidsenhoven, Y. Hermans, J. G. Lai, A. M. Vadiraj, M. Mongillo, D. Wan, J. De Boeck, A. Potočník, and K. De Greve, Advanced CMOS manufacturing of superconducting qubits on 300 mm wafers, *Nature* **634**, 74 (2024).
- [58] John M. Martinis and Michael R. Geller, Fast adiabatic qubit gates using only σ_z control, *Phys. Rev. A* **90**, 022307 (2014).
- [59] Michiel A. Rol, Livio Ciorciaro, Filip K. Malinowski, Brian M. Tarasinski, Ramiro E. Sagastizabal, Cornelis Christiaan Bultink, Yves Salathe, Niels Haandbæk, Jan Sedivy, and Leonardo DiCarlo, Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor, *Appl. Phys. Lett.* **116**, 054001 (2020).
- [60] Shelby Kimmel, Guang Hao Low, and Theodore J. Yoder, Robust calibration of a universal single-qubit gate set via robust phase estimation, *Phys. Rev. A* **92**, 062315 (2015).
- [61] Kenneth Rudinger, Shelby Kimmel, Daniel Lobser, and Peter Maunz, Experimental demonstration of a cheap and accurate phase estimation, *Phys. Rev. Lett.* **118**, 190502 (2017).
- [62] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Andreas Bengtsson, Sergio Boixo, Michael Broughton, Bob B. Buckley *et al.*, Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, [arXiv:2010.07965](#).
- [63] Colin Scarato, Kilian Hanke, Ants Remm, Stefania Lazăr, Nathan Lacroix, Dante Colao Zanuz, Alexander Flasby, Andreas Wallraff, and Christoph Hellings, Realizing a continuous set of two-qubit gates parameterized by an idle time, *PRX Quantum* **6**, 040317 (2025).
- [64] Jonathan A. Gross, Élie Genois, Dripto M. Debroy, Yaxing Zhang, Wojciech Mruczkiewicz, Ze-Pei Cui, and Zhang Jiang, Characterizing coherent errors using matrix-element amplification, *npj Quantum Inf.* **10**, 123 (2024).
- [65] Tahereh Abad, Yoni Schattner, Anton Frisk Kockum, and Göran Johansson, Impact of decoherence on the fidelity of quantum gates leaving the computational subspace, *Quantum* **9**, 1684 (2025).
- [66] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University, 2010).
- [67] Christiane P. Koch, Ugo Boscarino, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm, Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in Europe, *EPJ Quantum Technol.* **9**, 19 (2022).
- [68] Donald R. Jones, Matthias Schonlau, and William J. Welch, Efficient global optimization of expensive black-box functions, *J. Global Optim.* **13**, 455 (1998).
- [69] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka, Batched large-scale Bayesian optimization in high-dimensional spaces, *Proc. Int. Conf. AISTATS* **84**, 745 (2018).
- [70] F. Setiawan, Peter Groszkowski, and Aashish A. Clerk, Fast and robust geometric two-qubit gates for superconducting qubits and beyond, *Phys. Rev. Appl.* **19**, 034071 (2023).
- [71] Luca Petit, Maximilian Russ, Gertjan H. G. J. Eenink, William I. L. Lawrie, James S. Clarke, Lieven M. K. Vandersypen, and Menno Veldhorst, Design and integration of single-qubit rotations and two-qubit gates in silicon above one Kelvin, *Commun. Mater.* **3**, 82 (2022).
- [72] Ming Ni, Rong-Long Ma, Zhen-Zhen Kong, Ning Chu, Sheng-Kai Zhu, Chu Wang, Ao-Ran Li, Wei-Zhu Liao, Gang Cao, Gui-Lei Wang, Guang-Can Guo, Xuedong Hu, Hai-Ou Li, and Guo-Ping Guo, Diverse set of two-qubit gates for spin qubits in semiconductor quantum dots, *Phys. Rev. Appl.* **23**, 024065 (2025).

- [73] Sarah Sheldon, Lev S. Bishop, Easwar Magesan, Stefan Filipp, Jerry M. Chow, and Jay M. Gambetta, Characterizing errors on qubit operations via iterative randomized benchmarking, *Phys. Rev. A* **93**, 012301 (2016).
- [74] Antoine Brilliant, Peter Groszkowski, Alireza Seif, Jens Koch, and Aashish A. Clerk, Randomized benchmarking with non-Markovian noise and realistic finite-time gates, *Phys. Rev. Lett.* **135**, 070601 (2025).
- [75] Haonan Xiong, Jiahui Wang, Juan Song, Jize Yang, Zenghui Bao, Yan Li, Zhen-Yu Mi, Hongyi Zhang, Hai-Feng Yu, Yipu Song *et al.*, Scalable low-overhead superconducting non-local coupler with exponentially enhanced connectivity, [arXiv:2502.18902](https://arxiv.org/abs/2502.18902).
- [76] Michael Renger, Jeroen Verjauw, Nicola Wurz, Amin Hosseinkhani, Caspar Ockeloen-Korppi, Wei Liu, Aniket Rath, Manish J. Thapa, Florian Vigneau, Elisabeth Wybo *et al.*, A superconducting qubit-resonator quantum processor with effective all-to-all connectivity, [arXiv:2503.10903](https://arxiv.org/abs/2503.10903).
- [77] Akel Hashim, Ravi K. Naik, Alexis Morvan, Jean-Loup Ville, Bradley Mitchell, John Mark Kreikebaum, Marc Davis, Ethan Smith, Costin Iancu, Kevin P. O'Brien, Ian Hincks, Joel J. Wallman, Joseph Emerson, and Irfan Siddiqi, Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor, *Phys. Rev. X* **11**, 041039 (2021).
- [78] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly nonlinear qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [79] Zijun Chen *et al.*, Measuring and suppressing quantum state leakage in a superconducting qubit, *Phys. Rev. Lett.* **116**, 020501 (2016).
- [80] Joel Wallman, Chris Granade, Robin Harper, and Steven T. Flammia, Estimating the coherence of noise, *New J. Phys.* **17**, 113020 (2015).
- [81] Guanru Feng, Joel J. Wallman, Brandon Buonacorsi, Franklin H. Cho, Daniel K. Park, Tao Xin, Dawei Lu, Jonathan Baugh, and Raymond Laflamme, Estimating the coherence of noise in quantum control of a solid-state qubit, *Phys. Rev. Lett.* **117**, 260501 (2016).
- [82] Christopher J. Wood and Jay M. Gambetta, Quantification and characterization of leakage errors, *Phys. Rev. A* **97**, 032306 (2018).
- [83] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 14 (2021).
- [84] M. A. Rol, F. Battistel, F. K. Malinowski, C. C. Bultink, B. M. Tarasinski, R. Vollmer, N. Haider, N. Muthusubramanian, A. Bruno, B. M. Terhal, and L. DiCarlo, Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits, *Phys. Rev. Lett.* **123**, 120502 (2019).
- [85] B. Efron, Bootstrap methods: Another look at the Jackknife, *Ann. Stat.* **7**, 1 (1979).
- [86] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, Fast and unconditional all-microwave reset of a superconducting qubit, *Phys. Rev. Lett.* **121**, 060502 (2018).
- [87] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic Purcell filter, *Phys. Rev. Appl.* **17**, 044016 (2022).
- [88] Yu Zhou, Zhenxing Zhang, Zelong Yin, Sainan Huai, Xiu Gu, Xiong Xu, Jonathan Allcock, Fuming Liu, Guanglei Xi, Qiaonian Yu, Hualiang Zhang, Mengyu Zhang, Hekang Li, Xiaohui Song, Zhan Wang, Dongning Zheng, Shuoming An, Yarui Zheng, and Shengyu Zhang, Rapid and unconditional parametric reset protocol for tunable superconducting qubits, *Nat. Commun.* **12**, 5924 (2021).
- [89] Jiayu Ding, Yulong Li, He Wang, Guangming Xue, Tang Su, Chenlu Wang, Weijie Sun, Feiyu Li, Yujia Zhang, Yang Gao, Jun Peng, Zhi Hao Jiang, Yang Yu, Haifeng Yu, and Fei Yan, Multipurpose architecture for fast reset and protective readout of superconducting qubits, *Phys. Rev. Appl.* **23**, 014012 (2025).
- [90] Gihwan Kim, Andreas Butler, Vinicius S. Ferreira, Xueyue (Sherry) Zhang, Alex Hadley, Eunjong Kim, and Oskar Painter, Fast unconditional reset and leakage reduction of a tunable superconducting qubit via an engineered dissipative bath, *Phys. Rev. Appl.* **24**, 014013 (2025).
- [91] Daniel Sank, Zijun Chen, Mostafa Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, T. White, J. Wenner, Alexander N. Korotkov, and John M. Martinis, Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation, *Phys. Rev. Lett.* **117**, 190503 (2016).
- [92] Mostafa Khezri, Alex Opremcak, Zijun Chen, Kevin C. Miao, Matt McEwen, Andreas Bengtsson, Theodore White, Ofer Naaman, Daniel Sank, Alexander N. Korotkov, Yu Chen, and Vadim Smelyanskiy, Measurement-induced state transitions in a superconducting qubit: Within the rotating-wave approximation, *Phys. Rev. Appl.* **20**, 054008 (2023).
- [93] György P. Gehér, Marcin Jastrzebski, Earl T. Campbell, and Ophelia Crawford, To reset, or not to reset—that is the question, *npj Quantum Inf.* **11**, 39 (2025).
- [94] L. Pereira, J. J. García-Ripoll, and T. Ramos, Parallel tomography of quantum non-demolition measurements in multi-qubit devices, *npj Quantum Inf.* **9**, 22 (2023).
- [95] Bryan T. Gard, Zachary Parrott, Kurt Jacobs, José Aumentado, and Raymond W. Simmonds, Fast high-fidelity quantum nondemolition readout of a superconducting qubit with tunable transverse couplings, *Phys. Rev. Appl.* **21**, 024008 (2024).
- [96] R. Dassonneville, T. Ramos, V. Milchakov, L. Planat, É. Dumur, F. Foroughi, J. Puertas, S. Leger, K. Bharadwaj, J. Delaforce, C. Naud, W. Hasch-Guichard, J. J. García-Ripoll, N. Roch, and O. Buisson, Fast high-fidelity quantum nondemolition qubit readout via a nonperturbative cross-Kerr coupling, *Phys. Rev. X* **10**, 011045 (2020).
- [97] Cyril Mori, Vladimir Milchakov, Francesca d'Esposito, Lucas Ruela, Shelender Kumar, Vishnu Narayanan Suresh, Waël Ardati, Dorian Nicolas, Quentin Ficheux, Nicolas Roch *et al.*, High-power readout of a transmon qubit using a nonlinear coupling, [arXiv:2507.03642](https://arxiv.org/abs/2507.03642).

- [98] Fabian Marxer, Jakub Mrozek, and Joona Andersson, Dataset for above 99.9% fidelity single-qubit gates, two-qubit gates, and readout in a single superconducting quantum device (2025), <https://doi.org/10.5281/zenodo.17052889>.
- [99] C. C. Bultink, B. Tarasinski, N. Haandbæk, S. Poletto, N. Haider, D. J. Michalak, A. Bruno, and L. DiCarlo, General method for extracting the quantum efficiency of dispersive qubit readout in circuit QED, *Appl. Phys. Lett.* **112**, 092601 (2018).
- [100] Ji Chu and Fei Yan, Coupler-assisted controlled-phase gate with enhanced adiabaticity, *Phys. Rev. Appl.* **16**, 054020 (2021).
- [101] Neill Lambert, Eric Giguère, Paul Menczel, Boxi Li, Patrick Hopf, Gerardo Suárez, Marc Gali, Jake Lishman, Rushiraj Gadhvi, Rochisha Agarwal, Asier Galicia, Nathan Shammah, Paul Nation, J. R. Johansson, Shahnawaz Ahmed, Simon Cross, Alexander Pitchford, and Franco Nori, Qutip 5: The quantum toolbox in Python, [arXiv:2412.04705](https://arxiv.org/abs/2412.04705).
- [102] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, Dynamics of superconducting qubit relaxation times, *npj Quantum Inf.* **8**, 132 (2022).
- [103] C. W. Clenshaw, A note on the summation of Chebyshev series, *Math. Comput.* **9**, 118 (1955).
- [104] Edward W. Ng, On the direct summation of series involving higher transcendental functions, *J. Comput. Phys.* **3**, 334 (1968).
- [105] M. M. Saffren and Edward W. Ng, Recursive algorithms for the summation of certain series, *SIAM J. Math. Anal.* **2**, 31 (1971).
- [106] Milton Abramowitz and Irene A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Applied Mathematics Series, Vol. 55 (U.S. National Bureau of Standards, Washington, D.C., 1964).
- [107] Kevin C. Miao, Matt McEwen, Juan Atalaya, Dvir Kafri, Leonid P. Pryadko, Andreas Bengtsson, Alex Opremcak, Kevin J. Satzinger, Zijun Chen, Paul V. Klimov *et al.*, Overcoming leakage in quantum error correction, *Nat. Phys.* **19**, 1780 (2023).
- [108] Michael A Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [109] Vinay Tripathi, Noah Goss, Arian Vezvaei, Long B. Nguyen, Irfan Siddiqi, and Daniel A. Lidar, Qudit dynamical decoupling on a superconducting quantum processor, *Phys. Rev. Lett.* **134**, 050601 (2025).
- [110] Lawrence R. Rabiner and Biing-Hwang Juang, An introduction to hidden Markov models, *IEEE ASSP Mag.* **3**, 4 (1986).
- [111] Zoubin Ghahramani, An introduction to hidden Markov models and Bayesian networks, *Int. J. Pattern Recognit. Artif. Intell.* **15**, 9 (2001).
- [112] E. A. Catchpole and Byron Morgan, Detecting parameter redundancy, *Biometrika*, **84**, 187 (1997).
- [113] Jie Ding, Vahid Tarokh, and Yuhong Yang, Model selection techniques: An overview, *IEEE Signal Process. Mag.* **35**, 16 (2018).
- [114] J. F. Marques, H. Ali, B. M. Varbanov, M. Finkel, H. M. Veen, S. L. M. van der Meer, S. Valles-Sanclemente, N. Muthusubramanian, M. Beekman, N. Haider, B. M. Terhal, and L. DiCarlo, All-microwave leakage reduction units for quantum error correction with superconducting transmon qubits, *Phys. Rev. Lett.* **130**, 250602 (2023).