Robust CZ gate against flux line memory

Yao Song(宋垚)[1] and Xiu-Hao Deng(邓修豪)[2,3*]

[1]Shenzhen Institute for Quantum Science and Engineering,
Southern University of Science and Technology, Shenzhen 518048, China

[2]International Quantum Academy, Shenzhen 518048, China

[3]Shenzhen Branch, Hefei National Laboratory, Shenzhen 518048, China

Abstract

High-fidelity CZ gates are central to superconducting quantum processors, but flux line distortion complicates the reuse of calibrated gate pulses. Continuous predistortion can compensate for pulse history by updating the input across gate boundaries, provided that the line model and initialization are accurate. To enable reliable reuse of fixed, gate-local predistorted CZ pulses without continuous compensation and reduce sensitivity to an uncertain initial line state, we model flux line distortion as the dynamics of a stateful classical actuator coupled to a quantum system. Using a first-order Dyson expansion, we derive the error generators induced by variations in the initial flux-line state. We then design a robust CZ gate by optimizing the flux pulse to suppress these generators and minimize the residual flux line state at the gate exit. In a one-pole flux line model, the optimized CZ pulse exhibits the expected first-order robustness plateau and, over ten consecutive gates, maintains high fidelity across a broader initial state range than continuous predistortion initialized at zero. For a more practical three-pole model, the optimized CZ pulse achieves $F_{\mathrm{avg}} = 99.998\%$ in closed-system simulations, suppresses all first-order error generators, and brings the residual flux close to zero. Reusing the fixed, gate-local predistorted pulse without additional waiting between gates yields $F_{\mathrm{avg}} = 99.97\%$ over ten gates and reduces infidelity by approximately 2.3×10^3 relative to the baseline under the same predistortion method. These results show that explicitly accounting for flux line memory enables high-fidelity reuse of fixed, gate-local predistorted CZ pulses across repeated gate sequences.

Keywords: superconducting quantum processors, robust CZ gate, flux line distortion, robust quantum control
PACS: 03.67.Lx, 85.25.Cp, 02.30.Yy

1. Introduction

High-fidelity CZ gates are essential to scalable superconducting quantum processors [1–3]. Tunable buses and couplers provide switchable two-qubit interactions and have enabled high-fidelity CZ gates [4–11]. In implementations based on flux tuning, the conditional phase is generated by bringing a computational state close to an avoided crossing with a noncomputational state. Both the accumulated phase and leakage are therefore sensitive to the flux pulse delivered to the chip. [6,12,13]. Before reaching the chip, the flux pulse

*Corresponding author. E-mail: dengxiuhao@iqasz.cn

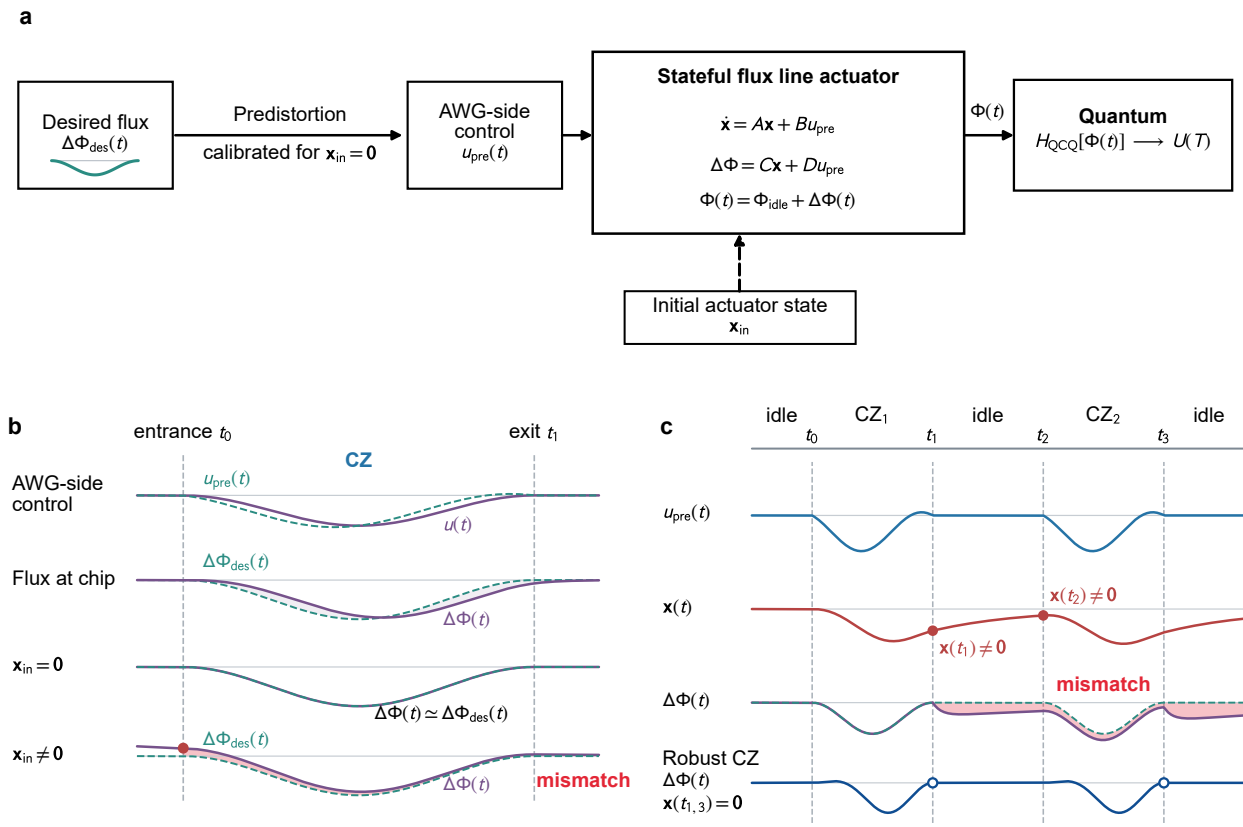


Fig. 1. Stateful actuator–quantum dynamics of flux line distortion. (a) Gate-local predistortion calibrated for the zero initial actuator state, $\boldsymbol{x}_{\rm in}=\mathbf{0}$, maps the desired dynamic flux $\Delta\Phi_{\rm des}(t)$ to the predistorted AWG control $u_{\rm pre}(t)$. The AWG control and initial actuator state jointly determine the delivered dynamic flux $\Delta\Phi(t)$ and hence the total flux $\Phi(t)$ that drives the quantum system. (b) The first row contrasts the uncorrected and predistorted AWG controls, $u(t)$ and $u_{\rm pre}(t)$; the second shows the distortion of the flux delivered at the chip when $u(t)$ is applied directly. The lower rows show that predistortion recovers the desired flux for $\boldsymbol{x}_{\rm in}=\mathbf{0}$, whereas a nonzero initial state adds the free response of the flux line and produces a flux mismatch. (c) In a CZ sequence, the actuator state left at one gate exit evolves during the idle interval and becomes the initial state of the next gate. Memory robustness suppresses the first-order gate sensitivity to this inherited state, while state closing suppresses the residual state generated by each gate, allowing the same predistorted pulse to be repeated. Red shading marks the flux mismatch; all pulses are schematic.

is distorted by the finite bandwidth, impedance mismatches, and long settling tails of the room-temperature electronics and cryogenic control line [14–18]. The flux line therefore has an internal dynamical state set by preceding controls, making the delivered flux history dependent. [12,19] We therefore model it as a classical dynamical actuator with memory rather than a memoryless input-output map. [20]

Measurements with the qubit itself or an embedded cryogenic transducer can reconstruct the pulse at or near the chip and determine the control line response [15–17]. Such characterization supports model-based predistortion [21], while calibration against measured gate errors can reduce the remaining control errors [17, 18,22–24]. Pulse shaping addresses complementary errors: derivative corrections suppress leakage, whereas Net-Zero symmetry reduces sensitivity to slow distortions in the flux line [12,13,25].

The distinction between an isolated pulse and a continuous sequence matters when calibrated gate inputs are reused. Gate-local predistortion constructs an input for a specified initial line state, taken here to be zero. Repeating that fixed input does not compensate for the state left by preceding gates. Continuous predistortion

instead applies an inverse filter across successive pulses while retaining its internal state, either during pulse generation or through continuous digital filtering [17]. With an accurate model and matched initialization, it can compensate for history-dependent distortion; a nonzero line state alone does not invalidate it. However, an unknown difference between the actual and assumed initial states produces a free response that the model-based inverse does not cancel. We therefore ask whether a fixed CZ input can remain accurate under this uncertainty, while limiting the residual state it generates for later gates.

This question shifts the control objective from reconstructing the flux pulse to suppressing the quantum error caused by actuator memory. Robust quantum control suppresses the effects of Hamiltonian perturbations through numerical pulse optimization or analytic constructions that cancel leading error terms [26–37]. The response of the control hardware can also be included directly in pulse optimization [20,38]. Flux line memory, however, is not represented by a static or quasistatic perturbation in the quantum Hamiltonian. It instead enters through the initial state of a classical dynamical actuator, and each component of that state produces a distinct flux perturbation over time. Robust control of a CZ gate must therefore suppress the errors induced by variations in the actuator state at the gate entrance and constrain the state remaining at the gate exit.

To impose these two conditions, we describe the flux line in state space and couple its dynamics to the quantum system, as shown in Fig. 1. Given the AWG input and the actuator state at the gate entrance, the coupled model determines the delivered flux, the quantum propagator, and the actuator state at the gate exit. A first-order Dyson expansion then associates each component of the initial actuator state with an error generator in the interaction picture. Memory robustness requires the computational and leakage components of every generator to vanish, whereas state closing requires the actuator state written by the pulse to vanish at the gate exit. The former removes the first-order sensitivity of the current gate to inherited memory, while the latter prevents the pulse from adding a new residual to the state passed to the next gate.

All numerical results reported here are obtained from closed-system simulations without decoherence, stochastic flux noise, or transfer-function drift. We first test the two boundary conditions in the one-pole model. The optimized pulse exhibits a clear first-order robustness plateau: the infidelity of a single gate remains near its value at $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$ up to $x_{\mathrm{c}} \simeq 3.5 \times 10^{-3}\Phi_0$ and then grows approximately as the fourth power of the initial state magnitude. By contrast, the baseline pulse retains the quadratic scaling expected without first-order cancellation [Fig. 3(d)]. The fixed robust input also retains high ten-gate fidelity over a broader initial-state range than continuous predistortion initialized at zero [Fig. 3(e)]. This advantage concerns initialization uncertainty: continuous predistortion gives lower errors when its assumed and actual initial states agree. For a sequence of 20 gates starting from zero state, the robust pulse has a cumulative infidelity of 1.2×10^{-4}, compared with 0.79 for the fixed baseline pulse [Fig. 3(f)]. We then test the method in a three-pole model fitted to a measured flux line step response [22]. The optimized pulse reaches $F_{\mathrm{avg}} = 99.998\%$ for a single gate and suppresses the error generators associated with all three actuator modes [Fig. 4(c)]. With no additional waiting time between gates, the complete sequence of ten gates reaches $F_{\mathrm{avg}} = 99.97\%$ [Fig. 5(f)]. Under the same gate-local predistortion protocol, this corresponds to a reduction in sequence infidelity by a factor of approximately 2.3×10^3 relative to both the baseline and Net-Zero pulses. Optimizations with only one boundary condition separate their roles: memory robustness suppresses errors caused by inherited actuator states, whereas state closing reduces the residual state at the gate exit generated by each gate [Fig. 6].

The remainder of the paper is organized as follows. Section 2 defines the qubit-coupler-qubit(QCQ) system and the stateful actuator–quantum forward dynamics. Section 3 derives a first-order error generator for each component of the initial actuator state, formulates memory robustness and state closing as separate conditions at the gate entrance and exit, and explains their roles in the repeated CZ gates. Section 4 presents the numerical optimization and validation of memory-robust CZ pulses in one-pole and three-pole flux line models. Section 5 summarizes the conclusions and scope.

2. Stateful actuator–quantum forward dynamics

The AWG control pulse is transmitted through the flux line, whose response determines the flux applied to the QCQ system. Because the line has internal dynamics, the delivered flux depends on both the AWG input and the actuator state at the gate entrance. The same input can therefore produce different flux pulses for different initial actuator states. To model this dependence, we first specify the QCQ system under flux control and then describe the flux line as a classical actuator using a state space representation. We couple the two models to track the quantum evolution and actuator state during a CZ gate duration. The coupled model, summarized in Fig. 1, provides the boundary quantities used to derive the robustness conditions in Sec. 3.

2.1. QCQ system under flux control

The QCQ system consists of two transmon qubits coupled through a tunable coupler. Treating the three transmons as weakly anharmonic oscillators, we write the Hamiltonian as

$$\begin{aligned} H_{\mathrm{QCQ}}(\Phi) = & \sum_{j\in\{1,2\}} \left[\omega_j n_j + \frac{\alpha_j}{2} n_j(n_j-1)\right] \\ & + \omega_c(\Phi) n_c + \frac{\alpha_c}{2} n_c(n_c-1) \\ & + \sum_{(j,k)\in\{(1,c),(c,2),(1,2)\}} g_{jk}\left(a_j^\dagger a_k + a_j a_k^\dagger\right), \end{aligned} \tag{1}$$

where a_j and $n_j = a_j^\dagger a_j$ are the annihilation and number operators of qubit j, respectively. The qubit frequencies ω_1 and ω_2 are fixed, while the coupler frequency $\omega_c(\Phi)$ depends on the applied flux [2,6]. The parameters $\alpha_\mu < 0$ denote the qubit anharmonicities, and g_{jk} are the fixed transverse coupling strengths. The total flux $\Phi(t)$ delivered by the line therefore controls the QCQ dynamics through $\omega_c[\Phi(t)]$, as illustrated in Fig. 2(a).

Using units with $\hbar = 1$, we model the flux dependence of an asymmetric SQUID coupler as

$$\begin{aligned} E_{J,c}(\Phi) &= E_{J\Sigma}\sqrt{\cos^2(\pi\Phi/\Phi_0) + d_c^2 \sin^2(\pi\Phi/\Phi_0)}, \\ \omega_c(\Phi) &\simeq \sqrt{8E_{C,c}E_{J,c}(\Phi)} - E_{C,c}. \end{aligned} \tag{2}$$

Here $E_{J\Sigma} = E_{J,L} + E_{J,R}$ is the total Josephson energy, $d_c = |E_{J,L} - E_{J,R}|\,/E_{J\Sigma}$ is the junction asymmetry, $E_{C,c}$ is the coupler charging energy, and Φ_0 is the magnetic flux quantum [39,40]. An experimentally measured coupler dispersion $\omega_c(\Phi)$ can replace this analytic form without changing the coupled model introduced below. For the parameters used here, Fig. 2(b) shows the nonlinear dependence of ω_c on Φ and the flux values corresponding to the CZ interaction region.

Because the static couplings remain nonzero at idle point, we define the computational basis using the four eigenstates of $H_{\mathrm{QCQ}}(\Phi_{\mathrm{idle}})$ that are adiabatically connected to the corresponding bare states. With ket

ordering (q_1, c, q_2), these dressed computational states are $\{|\widetilde{000}\rangle, |\widetilde{001}\rangle, |\widetilde{100}\rangle, |\widetilde{101}\rangle\}$. The projector P defines the computational subspace $\mathcal{H}_P = \mathrm{Ran}(P)$, while $Q = \mathbb{I} - P$ projects onto the orthogonal leakage subspace. Both projectors are defined at the idle point and held fixed throughout the time evolution and gate evaluation.

The spectrum of the QCQ system in the two-excitation subspace is shown in Fig. 2(c). At each ω_c, we label the six dressed eigenstates using the one-to-one assignment that maximizes the sum of their squared overlaps with the six bare states

$$\pi^\star(\omega_c) = \arg\max_{\pi} \sum_{\nu} \left| \langle \nu | \widetilde{\psi}_{\pi(\nu)}(\omega_c) \rangle \right|^2,$$

where ν labels a bare state in this subspace and π is a permutation of the six dressed eigenstates. The blue and purple segments in Fig. 2(c) identify the eigenstates assigned the labels $|101\rangle$ and $|020\rangle$, respectively. The assignment is repeated independently at each ω_c. The $|101\rangle$ and $|020\rangle$ labels can therefore switch between the two dressed eigenstates across an avoided crossing, even though their eigenenergies remain continuous. Within the flux range used for gate design, the avoided crossing between the states associated with $|101\rangle$ and $|020\rangle$ produces the conditional energy shift used to accumulate the conditional phase.

Equations (1) and (2) specify how the total flux applied to the coupler enters the QCQ Hamiltonian, but the AWG sets the input pulse rather than the flux delivered to the coupler. As shown in Fig. 1(a), gate-local predistortion determines the AWG-side input $u_{\mathrm{pre}}(t)$ needed to reproduce the desired flux $\Delta\Phi_{\mathrm{des}}(t)$. The flux line transforms the current $u_{\mathrm{pre}}(t)$ into the delivered dynamic flux $\Delta\Phi(t)$. For a nonzero flux-line state, this output also depends on earlier controls. The resulting total flux $\Phi(t) = \Phi_{\mathrm{idle}} + \Delta\Phi(t)$, rather than the desired flux pulse, drives the QCQ Hamiltonian. Thus, the AWG-side input and the flux-line state jointly determine the QCQ dynamics as seen in Eq. (4).

2.2. Coupled actuator and quantum system

For the bandwidth and signal amplitudes relevant to CZ control, we represent the flux line as a stable, finite-order, continuous-time linear system in state-space form [41,42]. Combining the actuator state and output equations with the quantum propagator $U(t)$ gives

$$\dot{\boldsymbol{x}}(t) = A\boldsymbol{x}(t) + Bu(t), \qquad \boldsymbol{x}(0) = \boldsymbol{x}_{\mathrm{in}}, \tag{3a}$$

$$\Delta\Phi(t) = C\boldsymbol{x}(t) + Du(t), \tag{3b}$$

$$\mathrm{i}\dot{U}(t) = H_{\mathrm{QCQ}}[\Phi(t)]U(t), \qquad U(0) = \mathbb{I}. \tag{3c}$$

We refer to Eq. (3) as the *stateful actuator–quantum forward dynamics*. The coupling is one way: the actuator determines the flux applied to the QCQ system, while backaction of the quantum system on the flux line is neglected.

The internal state $\boldsymbol{x}(t) \in \mathbb{R}^{n_x}$ carries the memory of earlier controls. We choose $\boldsymbol{x} = \mathbf{0}$ to represent the fully reset idle state, so $\boldsymbol{x}_{\mathrm{in}} = \boldsymbol{x}(0)$ is the actuator state at the gate entrance. The AWG input $u(t)$ and delivered dynamic flux $\Delta\Phi(t)$ are scalar, with $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times 1}$, $C \in \mathbb{R}^{1 \times n_x}$, and $D \in \mathbb{R}$. In the identified realization, A governs the internal dynamics, B couples the AWG input into the state dynamics, C maps the state to the delivered flux, and D describes direct feedthrough. The calibrated static offset is included in Φ_{idle}, and the model order n_x is chosen to reproduce the dynamics resolved by end-to-end measurements. We assume

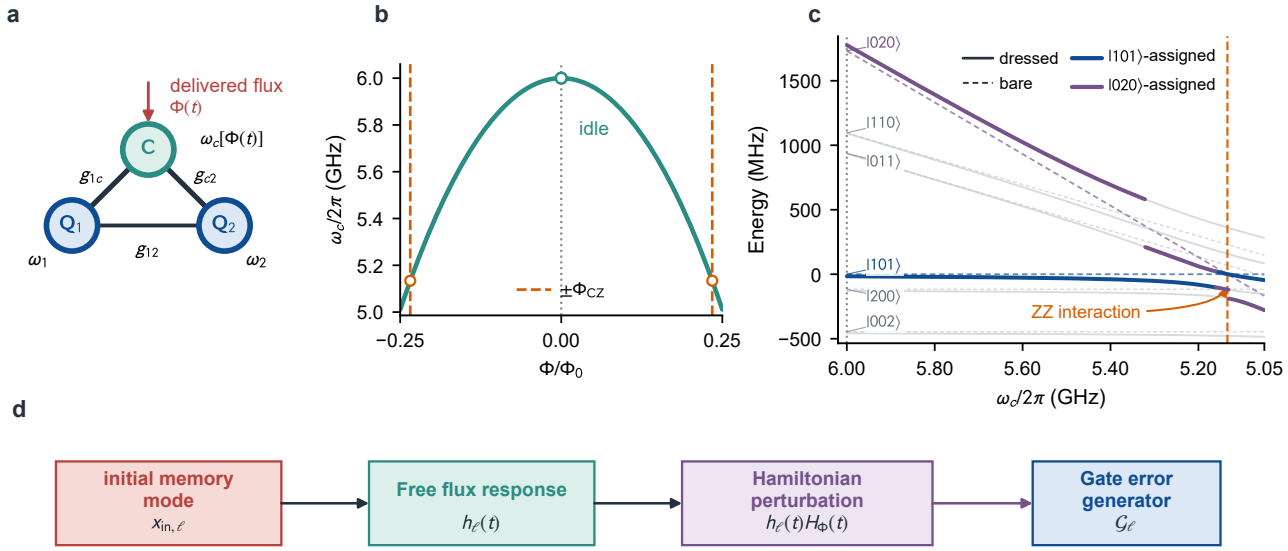


Fig. 2. QCQ flux control, the ZZ interaction used for CZ gates, and the mapping from initial actuator memory to quantum error. (a) Two fixed-frequency qubits coupled through a tunable coupler, with direct qubit–qubit coupling g_{12}. The delivered flux $\Phi(t)$ tunes the coupler frequency $\omega_c[\Phi(t)]$. (b) Flux dependence of $\omega_c/2\pi$ for the asymmetric SQUID, shown over $|\Phi| \leq 0.25\Phi_0$. The gray dotted line marks the idle point, and the orange dashed lines mark $\pm\Phi_{\mathrm{CZ}}$, where the bare $|101\rangle$ and $|020\rangle$ energies coincide. (c) Spectrum in the two-excitation manifold, with ket ordering (q_1, c, q_2), as $\omega_c/2\pi$ decreases from its 6.0 GHz idle value. The plotted energy is relative to the bare $|101\rangle$ level; solid and dashed curves denote dressed and bare energies, respectively. Blue and purple segments identify $|\widetilde{\psi}_{\pi^\star(101)}\rangle$ and $|\widetilde{\psi}_{\pi^\star(020)}\rangle$, respectively, using the bare-state overlap assignment defined in the text; the assigned labels can switch between continuous eigenenergy branches. The orange dashed line marks the resonance of the bare $|101\rangle$ and $|020\rangle$ states, while the arrow identifies the associated avoided crossing that contributes to the ZZ interaction. (d) The initial memory amplitude $x_{\mathrm{in},\ell}$ contributes the free flux perturbation $(x_{\mathrm{in},\ell}/\Phi_0)h_\ell(t)$ and the Hamiltonian perturbation $(x_{\mathrm{in},\ell}/\Phi_0)h_\ell(t)H_\Phi(t)$. First-order Dyson accumulation in the interaction frame gives the error generator $(x_{\mathrm{in},\ell}/\Phi_0)\mathcal{G}_\ell$. The last three nodes show the response functions and generator with the common amplitude factor $x_{\mathrm{in},\ell}/\Phi_0$ factored out.

that the identified realization remains time invariant during a gate sequence; slower drift requires recalibration and is not a noise source in this work.

Solving the linear state-space equations gives an exact decomposition of the state and output into contributions from the initial actuator state and the applied input

$$\begin{aligned} \boldsymbol{x}(t) &= \mathrm{e}^{At}\boldsymbol{x}_{\mathrm{in}} + \int_0^t \mathrm{e}^{A(t-t')}Bu(t')\,\mathrm{d}t', \\ \Delta\Phi(t) &= C\mathrm{e}^{At}\boldsymbol{x}_{\mathrm{in}} + C\int_0^t \mathrm{e}^{A(t-t')}Bu(t')\,\mathrm{d}t' \\ &\quad + Du(t) \\ &= \Delta\Phi_{\mathrm{memory}}(t) + \Delta\Phi_{\mathrm{driven}}(t). \end{aligned} \tag{4}$$

Here $\Delta\Phi_{\mathrm{memory}}(t) = C\mathrm{e}^{At}\boldsymbol{x}_{\mathrm{in}}$ is the free response; $\Delta\Phi_{\mathrm{driven}}(t)$ contains the convolution and direct feedthrough terms. Gate-local predistortion reconstructs the desired flux for $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, but its fixed input does not cancel an added free response. Continuous predistortion compensates for the history predicted by the line model, provided its initialization matches the actual state. A difference between these initial states leaves a free-response flux error. Our robust design suppresses the resulting quantum error rather than requiring that flux error to vanish. It also constrains $\boldsymbol{x}(T; \boldsymbol{x}_{\mathrm{in}} = \mathbf{0})$, the driven state left at the gate exit, which pulse reconstruction alone does not constrain.

3. Conditions for CZ gates robust to flux line memory

Within the coupled model of Eq. (3), flux line memory affects a CZ sequence at both gate boundaries. The initial actuator state changes the flux delivered during the current gate, while the driven residual at the gate exit is inherited by later gates. We first derive an error generator for each component of the initial state and use these generators to formulate memory robustness. We then formulate an independent state closing condition for the driven residual. Finally, a propagation map shows how the two conditions affect a sequence of CZ gates.

3.1. From initial actuator state perturbations to quantum error generators

Memory robustness is imposed on the quantum gate, whereas $\boldsymbol{x}_{\mathrm{in}}$ describes the initial state of the classical actuator. To connect them, we propagate a small perturbation of the initial actuator state through the delivered flux and into the quantum evolution. For a fixed AWG input, Eq. (4) shows that all dependence on $\boldsymbol{x}_{\mathrm{in}}$ enters through the free response. We expand $\boldsymbol{x}_{\mathrm{in}}$ in a basis $\{\boldsymbol{e}_\ell\}_{\ell=1}^{n_x}$. Each basis vector is scaled to a reference amplitude Φ_0, so that $\boldsymbol{x}_{\mathrm{in}} = \sum_{\ell=1}^{n_x}(x_{\mathrm{in},\ell}/\Phi_0)\boldsymbol{e}_\ell$ and $x_{\mathrm{in},\ell}$ has units of flux. In the modal realizations used below, $x_{\mathrm{in},\ell}$ is the initial stored flux of the ℓth state component.

This basis decomposition separates the reference flux trajectory from the flux perturbation associated with each component of the initial actuator state. Let $\bar{\Phi}(t) \equiv \Phi(t;\boldsymbol{x}_{\mathrm{in}}=\boldsymbol{0}) = \Phi_{\mathrm{idle}} + \Delta\Phi_{\mathrm{driven}}(t)$ be the reference total flux produced by the same AWG input for $\boldsymbol{x}_{\mathrm{in}} = \boldsymbol{0}$. Equation (4) then gives the exact decomposition $\Phi(t;\boldsymbol{x}_{\mathrm{in}}) - \bar{\Phi}(t) = \sum_{\ell=1}^{n_x}(x_{\mathrm{in},\ell}/\Phi_0)h_\ell(t)$, where

$$h_\ell(t) \equiv \Phi_0 \frac{\partial\Phi(t;\boldsymbol{x}_{\mathrm{in}})}{\partial x_{\mathrm{in},\ell}}\bigg|_{\boldsymbol{x}_{\mathrm{in}}=\boldsymbol{0}} = C\mathrm{e}^{At}\boldsymbol{e}_\ell. \tag{5}$$

Equation (5) defines $h_\ell(t)$ as the free flux response to the basis displacement $\boldsymbol{e}_\ell$. It has flux units, and the initial coordinate $x_{\mathrm{in},\ell}$ contributes the physical flux perturbation $x_{\mathrm{in},\ell}/\Phi_0 h_\ell(t)$.

This flux perturbation changes the quantum evolution. Let $U_0(t)$ be the reference propagator under $\bar{\Phi}(t)$, and define $H_\Phi(t) = \partial_\Phi H_{\mathrm{QCQ}}|_{\bar{\Phi}(t)}$ as the instantaneous Hamiltonian sensitivity to flux along the reference flux trajectory. A small flux perturbation $\delta\Phi(t)$ produces $\delta H(t) = \delta\Phi(t)H_\Phi(t) + \mathcal{O}(\delta\Phi^2)$. An initial state displacement therefore produces the time-dependent Hamiltonian perturbation $(x_{\mathrm{in},\ell}/\Phi_0)h_\ell(t)H_\Phi(t)$. At first order, the Dyson expansion in the interaction picture gives [43]

$$U(T;x_{\mathrm{in},\ell}) = U_0(T)\left[\mathbb{I} - \mathrm{i}\frac{x_{\mathrm{in},\ell}}{\Phi_0}\mathcal{G}_\ell\right] + \mathcal{O}\left[\left(\frac{x_{\mathrm{in},\ell}}{\Phi_0}\right)^2\right], \tag{6a}$$

$$\mathcal{G}_\ell = \int_0^T h_\ell(t)U_0^\dagger(t)H_\Phi(t)U_0(t)\,\mathrm{d}t. \tag{6b}$$

Here $\mathcal{G}_\ell$ is the first-order gate error generator in $x_{\mathrm{in},\ell}/\Phi_0$. The factors $h_\ell(t)$ and $H_\Phi(t)$ describe, respectively, the free flux response of the actuator and the sensitivity of the Hamiltonian to flux along the reference flux trajectory. The mapping from the initial actuator state perturbation to $\mathcal{G}_\ell$ is shown in Fig. 2(d).

Composite pulse and geometric control methods commonly treat systematic errors as static or quasistatic perturbations, although geometric approaches can also address driven noise [28,30–32]. Flux line memory instead generates the structured time-dependent perturbation $h_\ell(t)H_\Phi(t)$. Equations (6) map this perturbation to the first-order gate error generator $\mathcal{G}_\ell$. Appendix B derives these expressions using the Dyson expansion.

3.2. Memory robustness and state closing

Memory robustness suppresses the first-order logical and leakage errors induced by each component of a small initial actuator state. Here, "logical error" means a gate error within the computational subspace, not an error of an encoded logical qubit. State closing separately suppresses the actuator state written by the pulse at the gate exit.

3.2.1. Gate entrance: memory robustness

The first-order expansion in Eq. (6) is linear in the initial actuator state. For a general small $\boldsymbol{x}_{\mathrm{in}}$, the total first-order error generator is $\sum_{\ell=1}^{n_x}(x_{\mathrm{in},\ell}/\Phi_0)\mathcal{G}_\ell$. Because the components $x_{\mathrm{in},\ell}$ vary independently, robustness requires the relevant blocks of every $\mathcal{G}_\ell$ to vanish separately. For a reference evolution that realizes the target CZ gate and preserves $\mathcal{H}_P$, these are the traceless logical block and the leakage block. With $\dim(\mathcal{H}_P) = 4$, we therefore impose

$$\begin{aligned} \mathcal{G}_\ell^{\log} &\equiv P\mathcal{G}_\ell P - \frac{\mathrm{Tr}(P\mathcal{G}_\ell P)}{4}P = 0, \\ \mathcal{G}_\ell^{\mathrm{leak}} &\equiv Q\mathcal{G}_\ell P = 0, \qquad \ell = 1, \ldots, n_x. \end{aligned} \tag{7}$$

The first condition eliminates the traceless logical error generator, leaving only an irrelevant global phase, whereas the second eliminates first-order coupling from the computational subspace to leakage subspace. Equation (7) is imposed for every actuator state component through its free response $h_\ell(t)$. We use the Φ_0-normalized basis defined above. These conditions constrain the first-order error accumulated during the present gate and they do not constrain the actuator state left at its exit.

3.2.2. Gate exit: state closing

The gate exit condition follows directly from the actuator solution. Evaluating Eq. (4) at $t = T$ gives the exact map from the initial state to the terminal state,

$$\begin{aligned} \boldsymbol{x}(T; \boldsymbol{x}_{\mathrm{in}}) &= \mathrm{e}^{AT}\boldsymbol{x}_{\mathrm{in}} + \boldsymbol{r}[u], \\ \boldsymbol{r}[u] &= \int_0^T \mathrm{e}^{A(T-t')}Bu(t')\,\mathrm{d}t'. \end{aligned} \tag{8}$$

Equation (8) separates the inherited and driven contributions to the terminal state. The first term propagates the initial actuator state, whereas $\boldsymbol{r}[u]$ is the state written by the present AWG input. When $\boldsymbol{x}_{\mathrm{in}} = \boldsymbol{0}$, the terminal state is $\boldsymbol{r}[u]$, so state closing requires

$$\boldsymbol{x}(T; \boldsymbol{x}_{\mathrm{in}} = \boldsymbol{0}) = \boldsymbol{r}[u] = \boldsymbol{0}. \tag{9}$$

State closing prevents the AWG input from writing a residual actuator state, but it does not remove an arbitrary initial state.

3.3. Propagation of flux line memory in repeated CZ gates

In a sequence of consecutive CZ gates, the actuator state at the exit of one gate becomes the initial state of the next. Let $\boldsymbol{x}^{(n)}$ denote the state after n completed gates, and hence at the entrance to gate $n+1$, with

$\boldsymbol{x}^{(0)} = \boldsymbol{x}_{\rm in}$. Defining the state transition over one gate as $A_T \equiv \mathrm{e}^{AT}$, reuse of the same AWG input in Eq. (8) gives

$$\begin{aligned} \boldsymbol{x}^{(n+1)} &= A_T \boldsymbol{x}^{(n)} + \boldsymbol{r}[u], \\ \boldsymbol{x}^{(n)} &= A_T^n \boldsymbol{x}_{\rm in} + \sum_{q=0}^{n-1} A_T^q \boldsymbol{r}[u]. \end{aligned} \tag{10}$$

The recurrence separates two contributions to the actuator state. The first term propagates the state present before the sequence through n gates, whereas the sum accumulates the driven residual $\boldsymbol{r}[u]$ written by each gate.

When state closing holds, the driven residual does not accumulate, although the inherited state $A_T^n \boldsymbol{x}_{\rm in}$ can remain nonzero. Memory robustness suppresses the first-order gate error induced by each component of this inherited state, but it does not alter Eq. (10). Without state closing, a nonzero $\boldsymbol{r}[u]$ can accumulate and move the actuator state beyond the range in which the first-order expansion is accurate. Both conditions are therefore required to reuse the same pulse across a CZ gates sequence. Figure 6 tests their distinct roles using pulses optimized for only one condition.

4. Numerical validation of memory-robust CZ gates

We optimize finite-basis CZ pulses under the memory-robustness and state-closing conditions of Sec. 3. A one-pole model tests first-order cancellation in the simplest stateful line, and a three-pole model fitted to an experimental step response introduces multiple memory time scales. Finite initial-state scans, continuously propagated CZ sequences, and single-condition optimizations are reserved for validation and do not enter the pulse objective.

4.1. Numerical setup and pulse optimization

The numerical simulations use the Hamiltonian in Eq. (1) with representative parameters in the range reported for tunable coupler devices [2,6]:

$$\begin{aligned} (\omega_1, \omega_2)/2\pi &= (5.064, 4.904)\ \text{GHz}, \\ \omega_c(\Phi_{\rm idle})/2\pi &= 6.000\ \text{GHz}, \\ (\alpha_1, \alpha_c, \alpha_2)/2\pi &= (-0.276, -0.300, -0.286)\ \text{GHz}, \\ (g_{1c}, g_{2c}, g_{12})/2\pi &= (0.090, 0.090, 0.00796)\ \text{GHz}, \end{aligned}$$

The SQUID asymmetry is $d_c = 0.10$. Each transmon is truncated to three local levels. Because Eq. (1) conserves the total excitation number $N = n_1 + n_c + n_2$, numerical propagation uses the ten-dimensional $N \leq 2$ invariant subspace. The gate duration is $T = 50$ ns throughout.

We simulate unitary quantum evolution with fixed flux-line parameters. The reported fidelities quantify control accuracy within this model, including coherent gate errors and leakage, rather than predict experimental gate fidelities.

We parameterize the desired flux pulse as

$$\Delta\Phi_{\rm des}(t) = \Phi_{\rm max} \tanh\left[\sum_{m=1}^{M} c_m \sin(m\pi t/T)\right], \tag{11}$$

where M is the number of sine basis functions and $\boldsymbol{c} = (c_1, \ldots, c_M)^{\mathsf{T}}$ contains the dimensionless pulse coefficients. This parameterization gives $\Delta\Phi_{\mathrm{des}}(0) = \Delta\Phi_{\mathrm{des}}(T) = 0$ and enforces $|\Delta\Phi_{\mathrm{des}}(t)| < \Phi_{\max}$, with $\Phi_{\max} = 0.30\Phi_0$.

We compare two predistortion protocols using the same exact line model. Gate-local predistortion computes the AWG input for one desired pulse starting from $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$. The baseline and memory-robust sequences repeat their respective fixed inputs without recomputing the inverse at each gate. Continuous predistortion instead inverts the full baseline target flux sequence, retaining the inverse-filter state across gate boundaries. It is initialized for a zero physical line state. In the initial-state scan of Fig. 3(e), only the actual sequence-start state is varied; the filter is neither reinitialized nor supplied with that perturbation. Thus the comparison tests initialization uncertainty, not line-parameter error or compensation based on a measured state.

We calculate the average gate fidelity, including leakage, as

$$F_{\mathrm{avg}} = \frac{\mathrm{Tr}(K^\dagger K) + \left|\mathrm{Tr}(U_{\mathrm{tar}}^\dagger K)\right|^2}{20}. \tag{12}$$

Here $K = PU(T)P|_{\mathcal{H}_P}$ is the propagator restricted to the four-dimensional computational subspace. A virtual-Z correction is calibrated once at $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$ and incorporated into U_{tar}. For a sequence of n CZ gates, we propagate the coupled flux line and quantum dynamics continuously from $t = 0$ to nT, with no waiting interval or reset of either system between gates. The pulse coefficients and each pulse's nominal virtual-Z calibration remain fixed across all initial-state and sequence scans. We then evaluate Eq. (12) with $K = PU(nT)P|_{\mathcal{H}_P}$ and target U_{tar}^n; the resulting cumulative infidelity $1 - F_{\mathrm{avg}}$ is plotted for the one-pole model in Fig. 3(f). For the sequence tests in Fig. 3(e) and (f), we use 640 piecewise-constant AWG samples per gate. The line state is propagated exactly on each sample, and the quantum propagator uses the matrix exponential of the Hamiltonian at the sample's midpoint flux.

Table 1 summarizes the sampling settings, peak amplitudes, and estimated spectral bandwidths of the gate-local predistorted inputs. The effective sampling rates are 12.8 and 25.6 GSa/s for the one- and three-pole models, respectively. These rates are used for numerical accuracy, not as minimum AWG hardware requirements. We define the baseband spectral cutoff $B_{99.9}$ by

$$\int_{-B_{99.9}}^{B_{99.9}} |\widetilde{u}(f)|^2\, df = 0.999 \int_{-\infty}^{\infty} |\widetilde{u}(f)|^2\, df,$$

where $\widetilde{u}(f)$ is the Fourier transform of the held predistorted input $u(t)$, set to zero outside $[0, T]$. No additional AWG analog-bandwidth or quantization model is included. These spectral estimates do not establish the minimum hardware bandwidth needed to preserve gate fidelity, which requires validation with the actual AWG response.

To impose the memory robustness and state closing conditions in Eqs. (7) and (9), we minimize

$$\begin{aligned}\mathcal{J}(\boldsymbol{c}) = w_F\big[1 - F_{\mathrm{avg}}(\boldsymbol{c})\big] &+ \sum_{\ell=1}^{n_x} w_\ell^{\log} \left\|\mathcal{G}_\ell^{\log}\right\|_{\mathrm{F}}^2 \\ &+ \sum_{\ell=1}^{n_x} w_\ell^{\mathrm{leak}} \left\|\mathcal{G}_\ell^{\mathrm{leak}}\right\|_{\mathrm{F}}^2 \\ &+ \sum_{\ell=1}^{n_x} w_\ell^{r} \left|\frac{r_\ell[u]}{\Phi_0}\right|^2 .\end{aligned} \tag{13}$$

Here w_F, $w_\ell^{\log}$, w_ℓ^{leak}, and w_ℓ^r are the relative weights of the fidelity, logical generator, leakage generator, and state closing terms, respectively. The quantity $r_\ell[u]$ is the ℓth component of $\boldsymbol{r}[u]$ in Eq. (8). The generator

Table 1. Sampling, peak input, and estimated spectral bandwidth for the 50 ns pulses in Figs. 3(e,f) and 5. N is the number of piecewise-constant samples per gate. The sampling settings describe the numerical discretization, not minimum hardware requirements. The peak is $\max_t |u(t)|/\Phi_0$ after gate-local predistortion at zero initial state, expressed in flux-equivalent units rather than AWG voltage. It is distinct from the desired-flux bound $\Phi_{\max}$. $B_{99.9}$ encloses 99.9% of the isolated input pulse energy in $[-B_{99.9}, B_{99.9}]$; the table reports the positive cutoff, not the full two-sided width or the AWG −3 dB bandwidth.

Model	Pulse	N	Δt (ns)	Peak input / Φ_0	$B_{99.9}$ (MHz)
One-pole	Baseline	640	0.078125	0.23930	79
	Memory-robust	640	0.078125	0.27079	344
Three-pole	Baseline	1280	0.0390625	0.23930	74
	Net-Zero	1280	0.0390625	0.30098	210
	Memory-robust	1280	0.0390625	0.25160	243

norms are dimensionless, and each residual component is normalized explicitly by Φ_0. The weights absorb the relative numerical scales of these terms. The fidelity term is evaluated at $\boldsymbol{x}_{\text{in}} = \mathbf{0}$. We use the bound-constrained L-BFGS-B algorithm [44,45]. The weights and pulse initialization are given in Appendix A.

Increasing M enlarges the pulse search space and provides additional freedom to satisfy the robustness conditions, but does not itself ensure robustness. First-order memory robustness requires the generator cancellations in Eq. (7), while state closing separately requires Eq. (9). Nominal fidelity optimization enforces neither condition, irrespective of basis size: it may satisfy them incidentally, whereas the additional terms in Eq. (13) target them explicitly. The matched basis comparison in Appendix C illustrates this distinction.

4.2. One-pole demonstration of memory robustness

The one-pole flux line model provides the simplest setting in which to test memory robustness and state closing:

$$\dot{x}_1 = -\frac{x_1}{\tau_1} + \frac{\eta_1}{\tau_1} u, \qquad \Delta\Phi = ku + x_1, \tag{14}$$

with $(k, \eta_1, \tau_1) = (0.88, 0.12, 12\ \text{ns})$. Because the direct and pole weights satisfy $k + \eta_1 = 1$, the model has unit steady-state gain. Its unit step response is shown in Fig. 3(a). Since the model has a single actuator state component, the objective contains one logical generator norm, one leakage generator norm, and one state closing term. The baseline pulse is parameterized by 5 sine basis functions and is optimized only for CZ fidelity at $\boldsymbol{x}_{\text{in}} = \mathbf{0}$. The memory-robust pulse is parameterized by 20 sine basis functions and is optimized with Eq. (13). The corresponding desired flux pulses are shown in Fig. 3(b).

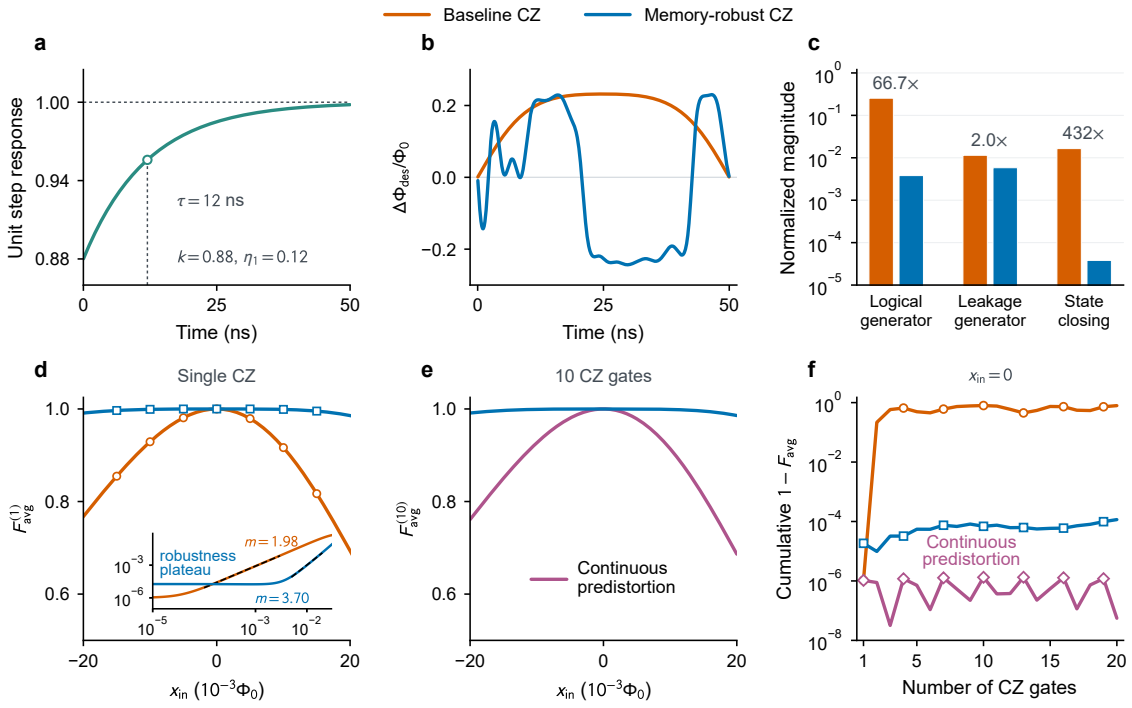


Fig. 3. One-pole validation of memory robustness, fixed-pulse reuse, and sensitivity to continuous-predistortion initialization. (a) Unit step response for $(k, \eta_1, \tau_1) = (0.88, 0.12, 12\ \text{ns})$. (b) Desired flux pulses for the baseline and memory-robust CZ gates, with duration $T = 50$ ns. (c) Normalized magnitudes of the logical and leakage error generators and the state closing residual. Labels above the bars give the reduction factors relative to the baseline. (d) Single-gate average fidelity versus the initial actuator state x_{in}. Open circles (baseline) and squares (memory-robust) are validation points from propagation in the full three-qutrit Hilbert space. The inset plots infidelity $1 - F^{(1)}_{\text{avg}}$ versus x_{in}/Φ_0 for $x_{\text{in}} > 0$ on log-log axes; black dashed lines are power law fits over $x_{\text{in}}/\Phi_0 \in [10^{-4}, 3 \times 10^{-3}]$ for the baseline and $[5 \times 10^{-3}, 2 \times 10^{-2}]$ for the memory-robust pulse. (e) Average fidelity of ten consecutive CZ gates versus the actual initial state x_{in}, comparing the fixed memory-robust input with the baseline under continuous predistortion (purple solid curve). The latter applies a model-based inverse filter across the full sequence, retaining its internal state between gates. For each sequence it assumes a zero initial line state, while the actual x_{in} is scanned. (f) Cumulative infidelity for 20 consecutive CZ gates starting from $x_{\text{in}} = 0$. The baseline and memory-robust methods repeat their respective gate-local predistorted inputs; continuous predistortion uses the baseline target flux sequence with correctly initialized filter state. In (e) and (f), the physical line and quantum dynamics evolve continuously, with no waiting interval or reset between gates; the continuous-predistortion filter is not reset either. Line parameters are exact, and pulse coefficients and nominal virtual-Z calibrations are held fixed throughout the scans. All curves are closed-system simulations.

The memory-robust pulse reduces the logical and leakage generator norms by factors of 66.7 and 2.0, respectively, and the state closing residual by a factor of 432 [Fig. 3(c)]. To characterize the plateau in Fig. 3(d), let $x = x_{\text{in}}$, $\epsilon(x) = 1 - F^{(1)}_{\text{avg}}(x)$, and $\epsilon_0 = \epsilon(0)$. When residual first-order errors and interference with the nominal gate error are negligible, $\epsilon(x) \simeq \epsilon_0 + a_4 x^4$, where $a_4 > 0$. The additional error equals the nominal error floor when

$$a_4 x_{\text{c}}^4 \simeq \epsilon_0, \qquad x_{\text{c}} \simeq \left(\frac{\epsilon_0}{a_4}\right)^{1/4}, \qquad \epsilon(x_{\text{c}}) \simeq 2\epsilon_0.$$

Motivated by this balance, we define the positive-side plateau boundary as the first crossing of $\epsilon(x) = 2\epsilon_0$. Linear interpolation between adjacent scan points gives $x_{\text{c}} \simeq 3.5 \times 10^{-3}\Phi_0$; the negative-side boundary need not have the same magnitude. Beyond this boundary, the robust pulse exhibits near-quartic scaling with $m \simeq 3.70$ over $x/\Phi_0 \in [5 \times 10^{-3}, 2 \times 10^{-2}]$. The baseline instead gives $m \simeq 1.98$ over $[10^{-4}, 3 \times 10^{-3}]$. The plateau boundary does not mark the breakdown of perturbation theory, and the observed near-quartic scaling is restricted to the

fitted interval.

Figure 3(e) tests whether the single-gate protection survives ten consecutive uses of the fixed robust input. Continuous predistortion provides the comparison: its model is exact, but its assumed zero initial state need not match the actual x_{in}. The fixed robust input retains high sequence fidelity over a broader initial-state range. Continuous predistortion gives lower errors near $x_{\mathrm{in}} = 0$, where its initialization is correct, but its fidelity drops more rapidly as the mismatch increases. For this one-pole line, the uncompensated free response is $x_{\mathrm{in}}\mathrm{e}^{-t/\tau_1}$, with t measured from the sequence start. Although this flux error decays, the quantum error accumulated during the early gates need not disappear. Generator suppression reduces that accumulated error; state closing limits the additional memory written by each reused pulse.

Figure 3(f) separates initialization uncertainty from memory generated within a sequence. All three methods start from $x_{\mathrm{in}} = 0$. Repeating the fixed baseline input produces a cumulative infidelity of 0.79 after 20 gates, whereas the fixed robust input keeps it at 1.2×10^{-4}. Correctly initialized continuous predistortion gives still lower errors. It compensates for the evolving line state by changing the AWG input across the sequence; the robust method instead reuses one input while suppressing its quantum sensitivity and driven residual. Together, panels (d)–(f) support fixed-pulse reuse under the tested initial-state uncertainty, not superiority over correctly initialized continuous predistortion.

4.3. Memory-robust CZ gate in a three-pole flux line model

To represent multiple memory time scales, we reduce a reference response reconstructed from the experimental fit parameters reported by Li *et al.* [22] (Tables C1 and C2). We add their short-time exponential distortion terms to a unit step and cascade the resulting transfer function with their long time response model. We approximate its zero-state response to a unit input step by

$$s_3(t) = k + \sum_{\ell=1}^{3} \eta_\ell \left(1 - \mathrm{e}^{-t/\tau_\ell}\right), \qquad t \geq 0,$$

where k is the instantaneous response, η_ℓ is the contribution of each relaxation mode, and τ_ℓ is its time constant. The three-pole form is a reduced model, not a claim that the physical line contains exactly three modes.

We fit all seven parameters by bounded nonlinear least squares, using logarithmic parameterization for positive η_ℓ and τ_ℓ. The fit combines 401 uniformly spaced points over 0–200 ns with 700 logarithmically spaced points over 0.42 ns–40 μs; the mean-square residual of the shorter window receives four times the weight of the wider window. On independent 2001-point grids with the same respective sampling patterns, the RMS errors are 4.04×10^{-5} and 3.34×10^{-4} in units of the input step amplitude. These quantify reduction error relative to the reconstructed reference, not residuals against experimental data. The fitted parameters are

$$\begin{aligned} k &= 0.9528, \\ (\eta_1, \eta_2, \eta_3) &= (0.02487, 0.01290, 0.01955), \\ (\tau_1, \tau_2, \tau_3) &= (18.47\ \mathrm{ns}, 199.4\ \mathrm{ns}, 11.80\ \mu\mathrm{s}). \end{aligned}$$

The short-time section has unit asymptotic gain, whereas the long-time section retains the reported asymptote 1.0127. Neither the complete reference nor its reduction is subsequently normalized to unity, and the fit imposes no DC-gain constraint. The resulting $s_3(\infty) = k + \sum_\ell \eta_\ell \simeq 1.01016$ is retained in both the actuator propagation

and its inverse; it is not a rounding artifact. The displayed parameters are rounded, while the simulations use full precision. The quoted reduction errors apply only to the finite fitting windows, not to the infinite-time limit.

For $\boldsymbol{x} = (x_1, x_2, x_3)^\mathsf{T}$, the corresponding state space matrices in Eqs. (3a) and (3b) are

$$A = -\operatorname{diag}\left(\frac{1}{\tau_1}, \frac{1}{\tau_2}, \frac{1}{\tau_3}\right), \qquad B = \begin{pmatrix} \eta_1/\tau_1 \\ \eta_2/\tau_2 \\ \eta_3/\tau_3 \end{pmatrix},$$

$$C = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}, \qquad D = k,$$

In this diagonal realization, x_ℓ is the stored flux contributed by the ℓth actuator mode. Figure 4(a) shows the unit step response and the three modal contributions. Relative to the gate duration $T = 50$ ns, the free response of the fast mode largely decays within one gate, the intermediate response persists over several gates, and the slow response changes little over the simulated sequence.

Each mode contributes one independent component to the initial actuator state. Equation (13) therefore includes one logical generator term, one leakage generator term, and one state closing term for each mode. We compare the resulting memory-robust pulse with two high-fidelity references under the same quantum and flux line models. The baseline pulse uses 5 sine basis functions and is also shown in Fig. 3(b). The Net-Zero pulse uses 10 even-index sine basis functions and satisfies $\Delta\Phi_{\mathrm{des}}(T - t) = -\Delta\Phi_{\mathrm{des}}(t)$, following the antisymmetric construction of Ref. [12]. Both reference pulses are optimized only for CZ fidelity at $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$ and exceed 99.999%. The memory-robust pulse uses 20 sine basis functions and reaches $F_{\mathrm{avg}} = 99.998\%$. Figure 4(b) compares the three optimized flux pulses.

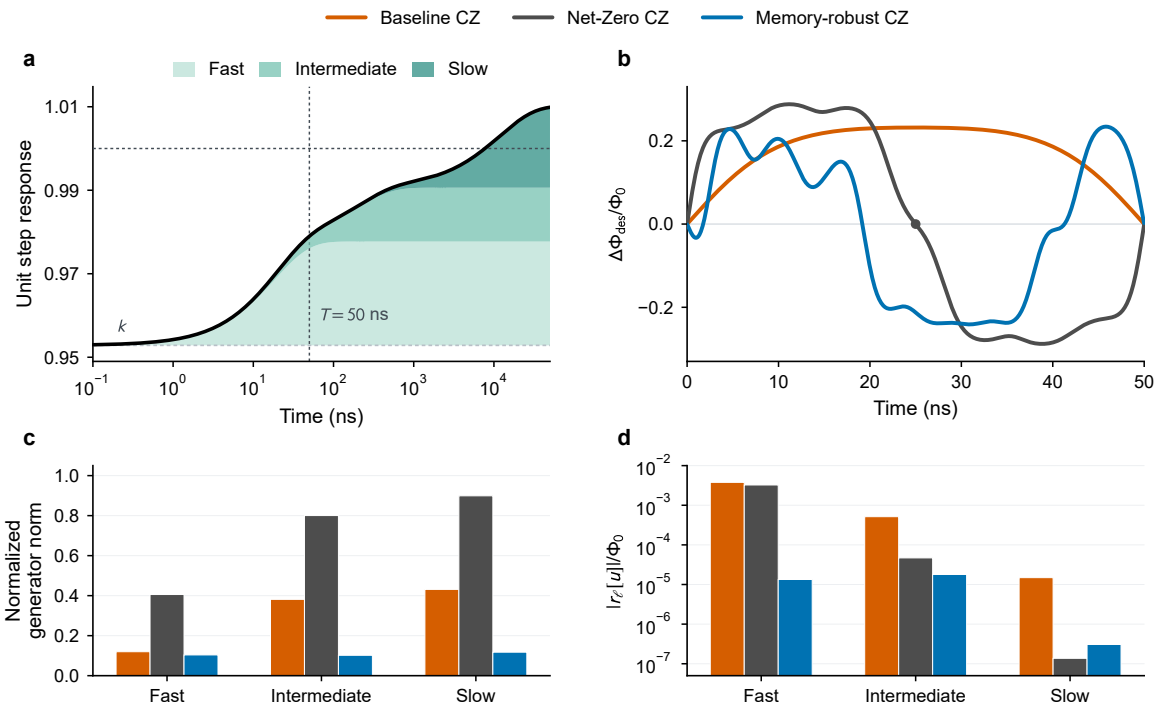


Fig. 4. Memory-robust CZ pulse in a three-pole flux line model. (a) Unit step response of the fitted model (black) and the fast, intermediate, and slow modal contributions (shaded). The horizontal dashed lines mark the direct feedthrough k and unity; the vertical dashed line marks the gate duration $T = 50$ ns. (b) Desired flux pulses for the baseline, Net-Zero, and memory-robust CZ gates. The marker at $t = T/2$ identifies the zero crossing imposed by Net-Zero antisymmetry. (c) The normalized generator norm is $10^{-3}\sqrt{\left\|\mathcal{G}_\ell^{\log}\right\|_{\mathrm{F}}^2 + \left\|\mathcal{G}_\ell^{\mathrm{leak}}\right\|_{\mathrm{F}}^2}$, where $\|\cdot\|_{\mathrm{F}}$ is the Frobenius norm and the factor 10^{-3} corresponds to $x_{\mathrm{in},\ell} = 10^{-3}\Phi_0$. This dimensionless quantity is not normalized to the baseline. (d) Magnitudes of the gate exit actuator state components for $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, normalized by Φ_0.

The memory-robust pulse gives the smallest combined logical and leakage generator norm for each actuator mode [Fig. 4(c)]. Relative to the baseline, the reduction is modest for the fast mode and much larger for the intermediate and slow modes; the Net-Zero pulse gives no comparable suppression. The residual components in Fig. 4(d) correspond to a gate exit actuator state norm of $2.23 \times 10^{-5}\Phi_0$. Thus, the optimization suppresses both the sensitivity to the initial actuator state and the gate exit residual across all three memory time scales within a single 50-ns gate.

4.4. Benchmarks under initial state perturbations and repeated CZ gates

The local generator norms and gate exit residual components in Figs. 4(c) and (d) do not determine the range of initial actuator states over which high fidelity is retained or the performance of repeated CZ gates. We therefore test the three optimized pulses under finite perturbations of the initial actuator state and in consecutive gate sequences. Each fixed-pulse sequence reuses its gate-local predistorted AWG input without an additional waiting interval. Continuous predistortion provides a matched-model reference: the inverse uses the exact line model and the correct zero sequence-start state. This is the three-pole counterpart of Fig. 3(f), whereas Fig. 3(e) tests a mismatch in that initialization. No real-time state measurement or feedback is used in either comparison.

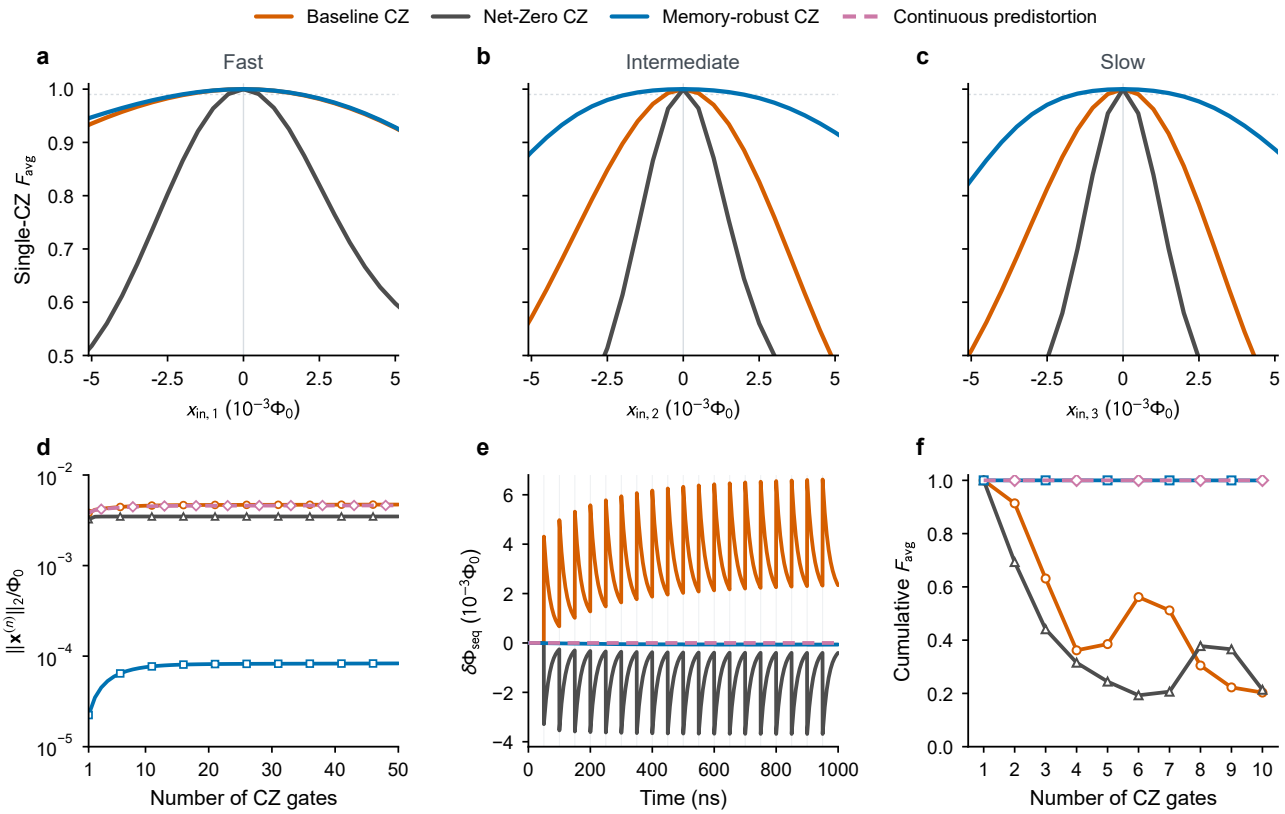


Fig. 5. Validation under finite perturbations of the initial actuator state and repeated CZ gates. (a)–(c) Single-CZ average fidelity versus the fast, intermediate, and slow components of $\boldsymbol{x}_{\mathrm{in}}$, respectively, with the other components set to zero. Horizontal dashed lines mark $F_{\mathrm{avg}} = 0.99$. (d) Actuator state norm after each of 50 consecutive gates starting from $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$. (e) Sequence flux mismatch $\delta\Phi_{\mathrm{seq}}(t) = \Delta\Phi(t) - \Delta\Phi_{\mathrm{des,seq}}(t)$ over 20 gates; vertical lines mark the gate boundaries. The three fixed pulses reuse their respective predistorted inputs, whereas continuous predistortion inverts the full baseline target sequence with the exact line model and correct zero initial state. Its filter state is retained between gates. (f) Cumulative average gate fidelity after each of the first ten gates. In (d)–(f), the line and quantum dynamics evolve continuously without additional waiting or resets; nominal virtual-Z calibrations are held fixed.

The free response associated with the fast component largely decays during a 50-ns gate, so the memory-robust pulse provides only a modest improvement over the baseline for perturbations in $x_{\mathrm{in},1}$. The intermediate and slow responses persist beyond one gate, and the memory-robust pulse retains high fidelity over much wider ranges of $x_{\mathrm{in},2}$ and $x_{\mathrm{in},3}$ than either reference pulse [Figs. 5(a)–(c)]. This behavior agrees with the generator norms in Fig. 4(c).

Unlike the one-pole example, the three-pole results do not exhibit a comparably clear plateau followed by near-quartic scaling. The optimized pulse reduces, but does not eliminate, the first-order generator norms for all three modes while also maintaining nominal fidelity and state closing. Accordingly, the finite-perturbation performance is assessed directly from the scans in Figs. 5(a)–(c), without assuming quartic scaling for each mode.

Figures 5(d)–(f) connect the two gate boundary conditions to sequence fidelity. Starting from $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, all three fixed inputs reconstruct their first desired flux pulse to numerical precision. State closing then keeps the actuator state at the exit of each memory-robust gate below $10^{-4}\Phi_0$ over 50 gates [Fig. 5(d)]. Consequently, its RMS flux mismatch is 59 times smaller than that of the baseline and remains close to zero on the scale of

Fig. 5(e). The continuous predistortion reference also cancels the mismatch, but does so by varying the AWG input according to the model-predicted state. Its small flux mismatch and high sequence fidelity do not require a small physical actuator state [Figs. 5(d)–(f)]. The suppressed error generators of the memory-robust pulse further reduce the quantum error caused by its remaining mismatch. After ten gates, its cumulative fidelity is 99.97%, compared with 20.3% for the baseline and 21.4% for the Net-Zero pulse [Fig. 5(f)], corresponding to an approximately 2.3×10^3 reduction in sequence infidelity relative to both fixed-pulse references, not to continuous predistortion. Because the reference pulses have higher fidelity for an isolated gate at $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, their sequence deterioration arises from actuator memory rather than from the calibration of the first gate.

4.5. Necessity of both gate-boundary conditions

To isolate the role of each boundary condition, we optimize two auxiliary pulses that retain either the six generator penalties or the three state closing penalties, together with the fidelity term. The auxiliary and combined pulses share the gate duration, flux bound, 20 sine basis functions, gate-local predistortion, and initial coefficients.

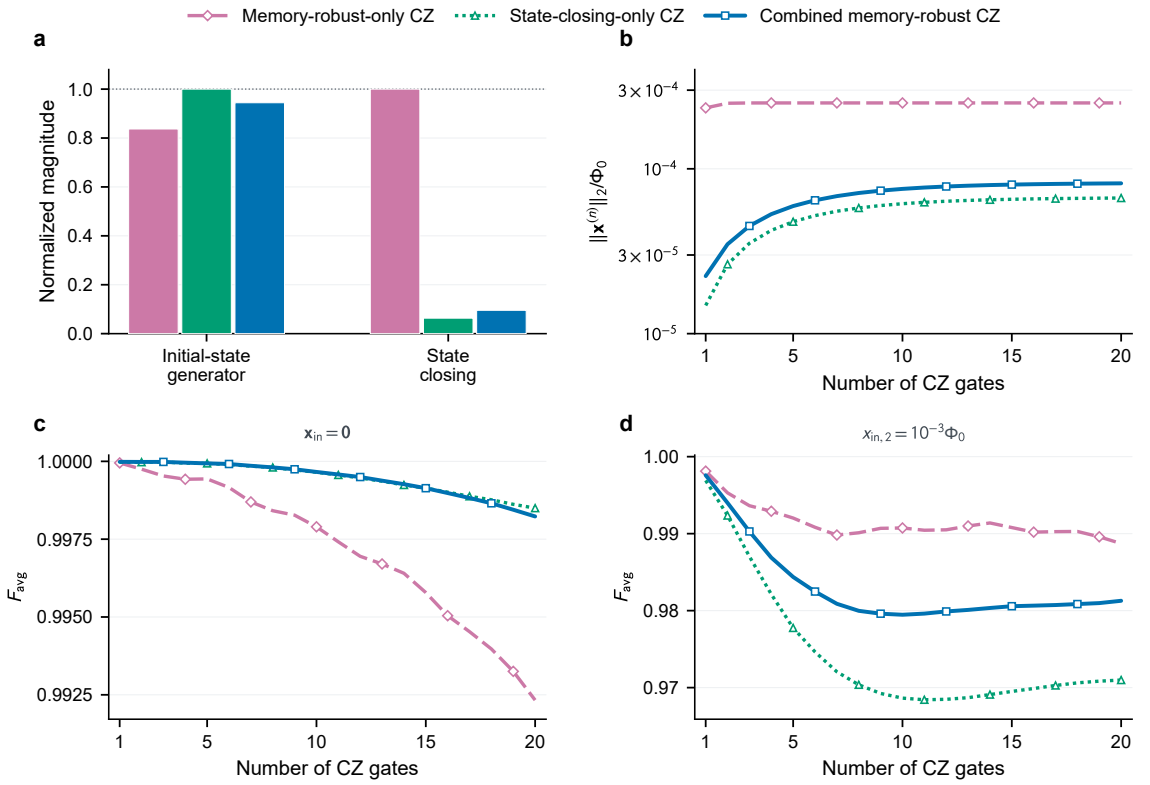


Fig. 6. Comparison of pulses optimized for memory robustness, state closing, or both in the three-pole model. All pulses use a 50-ns duration, a $0.30\Phi_0$ flux bound, 20 sine basis functions, the same initialization, and calibrated gate-local predistortion. At $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, all pulses reach $F_{\mathrm{avg}} > 99.994\%$. (a) Aggregate initial state generator magnitude, normalized to the pulse optimized for state closing alone, and gate exit residual norm, normalized to the pulse optimized for memory robustness alone. (b) Actuator state norm and (c) cumulative average gate fidelity over 20 gates starting from $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$. (d) Cumulative average gate fidelity for $x_{\mathrm{in},2} = 10^{-3}\Phi_0$, with the other components set to zero. Each sequence reuses its fixed predistorted input without an additional waiting interval. The sequence trajectories are excluded from optimization.

Figure 6(a) separates the two conditions. Memory robustness alone minimizes the aggregate generator

magnitude but leaves the largest gate exit residual, whereas state closing alone does the reverse. The combined pulse suppresses both quantities without reaching either single-condition minimum.

Panels (b)–(d) show the consequences for a CZ sequence. At $\boldsymbol{x}_{\mathrm{in}} = \mathbf{0}$, the first gate has no inherited actuator state; the CZ sequence therefore reveals how gate exit residuals generate memory. The two pulses that include state closing maintain smaller boundary states and higher cumulative fidelities than memory robustness alone [Figs. 6(b) and (c)]. At $x_{\mathrm{in},2} = 10^{-3}\Phi_0$, state closing does not suppress the quantum error induced by the inherited intermediate component. Memory robustness alone therefore gives the highest fidelity, while the combined pulse gives an intermediate result [Fig. 6(d)]. Thus, neither condition alone controls both terms in Eq. (10). The combined pulse limits the driven residual while reducing sensitivity to an inherited actuator state.

5. Conclusion and outlook

We treat flux line memory as part of the coupled dynamics of a classical actuator and a quantum system. The resulting first-order analysis yields two complementary requirements at the gate boundaries. The quantum evolution must be first-order insensitive to each component of the initial actuator state, while the control pulse must leave a small driven actuator state at the gate exit. The first condition limits how inherited memory affects the current gate; the second limits the residual actuator state generated by that gate and passed to the next.

The one-pole model exhibits the predicted first-order robustness plateau. Repeating its fixed robust input for ten gates retains high fidelity over a broader initial-state range than continuous predistortion initialized at zero. The latter remains more accurate when the initialization matches the actual state. The three-pole model shows that generator suppression and state closing can be achieved together for several memory time scales. The same gate-local predistorted pulse can then be reused without an additional settling interval: the complete ten-gate sequence retains an average gate fidelity of 99.97%, and its infidelity is smaller by a factor of about 2.3×10^3 than those of the fixed baseline and Net-Zero pulses under the same gate-local protocol. The single-condition comparison further shows that neither generator suppression nor state closing alone controls both routes by which memory affects the sequence.

Within a calibrated finite-dimensional linear model, these results support fixed-pulse reuse with reduced sensitivity to an uncertain initial actuator state. This approach complements continuous predistortion: it suppresses the quantum consequences of a residual flux error rather than requiring exact pulse reconstruction throughout the sequence. It still requires a calibrated line model and gate-local predistortion, and does not establish robustness to arbitrary model errors or schedules. Experimental validation should next determine how model identification errors, drift, nonlinear line response, flux noise, and decoherence affect the achieved sequence fidelity. Experiments with variable gate spacings and interleaved operations would further test whether the protection persists in less regular schedules. The same formulation may also apply to other flux control operations whose actuator response persists beyond the nominal gate window.

Appendix A: Optimization procedure

The coefficients c_m in Eq. (11) are optimized by L-BFGS-B with gradients from automatic differentiation and bounds $-6 \leq c_m \leq 6$. The baseline search starts from the lowest-order sine basis function, whereas the Net-Zero search is restricted to even indices to enforce antisymmetry. The memory-robust searches vary all coefficients and use random vectors or previously converged pulses as initial guesses.

All terms retained in Eq. (13) are active throughout each memory-robust search. Finite initial actuator state perturbations and repeated gate trajectories are used only for validation. The single-condition searches follow the same procedure with either the generator or state closing penalties omitted.

Table A1 lists the weights used in the optimization stages that produced the pulses reported in Figs. 3–6. For the one-pole search, the generator norms are evaluated at $x_{\mathrm{in},1} = 4 \times 10^{-3}\Phi_0$, which also normalizes the gate exit state. The three-pole searches use a scale-compressed form of Eq. (13), with fidelity, generator, and state closing contributions given by $\log[1 + (1 - F_{\mathrm{avg}})/\epsilon_F]$, $\log[1 + (10^{-3}\,\|\mathcal{G}_\ell^s\|_{\mathrm{F}}\,/g_0)^2]$, and $\log[1 + (r_\ell[u]/r_0)^2]$, respectively. Here $s \in \{\log, \mathrm{leak}\}$, $\epsilon_F = 5 \times 10^{-5}$, $g_0 = 0.05$, and $r_0 = 5 \times 10^{-5}\Phi_0$. A fidelity barrier and weak smoothness terms stabilize the search without changing the gate boundary conditions. Weight vector entries are ordered by the fast, intermediate, and slow actuator modes.

Table A1. Loss weights used in the reported pulse optimizations.

Pulse	w_F	$\boldsymbol{w}^{\log}$	$\boldsymbol{w}^{\mathrm{leak}}$	$\boldsymbol{w}^{r}$
One-pole robust	10	20	20	0.1
Three-pole robust	5	$(1, 3, 5)$	$(3, 5, 6)$	$(2, 1, 0.5)$
Memory robustness only	1.256	$(0.165, 0.772, 1.823)$	$(1.238, 1.724, 2.014)$	$(0, 0, 0)$
State closing only	2.5	$(0, 0, 0)$	$(0, 0, 0)$	$(1.167, 2.334, 0.367)$

Table A2 gives the dimensionless coefficients of the primary three-pole memory-robust pulse. Together with Eq. (11), they define the desired flux pulse in Figs. 4–6.

Table A2. Dimensionless coefficients of the three-pole memory-robust CZ pulse.

m	c_m	m	c_m
1	−0.36880843	11	0.01474042
2	0.50662224	12	−0.01618463
3	0.68738861	13	0.06440639
4	−0.27119782	14	0.00600673
5	0.23005490	15	−0.12386559
6	−0.13206655	16	−0.06142061
7	0.19703598	17	−0.08503133
8	0.00661433	18	−0.02531204
9	0.03337576	19	−0.01808988
10	−0.12097157	20	−0.05498015

Appendix B: Derivation of the initial-state error generator

Vary one component $x_{\mathrm{in},\ell}$ of the initial actuator state about zero while keeping all other components fixed at zero. Define the corresponding derivative of the propagator with respect to the dimensionless component $x_{\mathrm{in},\ell}/\Phi_0$ as

$$V_\ell(t) \equiv \Phi_0 \frac{\partial U(t;\boldsymbol{x}_{\mathrm{in}})}{\partial x_{\mathrm{in},\ell}}\bigg|_{\boldsymbol{x}_{\mathrm{in}}=\mathbf{0}}.$$

Let $H_0(t) = H_{\mathrm{QCQ}}[\bar{\Phi}(t)]$, and let $U_0(t)$ denote the corresponding reference propagator. We also define $H_\Phi(t) = \partial_\Phi H_{\mathrm{QCQ}}|_{\bar{\Phi}(t)}$. Using the flux sensitivity $h_\ell(t) = C\mathrm{e}^{At}\boldsymbol{e}_\ell$ from Eq. (5), differentiation of the Schrödinger equation gives

$$\mathrm{i}\dot{V}_\ell(t) = H_0(t)V_\ell(t) + H_\Phi(t)h_\ell(t)U_0(t), \qquad V_\ell(0) = 0. \tag{B1}$$

Transforming Eq. (B1) to the interaction picture and integrating from 0 to T gives

$$\begin{aligned} U_0^\dagger(T)V_\ell(T) &= -\mathrm{i}\mathcal{G}_\ell, \\ \mathcal{G}_\ell &= \int_0^T h_\ell(t)U_0^\dagger(t)H_\Phi(t)U_0(t)\,\mathrm{d}t. \end{aligned} \tag{B2}$$

Equation (B2) identifies $\mathcal{G}_\ell$ as the first-order error generator. Its contribution to the interaction picture propagator is $-\mathrm{i}\mathcal{G}_\ell$. Expanding the propagator in $x_{\mathrm{in},\ell}/\Phi_0$ therefore gives

$$U(T;x_{\mathrm{in},\ell}) = U_0(T)\left[\mathbb{I} - \mathrm{i}\frac{x_{\mathrm{in},\ell}}{\Phi_0}\mathcal{G}_\ell\right] + \mathcal{O}\left[\left(\frac{x_{\mathrm{in},\ell}}{\Phi_0}\right)^2\right], \tag{B3}$$

Therefore, first-order memory robustness requires $\mathcal{G}_\ell^{\log} = 0$ and $\mathcal{G}_\ell^{\mathrm{leak}} = 0$ for every retained component of the initial actuator state, as stated in Eq. (7).

Appendix C: Matched basis one-pole control comparison

To test whether increasing the number of control parameters alone accounts for memory robustness, we optimize an additional fidelity-only reference in the one-pole model. This reference and the unchanged robust pulse both use 20 sine basis functions in Eq. (11), with $T = 50$ ns, $\Phi_{\max} = 0.30\Phi_0$, and coefficient bounds $-6 \le c_m \le 6$. The original five component baseline is retained for comparison. All three pulses use the same line model, gate-local predistortion, and evaluation protocol of Sec. 4.2. Only the nominal isolated gate infidelity is minimized for the new reference; generator, state closing, and smoothness penalties are absent.

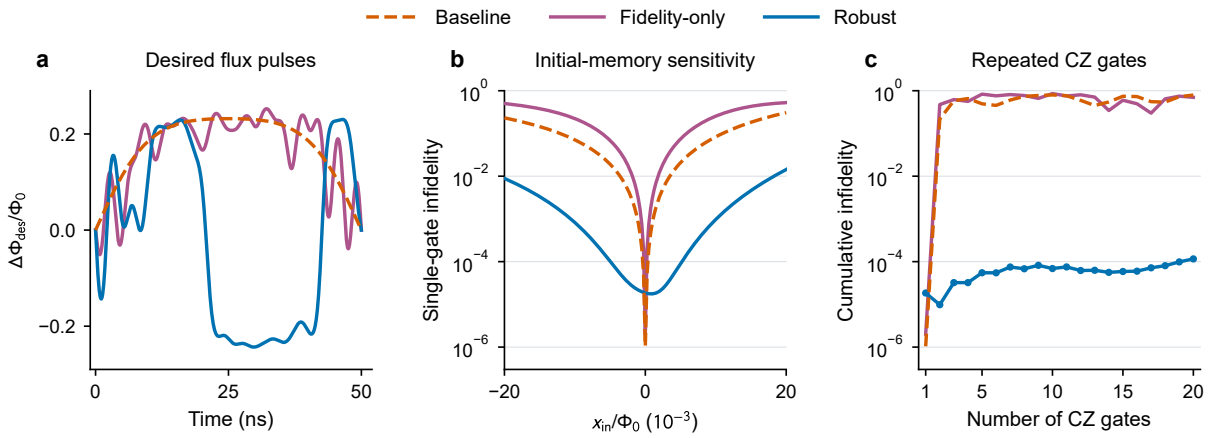


Fig. C1. One-pole comparison at matched control-basis size. The original baseline uses 5 sine components; the new fidelity-only reference and the unchanged memory-robust pulse each use 20. (a) Desired flux pulses. (b) Single-gate infidelity versus the initial actuator state, with fixed nominal virtual-Z calibration. (c) Cumulative infidelity of up to 20 consecutive CZ gates starting from $x_{\mathrm{in}} = 0$. Each sequence repeats its fixed, gate-local predistorted input without waiting, reinitialization, or continuous predistortion. Panels (b) and (c) use logarithmic vertical axes and all results are closed-system simulations.

Figure C1 shows that the new reference reaches a nominal infidelity of 1.99×10^{-6} but remains sensitive to incoming memory. Its zero-initial-state gate exit residual is $1.67 \times 10^{-2}\Phi_0$, compared with $3.81 \times 10^{-5}\Phi_0$ in magnitude for the robust pulse. The corresponding logical and leakage generator norms, evaluated for a reference initial displacement of $10^{-3}\Phi_0$, are $(9.66 \times 10^{-2}, 6.86 \times 10^{-2})$ for the new reference and $(9.51 \times 10^{-4}, 1.46 \times 10^{-3})$ for the robust pulse. After 20 gates, the original baseline, new reference, and robust pulse have cumulative infidelities of 0.792, 0.700, and 1.16×10^{-4}, respectively.

Thus a larger control basis does not by itself guarantee memory robustness or state closing. The matched basis reference attains a lower nominal gate error but does not reproduce the robust pulse's finite memory tolerance or sequence performance. This comparison complements the fixed-basis three-pole condition tests in Fig. 6. It does not establish global optimality or exclude the possibility that fidelity-only optimization could find a pulse with incidental robustness.

Data availability statement

No experimental data were generated in this study. The experimental step response used to parameterize the three-pole flux line model was obtained from Ref. [22]. The coefficients of the primary memory-robust pulse are reported in Appendix A. Additional numerical data supporting the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgment

This work was supported by the Key-Area Research and Development Program of Guangdong Province (Grant No. 2018B030326001), the Science, Technology and Innovation Commission of Shenzhen Municipality (Grant Nos. JCYJ20170412152620376 and KYTDPT20181011104202253), the Innovation Program for

Quantum Science and Technology (2021ZD0301703), Guangdong Major Project of Basic Research (2025B0303000007), the Shenzhen Science and Technology Program (Grant No. KQTD20200820113010023), and CCF-QuantumCtek Superconducting Quantum Computing Special Cooperation Program (No. CCF-QC2025002).

References

[1] Barends R, Kelly J, Megrant A *et al.* 2014 *Nature* **508** 500

[2] Yan F, Krantz P, Sung Y, Kjaergaard M, Campbell D L, Orlando T P, Gustavsson S and Oliver W D 2018 *Phys. Rev. Applied* **10** 054062

[3] Stehlik J, Zajac D M, Underwood D L, Phung T, Blair J, Carnevale S, Klaus D, Keefe G A, Carniol A, Kumph M, Steffen M and Dial O E 2021 *Phys. Rev. Lett.* **127** 080505

[4] McKay D C, Filipp S, Mezzacapo A, Magesan E, Chow J M and Gambetta J M 2016 *Phys. Rev. Applied* **6** 064007

[5] Mundada P, Zhang G, Hazard T and Houck A 2019 *Phys. Rev. Applied* **12** 054023

[6] Sung Y, Ding L, Braumüller J, Vepsäläinen A, Kannan B, Kjaergaard M, Greene A, Samach G O, McNally C, Kim D, Melville A, Niedzielski B M, Schwartz M E, Yoder J L, Orlando T P, Gustavsson S and Oliver W D 2021 *Phys. Rev. X* **11** 021058

[7] Xu Y, Chu J, Yuan J, Qiu J, Zhou Y, Zhang L, Tan X, Yu Y, Liu S, Li J, Yan F and Yu D 2020 *Phys. Rev. Lett.* **125** 240503

[8] Mitchell B K, Naik R K, Morvan A, Hashim A, Kreikebaum J M, Marinelli B, Lavrijsen W, Nowrouzi K, Santiago D I and Siddiqi I 2021 *Phys. Rev. Lett.* **127** 200502

[9] Marxer F, Vepsäläinen A, Jolin S W *et al.* 2023 *PRX Quantum* **4** 010314

[10] Ding L, Hays M, Sung Y, Kannan B, An J, Di Paolo A, Karamlou A H, Hazard T M, Azar K, Kim D K, Niedzielski B M, Melville A, Schwartz M E, Yoder J L, Orlando T P, Gustavsson S, Grover J A, Serniak K and Oliver W D 2023 *Phys. Rev. X* **13** 031035

[11] Ma X, Zhang G, Wu F *et al.* 2024 *Phys. Rev. Lett.* **132** 060602

[12] Rol M A, Battistel F, Malinowski F K, Bultink C C, Tarasinski B M, Vollmer R, Haider N, Muthusubramanian N, Bruno A, Terhal B M and DiCarlo L 2019 *Phys. Rev. Lett.* **123** 120502

[13] Negîrneac V, Ali H, Muthusubramanian N, Battistel F, Sagastizabal R, Moreira M S, Marques J F, Vlothuizen W J, Beekman M, Zachariadis C, Haider N, Bruno A and DiCarlo L 2021 *Phys. Rev. Lett.* **126** 220502

[14] Krinner S, Storz S, Kurpiers P, Magnard P, Heinsoo J, Keller R, Lütolf J, Eichler C and Wallraff A 2019 *EPJ Quantum Technol.* **6** 2

[15] Gustavsson S, Zwier O, Bylander J, Yan F, Yoshihara F, Nakamura Y, Orlando T P and Oliver W D 2013 *Phys. Rev. Lett.* **110** 040502

[16] Foxen B, Mutus J Y, Lucero E *et al.* 2019 *Supercond. Sci. Technol.* **32** 015012

[17] Rol M A, Ciorciaro L, Malinowski F K, Tarasinski B M, Sagastizabal R E, Bultink C C, Salathe Y, Haandbaek N, Sedivy J and DiCarlo L 2020 *Appl. Phys. Lett.* **116** 054001

[18] Hellings C, Lacroix N, Remm A, Boell R, Herrmann J, Lazăr S, Krinner S, Swiadek F, Andersen C K, Eichler C and Wallraff A 2025 *Phys. Rev. Research* **7** 043142

[19] Gao Y Y, Rol M A, Touzard S and Wang C 2021 *PRX Quantum* **2** 040202

[20] Hincks I N, Granade C E, Borneman T W and Cory D G 2015 *Phys. Rev. Applied* **4** 024012

[21] Guo L L, Duan P, Zhang S, Yang X X, Zhang C, Du L, Zhang H F, Tao H R, Wang T L, Jia Z L, Chen Z Y and Guo G P 2024 *Phys. Rev. Applied* **21** 064060

[22] Li T M, Zhang J C, Chen B J *et al.* 2025 *Phys. Rev. Applied* **23** 024059

[23] Kelly J, Barends R, Campbell B *et al.* 2014 *Phys. Rev. Lett.* **112** 240504

[24] Egger D J and Wilhelm F K 2014 *Phys. Rev. Lett.* **112** 240503

[25] Motzoi F, Gambetta J M, Rebentrost P and Wilhelm F K 2009 *Phys. Rev. Lett.* **103** 110501

[26] Khaneja N, Reiss T, Kehlet C, Schulte-Herbrüggen T and Glaser S J 2005 *J. Magn. Reson.* **172** 296

[27] Song Y, Li J, Hai Y J, Guo Q and Deng X H 2022 *Phys. Rev. A* **105** 012616

[28] Brown K R, Harrow A W and Chuang I L 2004 *Phys. Rev. A* **70** 052318

[29] Khodjasteh K and Viola L 2009 *Phys. Rev. Lett.* **102** 080501

[30] Zeng J, Yang C H, Dzurak A S and Barnes E 2019 *Phys. Rev. A* **99** 052321

[31] Buterakos D, Das Sarma S and Barnes E 2021 *PRX Quantum* **2** 010341

[32] Hai Y J, Song Y, Li J, Zeng J and Deng X H 2025 *Phys. Rev. Applied* **23** 054002

[33] Green T J, Sastrawan J, Uys H and Biercuk M J 2013 *New J. Phys.* **15** 095004

[34] Soare A, Ball H, Hayes D, Sastrawan J, Jarratt M C, McLoughlin J J, Zhen X, Green T J and Biercuk M J 2014 *Nat. Phys.* **10** 825

[35] Le I N M, Teske J D, Hangleiter T, Cerfontaine P and Bluhm H 2022 *Phys. Rev. Applied* **17** 024006

[36] Propson T, Jackson B E, Koch J, Manchester Z and Schuster D I 2022 *Phys. Rev. Applied* **17** 014036

[37] Poggi P M, De Chiara G, Campbell S and Kiely A 2024 *Phys. Rev. Lett.* **132** 193801

[38] Rasulov U and Kuprov I 2025 *J. Chem. Phys.* **162** 164107

[39] Koch J, Yu T M, Gambetta J, Houck A A, Schuster D I, Majer J, Blais A, Devoret M H, Girvin S M and Schoelkopf R J 2007 *Phys. Rev. A* **76** 042319

[40] Hutchings M D, Hertzberg J B, Liu Y, Bronn N T, Keefe G A, Brink M, Chow J M and Plourde B L T 2017 *Phys. Rev. Applied* **8** 044003

[41] Åström K J and Murray R M 2008 *Feedback Systems: An Introduction for Scientists and Engineers* (Princeton, NJ: Princeton University Press)

[42] Ljung L 1999 *System Identification: Theory for the User* 2nd ed (Upper Saddle River, NJ: Prentice Hall PTR)

[43] Dyson F J 1949 *Phys. Rev.* **75** 486

[44] Morales J L and Nocedal J 2011 *ACM Trans. Math. Softw.* **38** 1

[45] Virtanen P, Gommers R, Oliphant T E *et al.* 2020 *Nat. Methods* **17** 261