Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Reachability of complete quantum operations under dynamical error correction

Xiu-Hao Deng

dengxiuhao@iqasz.cn

Shenzhen International Quantum Academy

Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11814682/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Reachability of complete quantum operations under dynamical error correction

Xiu-Hao Deng[1,2]*

[1]International Quantum Academy, Shenzhen 518048, China

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

Abstract

Which complete quantum operations remain reachable when control errors in multi-qubit systems are cancelled to a prescribed order? We develop a manifold framework mapping corrected control sets to reachable full-register operations and distinguishing their logical action on a computational subspace. It covers separately corrected modules, joint cancellation across sequences and permitted subspace excursions. For signed, drift-free unitary controls with shared quasistatic errors and primitive-dependent linear gains, we classify the reachable set and prove its invariance across positive finite orders. In collective spin dynamics, relative phases distinguish correctable extensions of the same logical gate. Local fibers at regular controls describe implementation freedom. For a three-qubit target, we certify a continuous family of implementations with worst-case channel error below 0.01 throughout $\pm 1\%$ gain offsets, including bounded control perturbations and local dephasing. In four- and five-qubit simulations, first-order initialization lowers maximum sampled depth-32 errors in seven of eight matched comparisons with full-range training.

Introduction

A logical gate specifies an action on a computational subspace, while its control implementation evolves the full physical register. Error cancellation can be required at individual module boundaries or only after a complete sequence. When the available controls permit, the intermediate evolution can also leave and return to the computational subspace. These choices determine the constraints on the implementation. Which complete operations remain reachable under those constraints as the cancellation order increases?

A target is *nominally reachable* if an allowed control sequence implements it exactly at zero error. It is *reachable at order q* if one such open-loop sequence also cancels every channel derivative of total order one through q at zero error, including mixed derivatives with respect to different error parameters. The sequence is fixed independently of the unknown errors. A *complete operation* specifies the channel on the entire register, including inputs entangled with a reference. If the controls preserve several sectors, this specification fixes the relative phases that a superposition across sectors can reveal.

Composite pulses compensate amplitude and detuning errors by combining imperfect rotations [1, 2], with constructions extending cancellation to arbitrary prescribed order [3]. Dynamical decoupling suppresses unwanted interactions through controlled evolution [4]; dynamically corrected gates combine this protection with nontrivial operations [5–7]. Related compensation sequences address shared gain errors for specific multi-qubit generators [8].

Lie-algebraic ensemble control addresses gain dispersion [9–11], while robust optimization balances finite-range performance against control resources [12]. Geometric constructions represent accumulated coherent errors as curves whose closure enforces leading-order cancellation [13, 14]. Related work treats multiqubit control [15], correlated noise in superconducting circuits [16], and deterministic error suppression across circuit layers [17]. Hao and Bukov give first-order reachability criteria for static and selected waveform perturbations, including uniform Hamiltonian rescaling under full unitary control [18].

*Correspondence: dengxiuhao@iqasz.cn

We organize these questions through the map from corrected control sets to complete physical operations and their logical restrictions (Fig. 1). The task and completion boundary determine the response constraints; their solution set determines which physical targets can be implemented. Near a regular control, manifold coordinates separate changes of target from changes of implementation at fixed target. This framework accommodates separate module correction, joint cancellation across a sequence and subspace excursions permitted by the available generators.

For drift-free controls with signed pulse areas, shared quasistatic errors and primitive-dependent linear gain profiles, we classify the exact full-register targets at every positive finite order. Commutator-generated operations remain reachable, and a linear condition on pulse areas selects the additional central phases. Increasing the order restricts the controls without shrinking their target image; the required sequences and resources can depend on the order. In collective spin dynamics, an observable relative phase distinguishes correctable and excluded extensions of the same logical gate. We then quantify implementation freedom and finite-error performance through local charts, a certificate for a continuous implementation family and four- and five-qubit circuit simulations.

Results

Control paths, completion boundaries and target maps

Let a specify a control sequence or waveform chosen independently of the unknown error vector $\boldsymbol{\epsilon}$. For a complete unitary operation, let $\mathcal{M}_q$ contain controls whose channel responses vanish through order $q \geq 1$, and let $\pi(a) = U_0(a)$ be the nominal endpoint map, interpreted as a channel up to one common scalar. The complete reachable set is $\mathcal{K}_q = \pi(\mathcal{M}_q)$. Thus $U \in \mathcal{K}_q$ means that at least one allowed control implements U with the prescribed cancellation. Controls sharing a complete target form its fixed-target fiber; regularity conditions below give this fiber a local manifold structure.

Figure 1(a) also resolves the logical action on a computational subspace with projector P. If the nominal endpoint preserves P, a representative has the form $U = V \oplus W$. Fixing the logical action leaves the complementary action and its relative phase undetermined; fixing the complete target specifies both. The surfaces in the figure describe sets of controls. The subspace selected by P lies in the physical Hilbert space, and the dashed arrows map controls to completed operations.

Panel (b) distinguishes the correction boundary from the execution path. Separate correction requires cancellation at each module boundary, while joint correction requires cancellation only after the constituent responses combine. An execution can also leave P and return at completion when the available controls couple it to its complement. These choices impose different implementation constraints. Their effect on the reachable target image must be determined for the declared control alphabet and error model.

For a general task $\mathcal{I}$, specify the allowed inputs, the output to be preserved, the completion boundary, the allowed continuations, the error correlations and how leakage is counted. Its response $O_{\mathcal{I}}(a, \boldsymbol{\epsilon})$ takes values in a finite-dimensional normed linear space. For a complete operation it is the physical channel. Restricting the inputs to P still retains the physical output in its entirety, including leakage; a specified readout can define a weaker task with different continuation requirements.

Suppose $O_{\mathcal{I}}$ is jointly C^{q+1} near the nominal error. Collect the nonconstant Taylor coefficients through total order q in $R_{\mathcal{I},\leq q}(a)$. The controls that cancel those coefficients and their target map are

$$\begin{aligned} \mathcal{R}_{\mathcal{I},q} &= \{a \in C : R_{\mathcal{I},\leq q}(a) = 0\}, \\ \pi_{\mathcal{I},q}(a) &= O_{\mathcal{I}}(a, 0), \qquad a \in \mathcal{R}_{\mathcal{I},q}. \end{aligned} \tag{1}$$

Here C denotes the chosen controls. Each fixed word or waveform template can be parameterized in finitely many real variables.

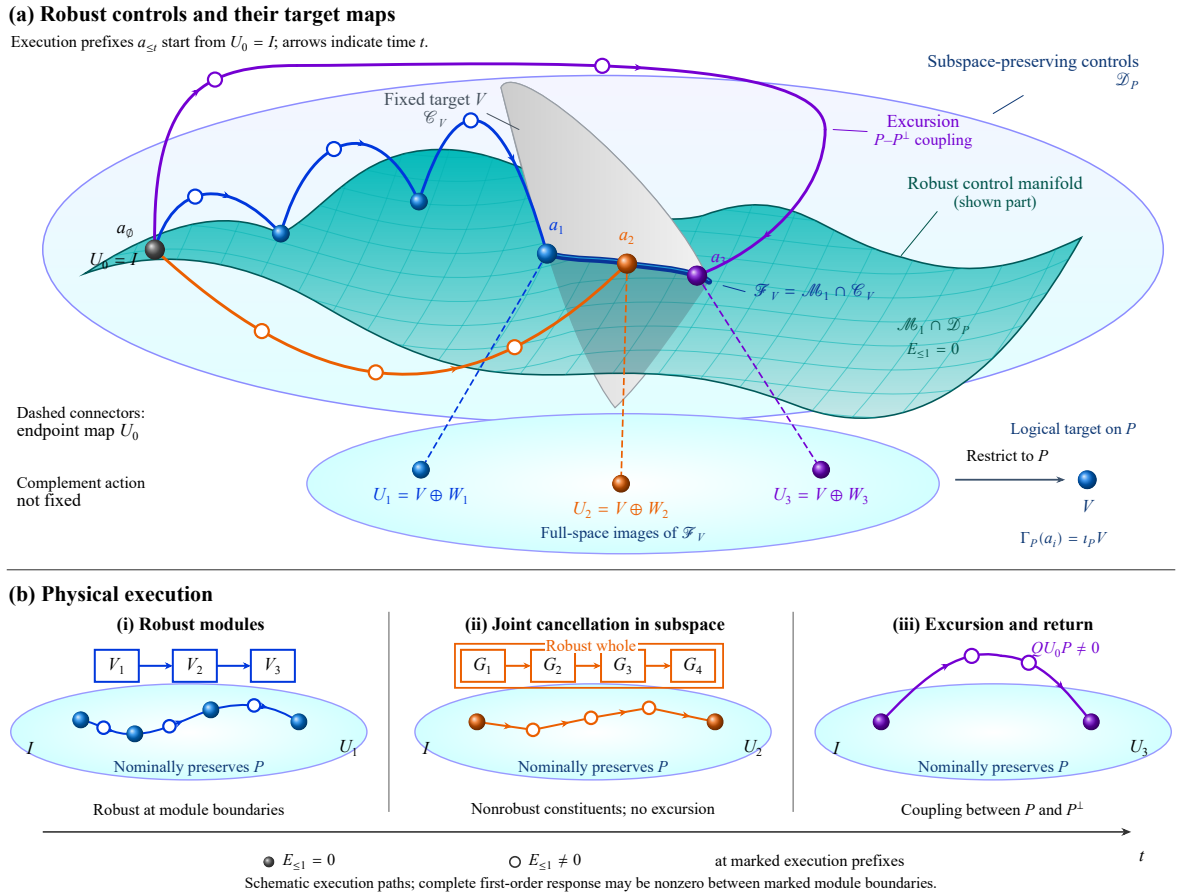


Figure 1: **From control paths to complete and logical operations.** Corrected controls map to complete physical operations and, at subspace-preserving endpoints, to logical actions. The execution paths indicate where correction is required. **(a)** Paths follow controls stopped at successive physical times. The set $\mathcal{D}_P$ contains controls whose nominal endpoints preserve the subspace P; $\mathcal{C}_V$ fixes their action on P to V. Its intersection $\mathcal{F}_V = \mathcal{M}_1 \cap \mathcal{C}_V$ contains controls with that logical action and zero complete first-order response, denoted $E_{\le 1} = 0$. The three completed controls can have different full endpoints $U_i = V \oplus W_i$. Dashed connectors denote the endpoint map. Restriction to inputs in P gives $\Gamma_P(a_i) = U_i \iota_P = \iota_P V$, where ι_P embeds the subspace. A fixed complete target also specifies W_i. **(b)** Blue: each completed module is corrected. Orange: constituent responses cancel jointly while the nominal evolution preserves P. Purple: the nominal evolution leaves P and returns at completion. Filled and open markers indicate zero and nonzero accumulated responses at selected prefixes, respectively. Paths denote execution time, not continuous redesign of a completed control. Purple requires coupling between P and its complement, which can occur within a conserved spin sector. The four- and five-qubit encoding below permits such motion under collective rotations. Supplementary Note 13 gives a separate exact example with an enlarged control alphabet. All surfaces and paths are schematic; smooth local geometry requires the regularity conditions stated below.

Definition 1 (Finite-order dynamical error correction)**.** An open-loop control a dynamically corrects a target $\mathcal{V}_\mathcal{I}$ through order q when $a \in \mathcal{R}_{\mathcal{I},q}$ and $\pi_{\mathcal{I},q}(a) = \mathcal{V}_\mathcal{I}$.

The guarantee refers to the declared completion boundary [5–7]. The definition also applies to coherent waveform perturbations with specified temporal profiles and to smooth families of open-system channels. The reachability theorem below concerns complete unitary channels with the same unknown gain vector throughout a sequence; this vector can change between runs.

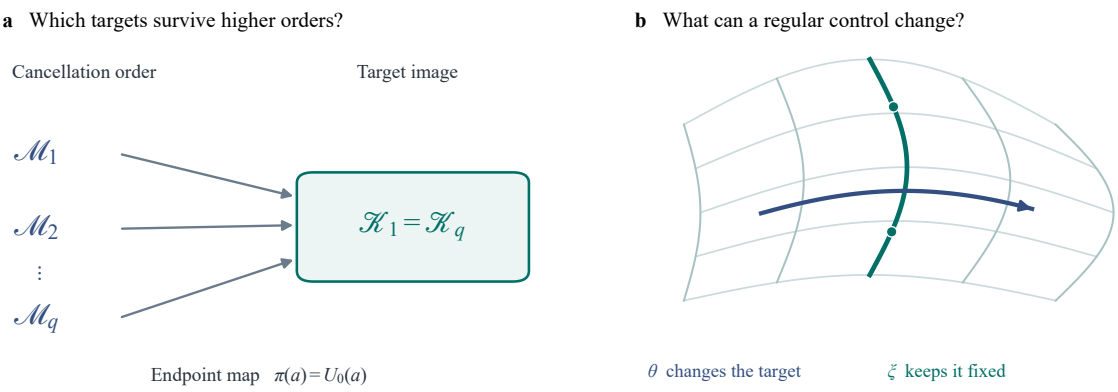


Figure 2: **Cancellation order and the nominal target map. (a)** The nested sets $\mathcal{M}_1 \supseteq \mathcal{M}_2 \supseteq \mathcal{M}_q$ for $q \geq 2$ contain finite signed words with progressively more vanishing derivatives. Under the signed, drift-free linear-gain model, each set maps onto the same reachable targets, $\mathcal{K}_q = \mathcal{K}_1$ for every finite $q \geq 1$. Different sequences may realize the same target at different orders. **(b)** Near a regular control in a fixed template, varying the local coordinate θ programs nearby targets. Varying ξ at fixed θ changes the implementation while preserving the target and cancellation. Both panels are schematic.

For a nominal endpoint U_0, the Hermitian input-frame responses are

$$K_\mu = iU_0^\dagger \partial_{\epsilon_\mu} U\big|_0. \tag{2}$$

Complete first-order channel insensitivity requires each K_μ to be scalar, hence zero for the traceless models considered here. For $U_c = G_m \cdots G_1$ and $S_{j-1} = G_{j-1,0} \cdots G_{1,0}$, the shared-error response is

$$K_{c,\mu} = \sum_{j=1}^{m} S_{j-1}^\dagger K_{j,\mu} S_{j-1}. \tag{3}$$

Separate first-order correction sets each $K_{j,\mu}$ to zero. Joint correction allows nonzero constituent responses whose transported sum vanishes. This is the mechanism behind the blue and orange executions in Fig. 1(b). Neither condition requires vanishing response at every intermediate time. The finite-depth consequences are developed below.

Increasing the cancellation order gives nested control sets $\mathcal{M}_{q+1} \subseteq \mathcal{M}_q$. Their endpoint images need not shrink: the theorem below proves that they coincide under the specified linear-gain model [Fig. 2(a)]. Figure 2(b) describes local variations among completed implementations, distinct from the physical execution paths in Fig. 1.

Reachable targets at every finite cancellation order

We first classify the complete endpoints in Fig. 1(a). The allowed controls are all finite sequences of the specified primitives with signed pulse areas (finite signed words), chosen independently of the unknown gains. Figure 2(a) asks whether increasing the cancellation order changes their target image.

Theorem 1 (Exact targets for linear multiplicative errors)**.** *Take a fixed finite family of traceless anti-Hermitian generators X_a in a drift-free system. In every occurrence of primitive a, let the generator be*

$$X_a(\boldsymbol{\epsilon}) = (1 + \nu_a \cdot \boldsymbol{\epsilon}) X_a, \qquad \nu_a \in \mathbb{R}^r, \tag{4}$$

for the same quasistatic error vector $\boldsymbol{\epsilon}$ and fixed profile ν_a. Real pulse areas of either sign are allowed. Write the actual nominal Lie algebra as $\mathfrak{g} = \mathrm{Lie}\{X_a\} = \mathfrak{z} \oplus \bigoplus_j \mathfrak{s}_j$, with center $\mathfrak{z}$ and simple factors $\mathfrak{s}_j$. Let Z_a be the central projection of X_a and let S_{der} be the channel image of the connected semisimple subgroup. On vectors of signed primitive areas $t = (t_a)$, set

$$C_0 t = \sum_a t_a Z_a, \qquad C_\mu t = \sum_a t_a \nu_{a\mu} Z_a, \qquad E = C_0\Big(\bigcap_{\mu=1}^{r} \ker C_\mu\Big).$$

For every finite $q \geq 1$, the exact complete-channel target set is

$$\boxed{\mathcal{K}_q = \mathcal{K}_1 = \exp(E)\, S_{\mathrm{der}}.} \tag{5}$$

The product is taken modulo one common global phase. Every target reachable at first order has an implementing word at each prescribed finite order. Its pulse areas and length may depend on the target and the order.

The central phases add under concatenation, so $C_\mu t = 0$ cancels their calibration dependence, while $C_0 t$ supplies the nominal central phase. On a simple factor, commutators generate independent motion in the nominal and error coordinates. To make this precise, let W_j span the profiles ν_a for which X_a has nonzero projection onto $\mathfrak{s}_j$. Let $\mathcal{A}_q(W_j)$ be the polynomials in those linear error forms truncated above degree q. The endpoint and all its Taylor coefficients through degree q form an order-q jet. Bracketing jets multiplies their error polynomials and brackets their generators. The jointly generated algebra is

$$\mathfrak{L}^{(q)} = \mathrm{span}_{\mathbb{R}}\{Z_a(1 + \nu_a \cdot \boldsymbol{\epsilon})\} \;\oplus\; \bigoplus_j \mathfrak{s}_j \otimes \mathcal{A}_q(W_j). \tag{6}$$

Independent constant semisimple terms realize all the semisimple targets. Supplementary Note 1 proves the direct-sum identity and constructs exact finite words by finding an open identity neighborhood in the generated group's intrinsic topology. The group in Eq. (5) is immersed and can be nonclosed. An exact exclusion can therefore coexist with approximation to arbitrary accuracy in the ambient channel space.

The linear profile assumption is essential to the order independence of the central test. For two commuting qubit primitives along $Z = -i\sigma_z/2$ with gains $1 + \epsilon$ and $1 + 2\epsilon + \epsilon^2$, the areas $(2a, -a)$ give $\exp[a(1 - \epsilon^2)Z]$. All phases $\exp(aZ)$ are reachable at first order. At second order, the quadratic condition sets the second area to zero and the linear condition sets the first to zero, so $\mathcal{K}_2 = \{I\}$. With two strictly linear gains $1 + \epsilon$ and $1 + 2\epsilon$, the same areas yield the error-independent phase $\exp(aZ)$.

For a fixed analytic family valued in its nominal algebra, the semisimple conclusion extends if the projected first-order response algebra is the full direct sum of independent nominal and linear error directions across all simple factors. The linear directions must span every error variable that occurs on the corresponding factor. Under this hypothesis, Lie brackets generate the higher-order coefficients. The central phases still require vanishing Taylor coefficients through the chosen order (Supplementary Note 2). A common unknown conjugation of all primitives illustrates how nominal controllability alone can leave the first-order condition unsatisfied.

Collective spin control

For $N \geq 2$ qubits and $\hbar = 1$, define

$$J_a = \frac{1}{2} \sum_{r=1}^{N} \sigma_a^{(r)}, \qquad H_{ZZ} = \sum_{r<s} Z_r Z_s = 2J_z^2 - \frac{N}{2} I. \tag{7}$$

Our principal model is

$$H(t;x,y) = (1+x)[u_x(t)J_x + u_y(t)J_y] + (1+y)g(t)H_{ZZ}, \tag{8}$$

where $x = \epsilon_R$ and $y = \epsilon_T$ are shared quasistatic fractional calibration errors in the drive and interaction strengths, respectively. Signed areas are realized with either amplitude sign and positive physical durations. The structural reachability analysis uses

$$X = -iJ_x, \qquad Y = -iJ_y, \qquad Z = -iH_{ZZ}, \tag{9}$$

without a prescribed time or bandwidth budget. The CNOT study later imposes amplitude and response-time limits. Here X, Y, Z denote collective generators, and Z_r is a single-qubit Pauli operator.

The total-spin decomposition is

$$(\mathbb{C}^2)^{\otimes N} \cong \bigoplus_{j\in\mathcal{J}_N} (V_j \otimes M_j), \qquad d_j = \dim V_j = 2j+1, \tag{10}$$

with $j = j_{\min}, j_{\min}+1, \ldots, N/2$, where $j_{\min} = 0$ for even N and $1/2$ for odd N. Collective controls act identically on the multiplicity copies M_j. The complete channel acts on the full register, including arbitrary inputs correlated with a reference. A superposition spanning two spin sectors can resolve their relative phase.

Symmetry and locality can restrict nominal unitary reachability [19], and invariant-subspace controllability has been studied for collective spin networks [20]. For the present model, the nominal controls generate independent special-unitary dynamics on the nonzero spin sectors and one additional sector-scalar direction [21]. The latter comes from twisting and carries the same first-order interaction gain as the twisting primitive. Theorem 1 removes that central direction under cancellation and gives the following closed group of robust matrix representatives.

$$\mathcal{K}_N^{\mathrm{col}} = \prod_{\substack{j\in\mathcal{J}_N \\ j\geq 1/2}} SU(d_j). \tag{11}$$

Each factor acts identically on all multiplicity copies M_j, and the spin-zero block is fixed to the identity. The robust channel group is the quotient of $\mathcal{K}_N^{\mathrm{col}}$ by its common scalar elements. These elements are $\{I\}$ for even N, by the spin-zero block, and $\{\pm I\}$ for odd N, when every d_j is even.

Theorem 2 (Collective finite-order robust reachability)**.** *For every $N \geq 2$, every fixed finite $q \geq 1$, and every target in $\mathcal{K}_N^{\mathrm{col}}$, there is a finite signed rectangular word with that nominal endpoint and with all mixed calibration-error derivatives through total order q equal to zero. Conversely, every nominal endpoint with those complete-interface derivatives zero lies in $\mathcal{K}_N^{\mathrm{col}}$.*

The pure constant generators in Eq. (S11) give the semisimple targets. The sector-trace condition of Lemma S1 makes every robust word have zero signed twisting area, including words that wind around the nominal phase.

An observable phase selects the correctable gates

The distinction between a logical action and a complete endpoint in Fig. 1(a) becomes observable under collective spin dynamics. Two gates can agree on every encoded input, yet differ in whether the complete operation admits dynamical correction. To see the difference, encode two logical qubits in the symmetric sector of three physical qubits, in the ordered basis $|111\rangle, |D_2\rangle, |D_1\rangle, |000\rangle$. Here $|D_k\rangle$ is the normalized Dicke state with k ones. Write V_D for the embedding, $P = V_D V_D^\dagger$ and $Q = I - P$. The family

$$U_\phi = e^{i\phi} V_D \mathrm{CNOT}_L V_D^\dagger + Q, \qquad U_{\mathrm{C}} = -U_{\pi/4} \tag{12}$$

acts identically on all inputs supported on P, including those entangled with a reference. In fact, any input without coherence between P and Q loses all sensitivity to ϕ, even if both sectors are populated. A complete operation also acts on coherent superpositions across the two sectors. Both the encoded and complete tasks count every output, including leakage from an encoded input.

For the controls of Eq. (8), every U_ϕ is nominally reachable. The sector-mean part $P - Q$ of the interaction lies in the nominal algebra, and $e^{-it(P-Q)}U_{\pi/4} = e^{it}U_{\pi/4-2t}$. Cancellation changes this freedom. The corrected target group is $SU(4) \times SU(2)$, with the same $SU(2)$ matrix on both copies of the spin-1/2 sector. One common phase $e^{i\alpha}$ must make both blocks special unitary. Since $\det(\mathrm{CNOT}_L) = -1$, this requires

$$e^{2i\alpha} = 1, \qquad -e^{4i(\alpha+\phi)} = 1 \quad \Longleftrightarrow \quad \alpha \in \pi\mathbb{Z}, \quad \phi \in \frac{\pi}{4} + \frac{\pi}{2}\mathbb{Z}. \tag{13}$$

Thus $U_{\pi/4}$ is reachable at every prescribed finite order, whereas $U_{\phi=0}$ already fails the first-order test. The four allowed phases per period all implement the same encoded gate. The sign in U_C fixes the representative for the numerical controls. The same allowed phases occur for an embedding in any register with $N \geq 3$ (Supplementary Note 6.1).

The distance to the corrected target group can also be found exactly on three qubits. Write $D(\mathcal{A}, \mathcal{B}) = \frac{1}{2}\|\mathcal{A} - \mathcal{B}\|_\diamond$. Operationally, D is the largest change in the probability of a measurement outcome over all inputs, references and output measurements. Let $\delta(\phi)$ be the angular distance to the nearest allowed phase. For any robust endpoint V, the means of the eigenangles in the two spin blocks of $U_\phi^\dagger V$ differ by at least $\delta(\phi)$ when lifted to a common spectral arc. A nearest allowed U_{ϕ_*} attains that arc width, giving

$$\min_{V \in \mathcal{K}_3^{\mathrm{col}}} D(\mathcal{U}_\phi, \mathcal{V}) = \sin\left(\frac{\delta(\phi)}{2}\right), \qquad \min_{V \in \mathcal{K}_3^{\mathrm{col}}} D(\mathcal{U}_0, \mathcal{V}) = \sin(\pi/8) \simeq 0.383. \tag{14}$$

This is the separation between ideal target channels and the exactly corrected group. The next application bounds the error of an implemented pulse.

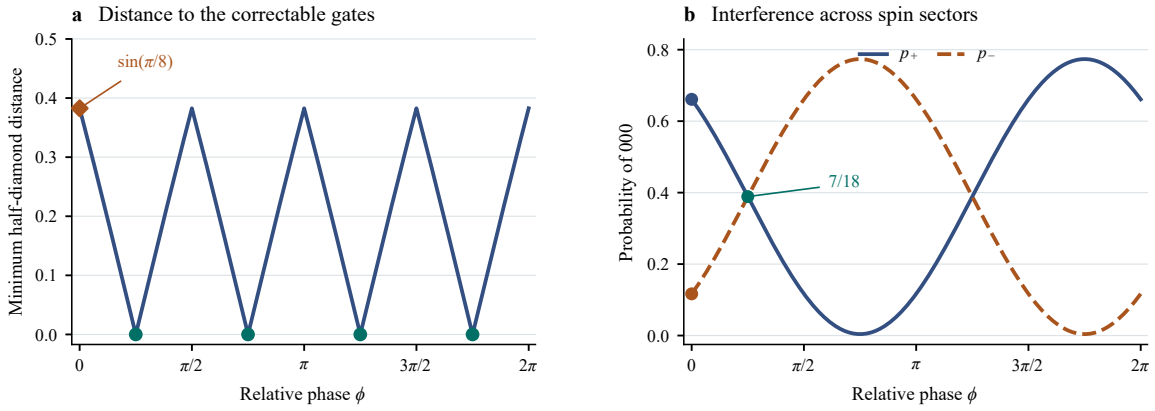


Figure 3: **A relative sector phase changes which complete gates admit correction. (a)** The minimum complete-channel half-diamond distance from U_ϕ to the corrected group on three qubits. Four zeros mark the allowed phases. At $\phi = 0$, the same logical CNOT has minimum distance $\sin(\pi/8)$ to the corrected group. **(b)** An addressed readout converts coherence between the symmetric sector and its complement into the probability of 000. The curves give the two analysis settings in Eq. (15). At $\phi = 0$ the probabilities are 0.661 and 0.117. At $\phi = \pi/4$ both equal $7/18$.

The phase difference has a direct experimental signature. Preparing $|100\rangle$ creates coherence across the two sectors. After the collective operation, two addressed analysis pulses with different phases convert that

Table 1: Complete-channel membership in the signed collective model. An allowed target has an exact realization with all error derivatives cancelled through any prescribed finite order. The stored encoded CNOT controls use the representative $U_C = -U_{\pi/4}$ of Eq. (12).

Register	Target representative	Complete channel	Logical action on P
Even N	$U^{\mathrm{MS}}_{\pi/2}$	Excluded	–
Odd $N \geq 3$	$e^{i\pi/8}U^{\mathrm{MS}}_{\pi/2}$	Allowed	–
Physical $N = 2$	CNOT local class	Excluded	–
Physical $N = 2$	i iSWAP	Allowed	–
Encoded $N = 3$	$\mathrm{CNOT}_L \oplus I_Q$	Excluded	CNOT_L
Encoded $N = 3$	$e^{i\pi/4}\mathrm{CNOT}_L \oplus I_Q$	Allowed	CNOT_L

coherence into unconditional probabilities p_+ and p_- of reading 000. Their contrast is

$$p_- - p_+ = \frac{4}{3\sqrt{3}} \sin\left(\phi - \frac{\pi}{4}\right). \tag{15}$$

For $\phi = 0$ and $\pi/4$ the predicted contrasts are $-4/(3\sqrt{6})$ and zero. Under ideal preparation and independent shots, 41 shots per setting suffice to distinguish these two hypotheses with a worst-case error probability below 5%. Resolving a contrast to 10^{-3} requires a different precision budget. Methods and Supplementary Note 8.5 give the pulses, estimator and bias assumptions.

Membership of standard gates

The same test applies to other collective targets. For a symmetry-compatible unitary $U = \bigoplus_j (U_j \otimes I_{M_j})$, membership requires one common phase α that makes every $e^{i\alpha}U_j$ special unitary, with the spin-zero block equal to the identity when present. For the Mølmer–Sørensen target $U^{\mathrm{MS}}_\theta = e^{-i\theta J_x^2}$ [22],

$$\det(e^{i\alpha}U^{\mathrm{MS}}_{\theta,j}) = \exp\left[id_j\alpha - i\theta\frac{d_j j(j+1)}{3}\right]. \tag{16}$$

Corollary 1 (Standard gates under collective control)**.** *At each prescribed finite order in model* (8)*, the channel of* U^{MS}_θ *is robustly reachable exactly for* $\theta \in \pi\mathbb{Z}$ *at even* $N \geq 2$ *and* $\theta \in (\pi/2)\mathbb{Z}$ *at odd* $N \geq 3$*. The allowed even-N members are collective rotations. For two physical qubits, the robustly reachable channel group excludes the full CNOT local-equivalence class and includes* i iSWAP*.*

The lowest two spin sectors give the necessary conditions on θ. Supplementary Note 6 checks every higher sector and derives the CNOT exclusion using two-qubit invariants [23]. Table 1 compares these complete targets with the encoded example. The test is for the effective spin Hamiltonian in Eq. (8). Robust trapped-ion constructions can use additional motional controls or oscillator dynamics [24–28].

Consequences of restricting interaction signs

The sector-scalar part of twisting fixes the signed area. Write $Z = Q + C$, with

$$Q = -2i(J_z^2 - J^2/3), \qquad C = -i(2J^2/3 - NI/2). \tag{17}$$

Here Q is the sector-traceless quadrupole and C generates the sector phases. For the input-frame interaction response $K_{T,j} = iU^\dagger_{0,j}\partial_y U_j|_0$, full cancellation implies

$$\begin{aligned} \mathrm{tr}\, K_{T,j} &= d_j\gamma_j S, \qquad \gamma_j = \frac{2j(j+1)}{3} - \frac{N}{2}, \\ S &= \int_0^T g(t)\, dt = 0. \end{aligned} \tag{18}$$

With one-sign g, $S = 0$ forces $g = 0$ almost everywhere. Equation (8) assumes that reversing the control amplitude reverses its gain-error term. An additional polarity-dependent term $\delta|g(t)|H_{ZZ}$ has a sector-trace response proportional to $\int |g(t)|\, dt$, so exact cancellation of that response forces zero twisting. A small polarity mismatch that preserves the fixed generator in each slot can be counted as an area perturbation under the finite-error budget below.

Supplementary Note 4 quantifies the error floor for one-sign interactions and gives a two-qubit example that attains it. Supplementary Note 5 applies the same phase test to a separate excitation-preserving network with signed local fields, exchange and one controlled ZZ pair.

Regular controls and local programming

Having identified the reachable targets, we next examine the freedom among controls that realize them. For a complete unitary operation, choose traceless Hamiltonian representatives and let $E_{\le q}(a)$ collect the nonconstant coefficients of the local logarithm of the error propagator. Channel-response cancellation is then equivalent to $E_{\le q} = 0$. Write

$$\begin{aligned} F_q(a) &= (U_0(a), E_{\le q}(a)), \\ \mathcal{M}_q &= \{a : E_{\le q}(a) = 0\}, \qquad \pi_q = U_0|_{\mathcal{M}_q}. \end{aligned} \tag{19}$$

In a fixed sequence template with m real control parameters, take independent cancellation coordinates $\mathcal{R} \in \mathbb{R}^c$ and target coordinates $\Theta \in \mathbb{R}^k$ on a declared reachable patch. At a regular point,

$$\operatorname{rank} D\mathcal{R} = c, \qquad \operatorname{rank} D(\mathcal{R}, \Theta) = c + k. \tag{20}$$

The regular-level-set and inverse-function theorems [29] give

$$\begin{aligned} a &= \Phi(\theta, \xi), \qquad \xi \in \mathbb{R}^{m-c-k}, \\ \mathcal{R}(\Phi(\theta, \xi)) &= 0, \qquad \Theta(\Phi(\theta, \xi)) = \theta. \end{aligned} \tag{21}$$

Here Φ is a local coordinate chart onto its image. The k target coordinates program nearby operations and the $m - c - k$ implementation coordinates vary a fixed-target control [Fig. 2(b)]. The same construction applies to independent task-response coordinates in Eq. (1). The ranks characterize the chosen control template and point. Global connectivity of the fixed-target set and regularity at singular controls or actuator boundaries require separate analysis.

Continuous families of controls implementing the same unitary have previously been studied through endpoint tracking [30]. Here the chart also preserves the specified error derivatives. Its geometry concerns the space of control implementations. Holonomic gates instead use geometric phases acquired along cyclic subspace evolution [31]; the chart construction imposes no holonomy condition.

For a target subspace with projector P, let $\mathcal{D}_P$ contain controls whose nominal endpoints preserve it and meet the declared interface conditions. With $\rho_P(a) = PU_0(a)P|_{\operatorname{ran} P}$, write

$$C_V = \{a \in \mathcal{D}_P : \rho_P(a) = V\}, \qquad \mathcal{F}_V = \mathcal{M}_q \cap C_V. \tag{22}$$

This is the logical-target fiber pictured in Fig. 1(a): its controls may realize different complete extensions $V \oplus W$. Specifying W selects a complete-target fiber within it. Moving within either fiber changes the completed implementation; the colored execution paths instead follow one control as it is applied in time. The restricted-input propagator is $U(a, \epsilon)\iota_P$, where ι_P embeds the input subspace. Its output includes leakage.

Finite signed words realizing the target classification can be endowed with regular charts. Finite conjugates of the primitive response generators yield a basis of the joint endpoint-response algebra. Products

along that basis give an invertible differential at the identity, and translation gives a chart at each reachable target (Supplementary Note 1). The collective joint space has dimension $D_{N,q}$ given by Eq. (S12), which scales as q^2N^3 for large q and N. At order two, $D_{2,2} = 49$ and $D_{3,2} = 100$. A chart onto that space needs at least $D_{N,q}$ independent parameters. One construction uses at most $D_{N,q}^2 - 4D_{N,q} + 6$ primitive exponentials for a local chart at the identity (Supplementary Note 7). Implementation of a distant target can require a separate word.

For the collective model, an exact $N = 2$ first-order chart uses 25 variable areas and a nonsingular symbolic Jacobian. The robust target patch has dimension eight in $SU(3)$ and the chart has zero fixed-target nullity at fixed smoothing width (Supplementary Note 3). Compact templates can instead be regular on the restricted coordinates that define their task.

An interval-certified $N = 3$ construction uses three modules of 15 twisting–rotation pairs each. On the control slice with the first rotation azimuth fixed in each module, there are 132 variables and 111 independent constraints for complete first-order cancellation and hidden-sector closure. The remaining 21-dimensional family has target rank 15 for the encoded $SU(4)$ with fixed complementary action. Fifteen directions program the target and six vary the implementation at fixed target. Releasing the three azimuths independently retains both full-rank minors at the certified point, yielding 24 local dimensions and nine fixed-target directions under the same constraints. These counts concern the regular construction in Supplementary Note 3.2. The CNOT controls below are a separate numerically obtained family assessed by their finite-error bounds. At the four- and five-qubit controls used for the circuits below, independently evaluated Jacobians give fixed-complete-target nullities of 47 and 59. These are numerical tangent-space counts at controls with finite residuals (Supplementary Note 18).

Regularity also allows waveform compensation. In the collective model, a positive-duration regular chart appended to an order-q signed word allows sufficiently narrow smoothing of the waveform at the same target and order. The leading effect of adjacent noncommuting generators is quadratic in the smoothing width, and the inverse joint differential gives the leading compensating control shift. Supplementary Note 7 gives the shift, width conditions and proof. Hao and Bukov also use regular endpoint maps to compensate smooth changes [18]. The CNOT certificate below uses a different realization, in which nonoverlapping fixed-generator slots preserve the prescribed calibrated areas.

Finite-error acceptance near a fixed target

Exact cancellation fixes derivatives at zero error. To assess an implementation over a finite error range, we must also bound the remaining response. For a fixed target define

$$\begin{aligned} e_{r,V}(a) &= \sup_{\|\boldsymbol{\epsilon}\| \le r} \|O_I(a,\boldsymbol{\epsilon}) - \mathcal{V}_I\|_I, \\ \mathcal{A}_{\delta,r,V} &= \{a \in C : e_{r,V}(a) \le \delta\}. \end{aligned} \tag{23}$$

The norm measures the declared task. We use half the diamond norm for channel error. The tolerance δ and error domain are fixed before selecting the control.

Proposition 1 (Local finite-error acceptance)**.** *Let $\mathcal{S}$ be a nonempty compact subset of the order-q fixed-target fiber. In fixed normalized control coordinates, suppose O_I extends to a neighborhood of $\mathcal{S}$ with uniform control derivative bound $L > 0$ and uniform directional error derivative bound M_{q+1} through an error ball of radius r_0, as specified in Supplementary Note 7.1. Then, for nearby admissible a and $0 \le r \le r_0$,*

$$e_{r,V}(a) \le B_q r^{q+1} + L \operatorname{dist}(a,\mathcal{S}), \qquad B_q = \frac{M_{q+1}}{(q+1)!}. \tag{24}$$

For $\delta > B_q r^{q+1}$, this estimate gives an accepted neighborhood of $\mathcal{S}$ of radius at most $(\delta - B_q r^{q+1})/L$, intersected with admissible controls and restricted to where the derivative bounds hold. It relates the

Taylor order in dynamical correction [5, 7] to a finite calibration range. The derivative bounds and nominal mismatch of stored controls determine the actual budget.

Finite-error acceptance for encoded CNOT implementations

The reachability theorem guarantees that a corrected sequence exists for the three-qubit target in Eq. (12). We first bound the error of individual stored controls over a continuous range of calibration offsets, retaining their numerical residuals. We then certify a continuous family around the jointly corrected control, including bounded control perturbations and local dephasing. For a unitary implementation a on the full eight-dimensional register, define

$$e_r(a) = \sup_{|x|,|y|\le r} \frac{1}{2}\|\mathcal{U}(a;x,y) - \mathcal{U}_{\mathrm{C}}\|_\diamond. \tag{25}$$

Both calibration offsets are shared across the entire gate. The supremum covers the continuous error square, and the half-diamond distance retains all physical outputs for arbitrary inputs entangled with a reference.

For the stored implementation corrected jointly at gate completion (C), an outward-rounded Taylor enclosure of the complete propagator gives

$$e_{0.018249}(a_{\mathrm{C}}) \le 0.01, \qquad e_{0.01}(a_{\mathrm{C}}) < 0.002897. \tag{26}$$

The calculation encloses coefficients through degree eight and bounds the remainder analytically (Supplementary Note 8.6). The control is designed for first-order cancellation, and the enclosure includes its nominal and first-order residuals. Figure 4(a) compares the five stored complete-target controls. The resulting radii give sufficient lower bounds on the acceptable error range. For R and M the upper bounds at $r = 0.01$ exceed 0.01, even though their sampled errors are below it (Supplementary Note 16).

To quantify implementation freedom, we hold the rotation axes fixed, construct four nearby numerically corrected controls around C, and join the five points by affine interpolation of their 90 pulse areas. An analytic interpolation bound covers every point of this path $\mathcal{P}$, giving

$$\sup_{a\in\mathcal{P}} e_{0.01}(a) < 0.005937122. \tag{27}$$

The bound includes nominal mismatch and first-order residuals throughout the path. The construction supplies an explicit subset of the finite-error acceptance set in Eq. (23).

For a base control a, changes of rotation areas, free azimuths, and twisting areas have the weighted budget

$$W(h;a) = \frac{3}{2}\sum_j |\Delta\theta_j| + \frac{3}{2}\sum_j |\theta_j|\,|\Delta\phi_j| + 3\sum_j |\Delta\chi_j|. \tag{28}$$

This weighted sum is a sufficient perturbation budget. In particular, an azimuth change costs nothing for a rotation of zero area. Duhamel's formula bounds the additional channel error from these perturbations by $(1 + r)W$, including calibration–control cross terms. One schedule covers the entire path and $W \le 0.0037$, with the actuator and switching costs specified in Methods. This common schedule has gate duration $T_g < 17.547134t_0$. Adding local Markov dephasing from Eq. (32) contributes at most $3\gamma T_g/2$ to half-diamond distance. Thus every control in this neighborhood satisfies

$$\boxed{\begin{array}{c} |x|, |y| \le 0.01, \quad W \le 0.0037, \quad \gamma t_0 \le 10^{-5} \\ \Longrightarrow \quad \frac{1}{2}\|\mathcal{E} - \mathcal{U}_{\mathrm{C}}\|_\diamond < 0.009938 < 0.01. \end{array}} \tag{29}$$

With axes fixed, the budget includes simultaneous deviations of up to 1.8×10^{-5} radians in each of the 90 areas. Both this joint box and the broader weighted budget retain all physical output sectors and

noise-induced leakage. The guarantee applies at the completed gate boundary. The gate time accounts for the pulse slots and switches, and the phase-readout cycle has additional costs (Methods). Here $\gamma = 1/T_\phi$ is the one-qubit pure-dephasing rate. The condition on γ requires $T_\phi \geq 10^5 t_0$, or 1–10 s if t_0 is 10–100 μs. These are physical requirements for that choice of clock.

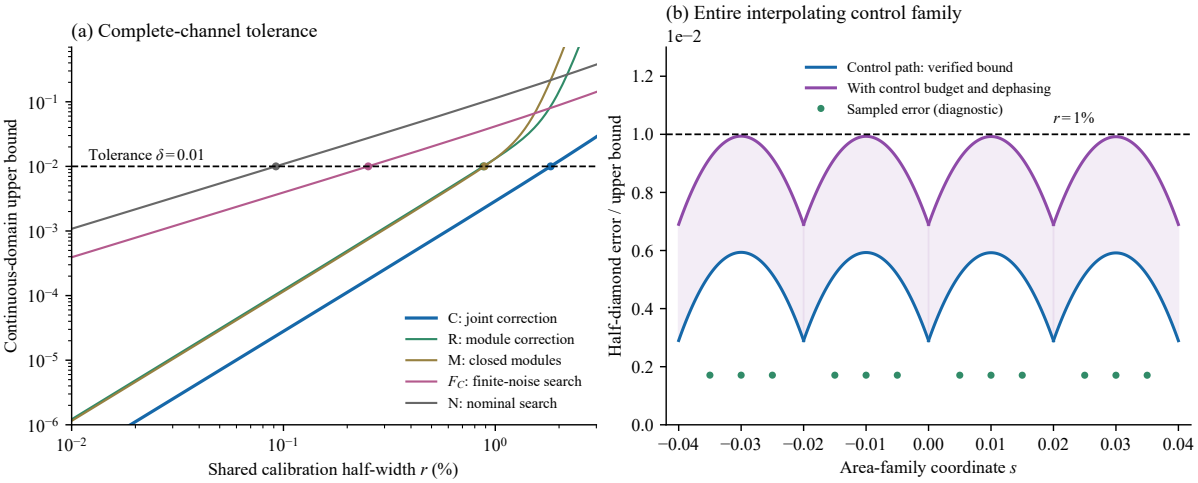


Figure 4: **Certified error bounds for encoded-CNOT implementations. (a)** Eighth-degree enclosures uniformly bound the complete-channel half-diamond error of each stored implementation under shared gain offsets alone. Intersections with $\delta = 0.01$ give sufficient shared-calibration half-widths for the fixed library. The correction design targets first order. The eighth-degree enclosure retains nonzero nominal and first-order residuals. **(b)** The five nearby C controls at $s = -0.04, -0.02, 0, 0.02, 0.04$ are joined by affine pulse-area segments. Blue curves bound every interpolating control for $|x|, |y| \leq 1\%$. Purple curves bound the error within $W \leq 0.0037$ of each path point, including local dephasing at $\gamma t_0 = 10^{-5}$ under the common schedule. Throughout this neighborhood $D < 0.009938$. Green points are simulated coherent maxima at sampled controls and gains. Supplementary Notes 8.6 and 16 give the proofs, arithmetic and data.

Protection boundaries and implementation cost

The blue and orange constructions in Fig. 1(b) impose cancellation at different completion boundaries and therefore define different feasible control sets. With the same 45-pair template and one common feasible start, the local searches reached $76.788 t_0$ for three separately corrected modules with hidden-sector closure (M), $61.699 t_0$ when the hidden-sector operations need only close at the end of the gate (R), and $18.445 t_0$ for joint correction across the complete gate (C). Each run had a cap of 120 outer iterations or 90 s. M shortened the common starting duration only slightly, from $76.809 t_0$ to $76.788 t_0$. A projected-response arm with the complete nominal target finished longer than C. Supplementary Table S1 gives the constraints and ranks. In later five-start feasibility checks, R reached $59.619 t_0$ and M reached $76.788138 t_0$, compared with $76.788413 t_0$ from the common start. These times are feasible results of finite local searches. Minimum gate times are open.

Separate correction covers gains that vary between completed modules. The joint-correction search exploits gains shared across the gate. In a separate library comparison with module-dependent gains and local dephasing, the preferred control varies with the task, gain correlations and duration. Supplementary Note 17 gives the selection results for mean entanglement infidelity, while Eq. (29) bounds worst-case channel distance over continuous shared gains.

Completion boundaries and composition

Equation (3) specifies how responses from successive operations combine. For inputs in P, including reference-assisted inputs, nominal preservation of P and $K_{c,\mu}P = c_\mu P$ suffice to cancel both logical distortion and leakage through first order. Separate correction supplies a guarantee at each completed module and admits different quasistatic errors in different modules. Joint correction uses the specified correlations across a sequence and supplies the guarantee at its completion. Geometric approaches to coherent-error suppression also exploit such interference at the circuit level [17]. Intermediate subspace excursions are compatible with the endpoint condition when the control alphabet permits them, as illustrated in Fig. 1(b).

Finite-depth bound. For a completed module with errors $(x, y) = rv_j$ and Hamiltonian linear in those quasistatic gains, let H_R and H_T denote the nominal drive and interaction terms. The interaction-frame error action is

$$\begin{aligned} A_j(v_j) &= \int_0^{T_j} \|E_{j,v_j}(t)\|\, dt, \\ E_{j,v_j}(t) &= U_{j,0}(t)^\dagger [v_{j,R} H_R(t) + v_{j,T} H_T(t)] U_{j,0}(t). \end{aligned} \tag{30}$$

If its complete responses vanish through order q, the Volterra remainder is at most $(|r|A_j)^{q+1}/(q+1)!$. For M completed modules with compatible interfaces and guarantees covering the states passed between them,

$$\boxed{D_{\text{circuit}} := \frac{1}{2}\|\widetilde{C} - C\|_\diamond \le \min\left\{1, \sum_{j=1}^{M} \frac{(|r|A_j(v_j))^{q+1}}{(q+1)!}\right\}.} \tag{31}$$

The Volterra and channel-telescoping bounds allow noncommuting gates and coherent accumulation [32–34]. With a common action bound, the sufficient gain radius at a fixed error budget scales as $M^{-1/(q+1)}$. A jointly corrected sequence counts as one completed module. Supplementary Note 7 includes the additional terms for nominal mismatch and nonzero response residuals in actual compiled gates. The triangle inequality adds the error magnitudes, so this expression can be loose when transported errors partially cancel. The circuit comparisons below use direct channel propagation.

Encoded circuits on four and five physical qubits

We now use completed gates as circuit interfaces and test their performance at finite depth. Do first-order controls provide good starting points for a compiler optimizing performance at finite errors? We tested this question on four and five physical qubits, keeping the logical task fixed. At each width, the four Dicke states with $N, N-1, N-2, N-3$ excitations encode two logical qubits. The complete target applies $e^{i\pi/4}\text{CNOT}_L$ there and the identity everywhere else. The determinant criterion establishes exact reachability at both widths. The simulations ask how accurately a fixed pulse template implements this target across a range of gains. We build noncommuting circuits by changing the drive azimuth of successive copies of one compiled gate. All gates act on the same encoded pair and preserve the same logical control bit. One pair of unknown gains applies throughout a circuit.

For $N = 4, 5$, the four-dimensional computational subspace is a proper subspace of the symmetric spin sector. Collective rotations can populate symmetric Dicke states outside this encoding while preserving total spin: for example, $\langle D^{(N)}_{N-4}|J_x|D^{(N)}_{N-3}\rangle = \sqrt{N-3} \neq 0$. This is the within-sector form of the subspace excursion in Fig. 1(b). The nominal target requires return at gate completion, and the finite-error analysis retains any residual leakage.

Figure 5 compares four controls from one development start per width. A direct finite-noise fit starts from a nominal gate. A second finite-noise fit starts from a gate whose two first-order responses have

been minimized. We also evaluate both starting gates. The four methods use the same physical schedule within each width and meet the same nominal accuracy requirement, $D \leq 10^{-6}$. In this first comparison the correction stage cost an extra 42 optimizer evaluations at $N = 4$ and 38 at $N = 5$.

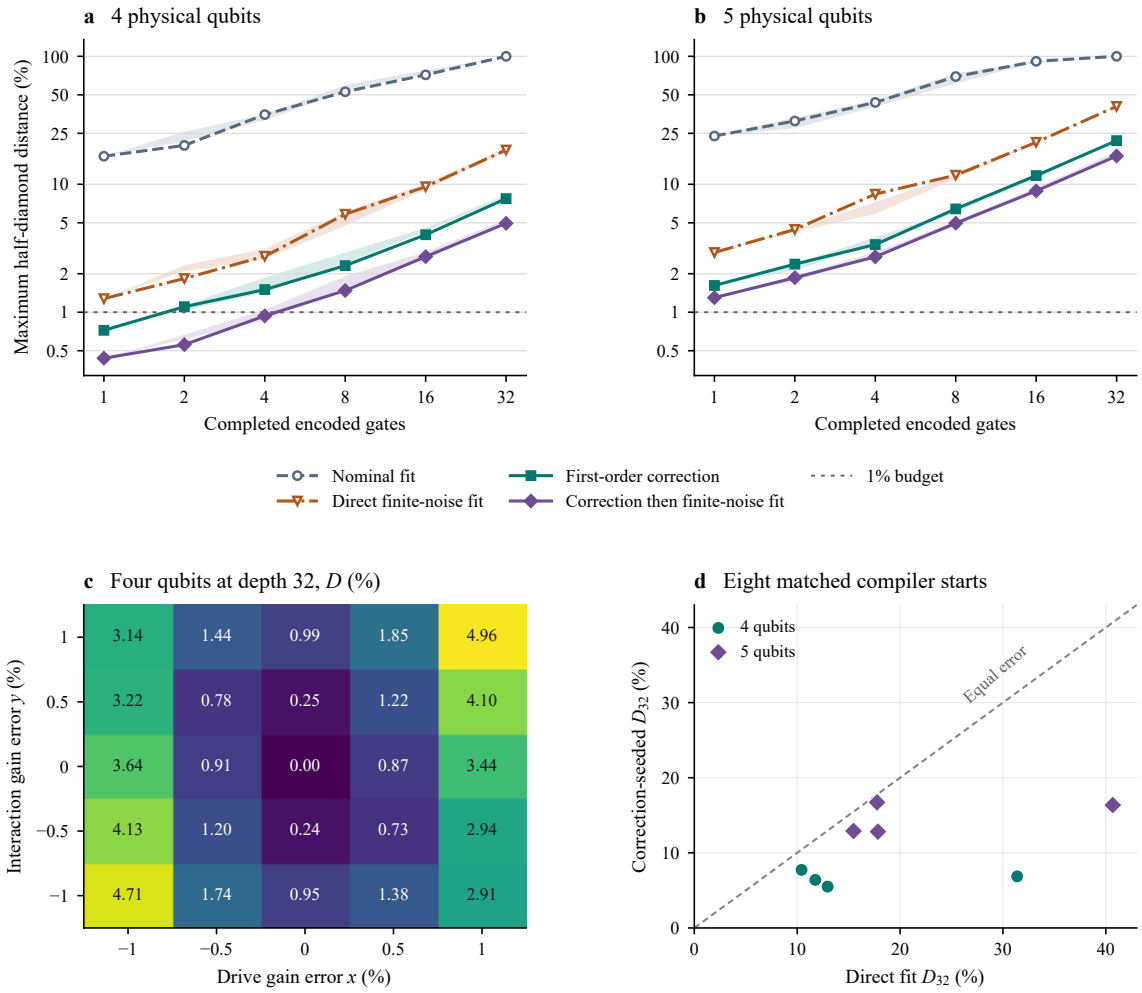


Figure 5: **Complete-channel error through noncommuting encoded circuits. (a,b)** Maximum half-diamond distance over a 5×5 grid of shared gains from -1% to 1%. Lines give prefixes of one fixed sequence of azimuth-conjugate gates. Each gate swaps the same pair of logical states with a phase, so two successive ideal gates act diagonally in the Dicke basis. Shading spans three further azimuth sequences using the same control fit. The dashed line marks the 1% distance budget. All methods have the same gate time, pulse count and nominal accuracy threshold within each width. The pulse count and duration change with the physical register size. **(c)** Individual gain settings for the refined four-qubit control at depth 32. The cells are discrete evaluations. **(d)** Worst sampled distances at depth 32 for eight new paired starts. Direct and correction-seeded searches received equal budgets with five training offsets. Points below the diagonal favor the correction-seeded arm. Supplementary Fig. S3 shows the separate comparison with 17 training offsets.

On the four-qubit circuit at depth 32, the maximum error on the 25-point gain grid was 18.53% for the direct fit, 7.72% for first-order correction alone and 4.96% after finite-noise refinement. The five-qubit values were 40.61%, 21.94% and 16.63%. Here 4.96% is the largest change in an outcome probability, in percentage points, attainable with the complete four-qubit channel at the sampled gains. The corresponding maximum leakage over encoded inputs and the 25 tested gains was 0.0154% (0.251% on five qubits). The

complete-channel error also sees distortions within the encoding and on other physical inputs. Neither width met the prespecified 1% distance budget at depths 16 and 32. At $N = 4$ refinement raised the drive-gain response norm from 8.83×10^{-8} to 0.00669, even as the depth-32 distance fell from 7.72% to 4.96%. The benefit here comes from the fit at finite gain as well as the slope at zero.

We then matched the compilation budgets in two studies, each with four fresh starts per width. With five training offsets, first-order initialization reduced the depth-32 maximum in all eight pairs on both the 25- and 81-point validation grids. The direct-to-seeded ratios on the 25-point grid ranged from 1.35 to 4.57 on four qubits and from 1.06 to 2.49 on five. In the second study both arms trained at 17 offsets spanning the full gain square, including the corners, and were scored at 80 disjoint settings. Seeding reduced the maximum in all four four-qubit pairs and three of the four five-qubit pairs. In the remaining pair the direct fit gave 13.09% and the seeded fit 13.34%. The initialization advantage thus depends on the optimization start. Supplementary Note 18.2 reports every pair, its cost and the validation procedure.

Local dephasing changes the comparison because it acts throughout the gate clock. On the original four-qubit schedule at $\gamma t_0 = 10^{-5}$, the depth-32 loss of a specified cross-sector readout rose from 0.0152% to 8.18% for the refined control and from 2.11% to 10.18% for the direct fit. The no-jump component also bounds the complete-channel distance. On the same 25 gain settings, the refined control has a noisy grid maximum at most 13.43%, and the direct fit has a maximum at least 13.78%. A shorter clock that obeys the same actuator limits for all four controls separates these bounds further (Supplementary Note 18.3).

Discussion

Dynamical error correction constrains a completed operation for a specified task, error model and completion boundary. Figure 1 relates these constraints to physical execution: constituent responses can cancel jointly, and controls can explore outside the computational subspace when the available couplings permit. The endpoint map turns the corrected control set into a set of reachable complete operations. Under the linear-gain hypotheses, that image is the same at every positive finite order, although its implementing sequences can change. The collective phase witness shows why the full physical extension of a logical gate matters for membership in this set.

Once the complete target is fixed, regular local fibers describe its implementation freedom, including compensating waveform changes. Completion boundaries determine the interfaces at which corrected operations can be reused. Finite-error bounds and circuit propagation then assess accuracy over a specified calibration range and depth. This separates exact reachability, local control geometry and the performance of stored implementations.

A continuous neighborhood of encoded CNOT implementations has certified worst-case half-diamond error below 0.01 over the specified shared gain offsets, control perturbations and local dephasing. This is a finite-error acceptance guarantee. Exact fixed-target fiber membership and any connection to the interval-certified three-module family remain open. The four- and five-qubit circuits address a separate compilation question. For the tested azimuth-conjugate family, first-order initialization reduced the sampled depth-32 maxima in eight of eight matched-budget comparisons with five training offsets and seven of eight with 17 training offsets spanning the gain square. The wider training set produced one reversal at $N = 5$. Both widths exceeded the 1% budget at depths 16 and 32 in the original depth study.

The classification concerns finite words over signed, drift-free primitives with a fixed shared gain realization. For the two independent collective gains, a single finite word with every Taylor coefficient zero can implement only the identity channel. Analyticity extends the constant channel to $x = y = -1$, where all controls switch off. The existence proof permits word lengths and durations that increase with the requested order. The relevant open control problem is how those resources scale under amplitude, bandwidth and sign constraints.

Calibrated positive and negative areas, the specified actuator law and quasistatic gains define the implementation model. Inhomogeneous gains or drift introduce additional response directions, potentially including motion between spin sectors. Time-dependent fluctuations require specifying their temporal structure; filter-function methods relate that structure to the control response [35, 36]. Cancellation order for quasistatic gains alone does not determine filtering performance at nonzero frequency. Local dephasing in the certificate contributes to the error budget throughout the specified schedule.

Code and subsystem protection are governed by quantum and operator quantum error-correction conditions [37, 38]. Here the Dicke states specify a logical embedding. They do not correct arbitrary single-site phase errors. On the three-qubit symmetric space, $PZ_iP = (2/3)PJ_zP$ is nonscalar, so the error pair I, Z_i fails the Knill–Laflamme condition. The simulated circuits retain coherent accumulation at a fixed gain realization. A conversion to effective stochastic Pauli noise by randomized compiling [39] would require additional randomizing operations compatible with the encoding and control alphabet. Establishing a fault-tolerance threshold would also require a specified code and recovery protocol.

Methods

Response algebra and regular control charts

For each primitive, we represent the endpoint and its mixed Taylor coefficients in the algebra of error polynomials truncated at degree q. Products implement control concatenation and brackets multiply the corresponding error polynomials. We decompose the resulting algebra into central and simple factors using the actual nominal generators, with scalar Hamiltonian parts subtracted. The central coefficients depend linearly on integrated primitive areas. On a simple factor the commutators generate the truncated polynomial algebra of its active profiles. A basis of finite conjugates of primitive responses gives an invertible endpoint-response Jacobian, yielding exact finite words and local charts in the reachable group. Supplementary Notes 1 and 2 give the proof and the conditions for extending it beyond linear profiles. Supplementary Notes 3–7 give chart verification, waveform continuation and finite-depth remainder estimates.

Complete-channel distance and phase readout

The gate and circuit distances in the main figures retain the complete physical output. For two unitary channels we compute D from the eigenvalues of $U^\dagger V$. If ω is the length in radians of the shortest spectral arc containing those eigenvalues, then $D = \sin[\min(\omega, \pi)/2]$. Equivalently, if d_{hull} is the distance from zero to their convex hull, $D = \sqrt{1 - d_{\mathrm{hull}}^2}$ [32–34]. The maximization defining D includes arbitrary physical inputs entangled with a reference. When evaluating the projected-input task, the input is restricted to its code and the output still includes all leakage.

For the phase readout, start in $|000\rangle$ and apply an addressed $R_x^{(1)}(\pi)$ before the collective gate. After the gate, apply $R_\beta^{(1)}(\pi/2) = \exp[-i\pi(\cos\beta\, X_1 + \sin\beta\, Y_1)/4]$ at the fixed azimuths $\beta_+ = -\pi/4$ and $\beta_- = 3\pi/4$. Read out all three qubits in the computational basis and count 000 without postselection. Before the analysis pulse only the $|000\rangle$ and $|100\rangle$ components can contribute to the 000 amplitude; their coefficients after U_ϕ have magnitudes $1/\sqrt{3}$ and $2/3$ and relative phase ϕ. The two settings give Eq. (15). The addressed pulses are additional preparation and analysis resources. Supplementary Note 8 includes their duration, sampling and state-preparation and measurement costs. The declared overhead is $1.16188t_0$ per shot, separate from the gate-only clock.

Actuator and physical gate time

The time model uses signed, nonoverlapping slots, with $t_0 = 1/g_{\max}$, drive limit $\Omega_{\max} = 10g_{\max}$ and a specified linear actuator response $\tau_b\dot{u} + u = v$, where $\tau_b = 0.02t_0$ and v is the command. A pulse area a in a slot of length τ is delivered as $u(t) = a[1 - \cos(2\pi t/\tau)]/\tau$. The compensated command $v = u + \tau_b\dot{u}$ must obey the appropriate amplitude limit. The smallest slot satisfying that limit and $\tau \geq 0.04t_0$ gives the area-dependent schedule. We add $0.001t_0$ at each switch. Each pulse stays along a fixed generator with a constant gain during the slot, so its propagator is determined by its signed area. The declared actuator delivers that area by compensated shaping. Gain dynamics within a pulse or an uncertain actuator would require a separate area or waveform error model. The three-qubit certificate uses a schedule common to the entire accepted control set and includes its 89 internal transitions.

For four and five physical qubits we use 48 and 72 twisting–rotation pairs per gate, respectively. For each width we set common slot times able to deliver rotation areas up to π and twisting areas up to 1.2. All controls at fixed width have the same time under that schedule. Including internal switches, the gate times are $145.83t_0$ and $218.74t_0$, and 32 gates with between-gate switches take $4666.46t_0$ and $6999.70t_0$. Separate area-dependent timings are given in Supplementary Note 18. Initialization, measurement and reset are outside the circuit gate times.

Circuit compilation and discrete validation

Each compiled gate applies alternating twisting and collective rotation slots in the model of Eq. (8). Every compilation arm parameterizes the twisting areas as $\chi_\ell = c_\ell - \overline{c}$ in dimensionless area units, so the total twisting area is zero before any response is fitted. Shifting all drive azimuths by φ_j implements conjugation by $\exp(-i\varphi_j J_z)$ while the twisting slots are unchanged. This rephasing uses no extra physical rotations. We use a predetermined noncommuting sequence with $\varphi_j = \pi(\sqrt{2}j^2 + \sqrt{3}j) \bmod 2\pi$, starting at $j = 0$, and three additional independently chosen sequences. We compare nominal fitting, a fit to nominal and first-order responses, and finite-noise fitting from each of these two starting controls. Least-squares residuals use multiplicity-weighted Frobenius norms over the spin blocks. The first-order fit reuses the propagator and pulse-parameter Jacobian used by the finite-noise objective. It approximates the gain responses using centered offsets $\pm 2 \times 10^{-5}$ and weight 0.12 per derivative. Independent checks evaluate those gain responses analytically. For finite-noise fitting we use the origin and four training offsets $(x, y) = (\pm 0.006, \pm 0.006)$ and restore nominal accuracy with a minimum-norm Jacobian step iterated to the recorded stopping rule. Methods share their nominal start within each paired search. Candidate selection uses the training points among candidates passing the nominal threshold. Validation uses $M = 1, 2, 4, 8, 16, 32$ and the 25 gain pairs in $\{-0.01, -0.005, 0, 0.005, 0.01\}^2$. The validation grid shares only the origin with training. For the comparison with matched compilation budgets, we fix four new seeds at each width and also evaluate a denser 81-point gain grid on the original sequence at $M = 32$. A separate matched-budget study trains on 17 offsets spanning the gain square and tests on 80 held-out offsets.

The primary calculations use the symmetry blocks, retaining all multiplicity copies in the metric. Independent checks use physical-basis Pauli tensors, analytic response derivatives, direct execution with the altered axes, and the convex-hull distance formula. Supplementary Notes 18 and 18.2 report the seeds, derivative residuals, target errors, timing, stopping rules and saved candidates. The controls and targets are specified independently of the validation offsets.

Local dephasing and finite-error verification

We model independent local Markovian dephasing by

$$\dot{\rho} = -i[H, \rho] + \frac{\gamma}{2} \sum_{r=1}^{N} (Z_r \rho Z_r - \rho). \tag{32}$$

It acts throughout the declared gate and idle intervals. For the three-qubit certificate we bound the complete-unitary error by outward-rounded Taylor coefficients through degree eight and an analytic remainder. The controls are designed for first-order correction. The higher Taylor degree sets the order of the enclosure, which retains the nominal and first-order residuals. With fixed rotation axes, an analytic interpolation bound encloses the four affine segments between five stored controls. Duhamel's formula covers the weighted neighborhood of each base control, including calibration and control cross terms. The dephasing contribution is bounded by $3\gamma T_g/2$. Supplementary Notes 8.6 and 16 give the arithmetic and the actuator inequalities.

In the separate four-qubit dephasing simulation we propagate an unprojected density matrix on the common clock. We initialize $(|D_2^{(4)}\rangle + |s_{12}\rangle|00\rangle_{34})/\sqrt{2}$, with $|s_{12}\rangle = (|01\rangle - |10\rangle)/\sqrt{2}$, and compute the probability for the projector onto the ideal output at each depth. This fixed input and measurement give a state-transfer probability loss. Twisting commutes with the dephasing superoperator. The simultaneous drive and dissipator during each rotation slot factorize into single-site channels, which we compute by solving the shaped Bloch equations with DOP853 (relative tolerance 2×10^{-11}, absolute tolerance 2×10^{-13}). An independently integrated one-site Liouville equation checks a test pulse. Numerical checks cover propagation at zero dephasing, integration tolerance, trace, Hermiticity and positivity. To bound the complete-channel distance on the same gain grid, we decompose the noisy channel into its coherent no-jump contribution of weight $\exp[-N\gamma T(M)/2]$ and the channel conditional on at least one jump. Convexity supplies an upper bound, and a measurement on a pure input attaining the coherent distance supplies a lower bound. We apply these analytical mixture bounds to the numerical coherent distances. This also allows comparison on a shorter clock common to all four control designs. Supplementary Note 18.3 contains the propagation and bracket data.

Three-qubit constraint and finite-library comparisons

The three-qubit library uses 45 twisting–rotation pairs, with 132 control parameters across three blocks. At the fixed feasible start we compare five constraint sets that differ in their nominal targets, projected or complete first-order conditions, and module-boundary conditions. A local search proposes a reduction in actuator-limited time along the tangent to the applicable constraints. Nonlinear constraint correction and backtracking check recorded feasibility and duration before accepting the step. Each matched-start run receives at most 120 outer iterations or 90 s. A separate finite library includes additional nominal controls and finite-noise refinement, with family-dependent search budgets. In the library comparison the gain in each module is a common offset plus an independent module offset. We compute reference-assisted mean entanglement infidelity with the local dephasing of Eq. (32) using a fixed confirmation ensemble. The control-selection criterion tests eligibility for the declared task before ranking mean scores. Supplementary Notes 8 and 17 specify the constraints, gains, budgets and uncertainty estimates.

Use of AI tools

Generative AI agents, including DeepSeek Harnese, Kilo Code and Kimi, assisted with literature organization, derivations, numerical code, analysis and manuscript preparation. The author is responsible for the scientific claims and verification.

Data availability

The numerical data supporting the figures and tables, final pulse parameters and interval-verification inputs and outputs are available in the RobustControl repository, revision `0e02087`.

Code availability

The simulation, plotting and interval-verification code and replay instructions are available in the RobustControl repository, revision `0e02087`. Start with `README.md` in that directory.

Acknowledgements

This work was supported by the Innovation Program for Quantum Science and Technology (2021ZD0301703) and the Quantum Science and Technology-National Science and Technology Major Project (2024ZD0300400). Additional support came from the Science, Technology and Innovation Commission of Shenzhen Municipality (grants JCYJ20170412152620376 and KYTDPT20181011104202253), the Shenzhen Science and Technology Program (grants KQTD20200820113010023 and ZDYJ20251211121039009), and the CCF-QuantumCtek Superconducting Quantum Computing Special Cooperation Program (CCF-QC2025002). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

X.-H.D. conceived the study, developed the theoretical framework, performed the analytical and numerical work, prepared the figures, and wrote and revised the manuscript.

Competing interests

The author declares no competing interests.

References

[1] Wimperis, S. Broadband, narrowband, and passband composite pulses for use in advanced NMR experiments. *J. Magn. Reson. Ser. A* **109**, 221–231 (1994).

[2] Cummins, H. K., Llewellyn, G. & Jones, J. A. Tackling systematic errors in quantum logic gates with composite rotations. *Phys. Rev. A* **67**, 042308 (2003).

[3] Brown, K. R., Harrow, A. W. & Chuang, I. L. Arbitrarily accurate composite pulse sequences. *Phys. Rev. A* **70**, 052318 (2004).

[4] Viola, L., Knill, E. & Lloyd, S. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.* **82**, 2417–2421 (1999).

[5] Khodjasteh, K. & Viola, L. Dynamically error-corrected gates for universal quantum computation. *Phys. Rev. Lett.* **102**, 080501 (2009).

[6] Khodjasteh, K. & Viola, L. Dynamical quantum error correction of unitary operations with bounded controls. *Phys. Rev. A* **80**, 032314 (2009).

[7] Khodjasteh, K., Lidar, D. A. & Viola, L. Arbitrarily accurate dynamical control in open quantum systems. *Phys. Rev. Lett.* **104**, 090501 (2010).

[8] Tomita, Y., Merrill, J. T. & Brown, K. R. Multi-qubit compensation sequences. *New J. Phys.* **12**, 015002 (2010).

[9] Li, J.-S. & Khaneja, N. Control of inhomogeneous quantum ensembles. *Phys. Rev. A* **73**, 030302(R) (2006).

[10] Agrachev, A., Baryshnikov, Y. & Sarychev, A. Ensemble controllability by Lie algebraic methods. *ESAIM Control Optim. Calc. Var.* **22**, 921–938 (2016).

[11] Zhang, W. & Li, J.-S. Ensemble control on Lie groups. *SIAM J. Control Optim.* **59**, 3805–3827 (2021).

[12] Kosut, R. L., Grace, M. D. & Brif, C. Robust control of quantum gates via sequential convex programming. *Phys. Rev. A* **88**, 052326 (2013).

[13] Hai, Y.-J., Song, Y., Li, J., Zeng, J. & Deng, X.-H. Geometric correspondence of noisy quantum dynamics and universal robust quantum gates. *Phys. Rev. Applied* **23**, 054002 (2025).

[14] Zeng, J., Deng, X.-H., Russo, A. & Barnes, E. General solution to inhomogeneous dephasing and smooth pulse dynamical decoupling. *New J. Phys.* **20**, 033011 (2018).

[15] Buterakos, D., Das Sarma, S. & Barnes, E. Geometrical formalism for dynamically corrected gates in multiqubit systems. *PRX Quantum* **2**, 010341 (2021).

[16] Yi, K. *et al.* Robust quantum gates against correlated noise in integrated quantum chips. *Phys. Rev. Lett.* **132**, 250604 (2024).

[17] Zeng, J., Liang, H., Hai, Y.-J. & Deng, X.-H. Bridging pulse and circuit levels: A geometric framework for deterministic error suppression. *Quantum Sci. Technol.* **11**, 045025 (2026).

[18] Hao, Y. & Bukov, M. Designing robust quantum gates. *Preprint at* https://arxiv.org/abs/2610.24039v2 (2026).

[19] Marvian, I. Restrictions on realizable unitary operations imposed by symmetry and locality. *Nature Physics* **18**, 283–289 (2022).

[20] Albertini, F. & D'Alessandro, D. Controllability of symmetric spin networks. *J. Math. Phys.* **59**, 052102 (2018).

[21] Kazi, S., Larocca, M. & Cerezo, M. On the universality of S_n-equivariant k-body gates. *New J. Phys.* **26**, 053030 (2024).

[22] Sørensen, A. & Mølmer, K. Quantum computation with ions in thermal motion. *Phys. Rev. Lett.* **82**, 1971–1974 (1999).

[23] Makhlin, Y. Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations. *Quantum Information Processing* **1**, 243–252 (2002).

[24] Hayes, D. *et al.* Coherent error suppression in multiqubit entangling gates. *Phys. Rev. Lett.* **109**, 020503 (2012).

[25] Shapira, Y., Shaniv, R., Manovitz, T., Akerman, N. & Ozeri, R. Robust entanglement gates for trapped-ion qubits. *Phys. Rev. Lett.* **121**, 180502 (2018).

[26] Webb, A. E. *et al.* Resilient entangling gates for trapped ions. *Phys. Rev. Lett.* **121**, 180501 (2018).

[27] Milne, A. R. *et al.* Phase-modulated entangling gates robust to static and time-varying errors. *Phys. Rev. Applied* **13**, 024022 (2020).

[28] Le, N. H. *et al.* Amplitude-noise-resilient entangling gates for trapped ions (2025). Version 2, 2407.03047.

[29] Lee, J. M. *Introduction to Smooth Manifolds*, vol. 218 of *Graduate Texts in Mathematics* (Springer, New York, 2012), 2 edn.

[30] Dominy, J. & Rabitz, H. Exploring families of quantum controls for generating unitary transformations. *J. Phys. A: Math. Theor.* **41**, 205305 (2008).

[31] Sjöqvist, E. *et al.* Non-adiabatic holonomic quantum computation. *New J. Phys.* **14**, 103035 (2012).

[32] Gilchrist, A., Langford, N. K. & Nielsen, M. A. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A* **71**, 062310 (2005).

[33] Sanders, Y. R., Wallman, J. J. & Sanders, B. C. Bounding quantum gate error rate based on reported average fidelity. *New J. Phys.* **18**, 012002 (2016).

[34] Kueng, R., Long, D. M., Doherty, A. C. & Flammia, S. T. Comparing experiments to the fault-tolerance threshold. *Phys. Rev. Lett.* **117**, 170502 (2016).

[35] Green, T. J., Sastrawan, J., Uys, H. & Biercuk, M. J. Arbitrary quantum control of qubits in the presence of universal noise. *New J. Phys.* **15**, 095004 (2013).

[36] Paz-Silva, G. A. & Viola, L. General transfer-function approach to noise filtering in open-loop quantum control. *Phys. Rev. Lett.* **113**, 250501 (2014).

[37] Knill, E. & Laflamme, R. Theory of quantum error-correcting codes. *Phys. Rev. A* **55**, 900–911 (1997).

[38] Kribs, D., Laflamme, R. & Poulin, D. Unified and generalized approach to quantum error correction. *Phys. Rev. Lett.* **94**, 180501 (2005).

[39] Wallman, J. J. & Emerson, J. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325 (2016).