

Article

Drift-Adapted Lattice Geodesics for Quantum Gate Synthesis: Exact Global Optima for Full-Isotropic $SU(n)$ and Weighted Commuting Sectors

Spyridon Talaganis 

Department of Physics, Lancaster University, Lancaster LA1 4YB, UK; s.talaganis@lancaster.ac.uk

Abstract

Finite-dimensional closed-system gate synthesis is a geometric optimal-control problem on a compact Lie group, but global solutions require careful treatment of logarithm branches, determinant-one constraints, drift, anisotropic penalties, and amplitude limits. This paper assembles and extends a self-contained family of exactly solved benchmarks under explicit hypotheses. For a positive right-invariant quadratic metric, smooth stationary curves obey the Euler–Arnold equation and possess Lax invariants. Under fully actuated isotropic control on $SU(n)$, the fixed-time minimum action is the squared Hilbert–Schmidt distance divided by twice the gate time and is generated by a minimum-norm skew-Hermitian logarithm. A finite eigenphase-unwrapping rule computes that logarithm, while the affine spectral-width cut wall and a branch-gap criterion distinguish stable selection from set-valued behavior. Drift is removed isometrically whenever the metric is invariant under the drift adjoint action. Weighted action and box-constrained time on maximal tori and arbitrary closed commuting subtori reduce to explicit period-lattice problems; the latter gives distinct strict determinant-one and projective period lattices for an ideal $fSim$ sector. Numerical validation uses 1506 Haar-random targets through dimension 128, a complete cost-bounded dynamic-program cross-check for every target, 300 direct Cartesian branch enumerations, 120 random-conjugation tests with analytically prescribed spectra, 240 independently cross-checked random weighted-torus instances, a separate $SU(256)$ stress target, large-cluster selector regressions, repeated timing trials, GRAPE–L-BFGS, a first-order sequential Krotov-type update with exact discrete propagation and Fréchet derivatives, randomized and drift-adapted CRAB–Powell bases, and cut-locus tests that exercise the implemented Schur/logarithm solver. When sparse actuation, decoherence, leakage, or uncertainty invalidates the ideal hypotheses, the exact results are positioned as ideal-model reference values, branch-aware seeds, and regression tests rather than as certificates for the enlarged objective.



Academic Editor: Janos Polonyi

Received: 11 May 2026

Revised: 15 July 2026

Accepted: 23 July 2026

Published: 24 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: quantum control; optimal control; geometric control; Lie groups; $SU(n)$; quantum speed limits; matrix logarithm; cut locus; maximal torus; lattice optimization; numerical benchmarks; robustness

1. Introduction

The synthesis of a target quantum gate by a controlled Hamiltonian is naturally an optimal-control problem on a Lie group. After removal of the physically irrelevant scalar phase, the propagator of an n -level closed system lies in $SU(n)$, while its instantaneous

traceless Hamiltonian is represented by an element of the Lie algebra $\mathfrak{su}(n)$. Controllability, reachable sets, minimum time, and minimum fluence can therefore be studied using geometric control theory and Riemannian or sub-Riemannian geometry [1–5]. These tools underlie analytic work on time-optimal spin dynamics and quantum brachistochrone problems [6–9], and they are closely related to geometric descriptions of quantum computation and circuit complexity [10,11].

Most device-level quantum-control design is numerical. Variational formulations, GRAPE, Krotov-type methods, chopped-random-basis (CRAB) optimization, and modern software packages can accommodate sparse Hamiltonians, amplitude and bandwidth constraints, leakage, robustness objectives, and large parameter spaces [12–20]. Exact solutions remain useful because they provide a global certificate that a numerical pulse cannot improve upon within the stated model. They also expose errors that are otherwise difficult to diagnose: an incorrect matrix-logarithm branch, an interaction-picture sign error, a failure to cross a phase-lattice boundary, or an apparent optimum that is only local.

The practical motivation is not that current hardware is fully actuated on $SU(2^k)$. Rather, experimentally calibrated native gates often admit reduced effective descriptions. A fixed-axis Mølmer–Sørensen interaction in trapped ions can generate a one-parameter or low-dimensional commuting sector, while echoed cross-resonance and tunable-coupler gates in superconducting circuits are commonly interpreted through a small number of effective Pauli couplings [21–25]. Moreover, the nonlocal core of an ideal two-qubit gate is a Cartan-torus problem generated by the mutually commuting operators XX , YY , and ZZ [6,26]. The exact formulas developed here can therefore be used directly in those ideal reduced sectors or as ideal-model references, initializations, and regression tests when the complete hardware model is noncommuting or open. They remain rigorous lower bounds for a more constrained closed plant only when the endpoint and objective are unchanged and the ideal admissible set contains the plant’s controls.

This paper develops such benchmarks under deliberately explicit hypotheses. Full actuation means that the control Hamiltonian can realize every element of $\mathfrak{su}(n)$; torus actuation means that it remains in a fixed Cartan subalgebra; and drift is cost-neutral only when the quadratic control metric is invariant under the drift adjoint action. These assumptions are restrictive, but they make the global optima mathematically sharp and reproducible. Standard matrix-analysis, matrix-function, and lattice terminology follows [27–30].

The geometric ingredients—Euler–Arnold dynamics, bi-invariant geodesics, compact-group cut loci, interaction pictures, and closest-vector formulations—are established theory. The contribution here is their control-specific synthesis into one auditable benchmark framework, together with the finite principal-phase selector and its stability margin, the pullback to arbitrary accessible closed commuting subtori, the strict/projective f Sim period calculation, and a reproducibility suite designed to expose branch and conditioning failures. The statements below distinguish global theorems from numerical evidence and from device-level interpretation.

The main contributions are as follows.

- (i) A precise L^2 and L^∞ admissibility framework is stated for finite-dimensional closed systems on $SU(n)$. Smooth stationary curves for general right-invariant quadratic metrics satisfy the right-trivialized Euler–Arnold equation; the sign convention, Lax invariants, and energy conservation are derived explicitly.
- (ii) The fully actuated isotropic fixed-time action is solved globally. The optimum is the squared bi-invariant Hilbert–Schmidt distance divided by $2T$ —equivalently the squared norm of a minimum-norm skew-Hermitian logarithm divided by $2T$. A finite principal-phase unwrapping theorem identifies the minimizing branch.

- (iii) The cut locus and conjugate-point structure relevant to branch selection are made explicit for $SU(n)$. A phase-gap theorem quantifies when the minimizing branch is stable under perturbations and when the minimizing logarithm becomes set-valued.
- (iv) Drift is handled by an interaction-picture isometry. When the metric is drift-adjoint invariant, the drifted and driftless fixed-time problems are equivalent. For the isotropic Hilbert–Schmidt metric, this yields an explicit laboratory-frame optimum and an exact norm-bounded minimum-time criterion.
- (v) The maximal-torus sector is solved with anisotropic quadratic weights and coordinate-wise amplitude limits. Weighted action minimization is a nearest-lattice problem, box-limited minimum time is a weighted lattice-Chebyshev problem, and fixed-time synthesis under both constraints is a feasible nearest-lattice problem. A pullback theorem extends these formulas to any accessible closed commuting subtorus and separates strict from projective period lattices.
- (vi) The connection to native trapped-ion and superconducting gates is stated at the effective-Hamiltonian level together with rigorous perturbation bounds and a clear account of what changes under parameter uncertainty, leakage, and Lindblad decoherence.
- (vii) The theory is validated computationally on 1506 Haar-random targets through dimension $n = 128$ with a complete cost-bounded Bellman cross-check for every target and direct Cartesian enumeration for 300 low-dimensional cases. A further 120 random-conjugation cases have analytically prescribed interior, near-cut, or exact-cut spectra; 240 random anisotropic torus instances are solved by independent sphere-decoding/Cartesian and interval-feasibility/Cartesian pairs; and a separate $n = 256$ target stress-tests the dense Schur selector and Bellman certificate. Large tolerance-cluster regressions test branch reporting above the explicit-enumeration cap, and two directional cut-locus tests invoke the implemented Schur selector. Controlled comparisons use compact GRAPE–L-BFGS, a first-order sequential Krotov-type method, randomized CRAB–Powell, and a drift-adapted CRAB basis.

The results are exact calibration cases rather than a replacement for hardware-level numerical optimal control. Their value is precisely that the assumptions, global optimum, branch certificate, and failure modes are all known in advance.

2. Mathematical Setting

The recurring geometric objects have direct control interpretations, which are summarized in Table 1. In particular, a “metric” is not an additional dynamical assumption: it is the positive quadratic form that assigns a cost to each instantaneous control direction. The adjoint action is simply the change of coordinates induced by a rotating interaction frame.

Throughout, $n \geq 2$. Let

$$\mathfrak{su}(n) = \{X \in M_n(\mathbb{C}) : X^\dagger = -X, \operatorname{Tr} X = 0\} \quad (1)$$

be the Lie algebra of $SU(n)$. We use the real Hilbert–Schmidt pairing

$$\langle A, B \rangle := -\operatorname{Tr}(AB), \quad A, B \in \mathfrak{su}(n), \quad (2)$$

which is symmetric, positive definite, and Ad-invariant. The associated norm is $\|X\|^2 = \langle X, X \rangle$. We set $\hbar = 1$ and write Hamiltonians in angular-frequency units, so $X = -iH$ for a traceless Hermitian H ; equivalently, for a physical Hamiltonian, $X = -iH_{\text{phys}}/\hbar$. In these units, (2) is $\operatorname{Tr}(H_A H_B)$.

The controlled propagator is written in right-invariant form as

$$\dot{U}(t) = (\Omega_0 + \Omega_c(t))U(t), \quad U(0) = I, \quad (3)$$

where $\Omega_0 \in \mathfrak{su}(n)$ is a fixed drift and $\Omega_c(t) \in \mathfrak{su}(n)$ is the controlled generator. If $\Omega_c \in L^2([0, T], \mathfrak{su}(n))$, then $\Omega_0 + \Omega_c \in L^1$ on the finite interval $[0, T]$, and (3) has a unique absolutely continuous Carathéodory solution. The solution remains in $SU(n)$ because the coefficient is skew-Hermitian and traceless almost everywhere.

Table 1. Guide to the principal notation and its control-theoretic meaning.

Object	Mathematical Role	Control Interpretation
$SU(n)$	Smooth group of determinant-one unitary matrices	Closed-system gates after removal of global phase
$\mathfrak{su}(n)$	Traceless skew-Hermitian tangent space at the identity	Instantaneous generators $-iH$ of traceless Hamiltonians
$\mathcal{G}$ or G	Positive quadratic metric	Direction-dependent energy or fluence penalty
$\text{Ad}_V X = VXV^\dagger$	Adjoint action of the group on its Lie algebra	The same generator expressed in a rotated frame
Maximal torus	Largest connected commuting subgroup	Simultaneously diagonalizable phase controls
L_n	Integer lattice of trace-zero 2π shifts	Different phase lifts that exponentiate to the same gate
Weighted Voronoi cell	Region where one lattice representative is nearest	Targets for which one logarithm branch is globally cheapest

Definition 1 (Full actuation). *The full-actuation model permits $\Omega_c(t)$ to be any measurable element of $\mathfrak{su}(n)$, subject only to the stated integral or pointwise constraints. Thus, the control distribution is the whole tangent space after right translation.*

Definition 2 (Torus actuation). *The torus-control model fixes the standard Cartan subalgebra*

$$\mathfrak{t} = \left\{ i \text{diag}(x_1, \dots, x_n) : x \in \mathbb{R}^n, \sum_{j=1}^n x_j = 0 \right\}. \quad (4)$$

We identify $\mathfrak{t}$ with the phase hyperplane

$$\mathbb{R}_0^n = \left\{ x \in \mathbb{R}^n : \sum_{j=1}^n x_j = 0 \right\} \quad (5)$$

through $x \mapsto i \text{diag}(x_1, \dots, x_n)$. A torus control has the form

$$\Omega[u(t)] = i \text{diag}(u_1(t), \dots, u_n(t)), \quad u(t) \in \mathbb{R}_0^n. \quad (6)$$

Only targets in this fixed maximal torus are reachable. A conjugate maximal torus is obtained by conjugating all formulas by the same fixed unitary.

For full actuation, a general quadratic control cost is determined by a positive self-adjoint operator $\mathcal{G} : \mathfrak{su}(n) \rightarrow \mathfrak{su}(n)$:

$$J_T[\Omega_c] = \frac{1}{2} \int_0^T \langle \Omega_c(t), \mathcal{G}\Omega_c(t) \rangle dt. \quad (7)$$

The isotropic case is $\mathcal{G} = \text{id}$:

$$J_T[\Omega_c] = \frac{1}{2} \int_0^T \|\Omega_c(t)\|^2 dt. \quad (8)$$

For torus control, let G be a symmetric positive definite bilinear form on $\mathbb{R}_0^n$. Equivalently, after choosing the Euclidean structure inherited from $\mathbb{R}^n$, G is a positive self-adjoint operator on $\mathbb{R}_0^n$. When an ambient symmetric matrix $A \in \mathbb{R}^{n \times n}$ is displayed, it denotes the restricted form $x^\top A y$ on $\mathbb{R}_0^n$; the associated operator on $\mathbb{R}_0^n$ is the compression $P_0 A|_{\mathbb{R}_0^n}$, where

$$P_0 = I_n - \frac{1}{n} \mathbf{1}\mathbf{1}^\top \quad (9)$$

is the orthogonal projection onto $\mathbb{R}_0^n$. The weighted torus action is

$$J_T[u] = \frac{1}{2} \int_0^T u(t)^\top G u(t) dt. \quad (10)$$

Coordinate-wise amplitude bounds are written

$$|u_j(t)| \leq \mu_j, \quad j = 1, \dots, n, \quad (11)$$

with all $\mu_j > 0$; admissibility is then in $L^\infty([0, T], \mathbb{R}_0^n)$.

The integer phase lattice in $\mathbb{R}_0^n$ is

$$L_n = \mathbb{Z}^n \cap \mathbb{R}_0^n = \left\{ m \in \mathbb{Z}^n : \sum_{j=1}^n m_j = 0 \right\}. \quad (12)$$

For $\theta \in \mathbb{R}_0^n$, define

$$D(\theta) = \text{diag}(e^{i\theta_1}, \dots, e^{i\theta_n}). \quad (13)$$

Then, $D(\theta + 2\pi m) = D(\theta)$ for all $m \in L_n$.

For a general $U \in \text{SU}(n)$, choose principal eigenphases $\vartheta_j \in (-\pi, \pi]$ such that

$$U = V \text{diag}(e^{i\vartheta_1}, \dots, e^{i\vartheta_n}) V^\dagger. \quad (14)$$

Because $\det U = 1$, the sum $\Theta = \sum_j \vartheta_j$ is an integer multiple of 2π . A skew-Hermitian logarithm with eigenvalues $i(\vartheta_j + 2\pi m_j)$ lies in $\mathfrak{su}(n)$ exactly when

$$\sum_{j=1}^n (\vartheta_j + 2\pi m_j) = 0. \quad (15)$$

3. Quadratic Right-Invariant Variational Geometry

The next result fixes the sign convention for the Euler–Arnold equation in the finite-dimensional right-invariant setting, following the geometric viewpoint initiated by Arnold [31]. It is a local stationarity statement; the global optima in later sections require additional compactness and branch arguments.

Proposition 1 (Right-trivialized Euler–Arnold equation). *Consider the driftless full-actuation problem*

$$\dot{U}(t) = \Omega(t)U(t), \quad U(0) = I, \quad U(T) = U_f, \quad (16)$$

with action

$$S[\Omega] = \frac{1}{2} \int_0^T \langle \Omega(t), \mathcal{G}\Omega(t) \rangle dt. \quad (17)$$

Let $M = \mathcal{G}\Omega$. Every smooth stationary curve for fixed endpoints satisfies

$$\dot{M} = [\Omega, M]. \quad (18)$$

In the isotropic case $\mathcal{G} = \text{id}$, this reduces to $\dot{\Omega} = 0$.

Proof. Let $U_\varepsilon(t)$ be a smooth endpoint-fixed variation of $U(t)$ and define the right-trivialized variation field $\eta(t) = \partial_\varepsilon U_\varepsilon(t)|_{\varepsilon=0} U(t)^{-1}$. Then, $\eta(0) = \eta(T) = 0$, $\delta U = \eta U$, and, since $\Omega = \dot{U}U^{-1}$,

$$\delta\Omega = \dot{\eta} + [\eta, \Omega]. \quad (19)$$

Therefore,

$$\delta S = \int_0^T \langle M, \dot{\eta} + [\eta, \Omega] \rangle dt \quad (20)$$

$$= - \int_0^T \langle \dot{M}, \eta \rangle dt + \int_0^T \langle M, [\eta, \Omega] \rangle dt. \quad (21)$$

The Ad-invariance of (2) gives

$$\langle M, [\eta, \Omega] \rangle = \langle [\Omega, M], \eta \rangle. \quad (22)$$

Thus,

$$\delta S = - \int_0^T \langle \dot{M} - [\Omega, M], \eta \rangle dt. \quad (23)$$

The fundamental lemma of the calculus of variations gives (18). If $\mathcal{G} = \text{id}$, then $M = \Omega$ and $\dot{\Omega} = [\Omega, \Omega] = 0$. $\square$

Corollary 1 (Lax invariants and energy conservation). *Along every smooth solution of (18), the spectrum of $M(t)$ is independent of t . Moreover,*

$$E(t) = \frac{1}{2} \langle \Omega(t), \mathcal{G}\Omega(t) \rangle = \frac{1}{2} \langle \Omega(t), M(t) \rangle \quad (24)$$

is constant.

Proof. Equation (18) is a Lax equation, so $\text{Tr}(M(t)^k)$ is constant for every positive integer k . Since $M(t)$ is skew-Hermitian, these trace invariants determine its spectrum with multiplicity. For the energy, the self-adjointness of $\mathcal{G}$ gives

$$\dot{E} = \langle \dot{M}, \Omega \rangle = \langle [\Omega, M], \Omega \rangle = \langle M, [\Omega, \Omega] \rangle = 0, \quad (25)$$

again using the Ad-invariance of (2). $\square$

Remark 1 (Normality). *For full actuation, the control map onto the tangent space is an isomorphism after right translation. Thus, the quadratic problem is an ordinary right-invariant Riemannian variational problem on $\text{SU}(n)$, and smooth local minimizers are normal geodesics. For a generic anisotropic $\mathcal{G}$, (18) is a nonlinear boundary-value problem. The global closed forms below rely either on the isotropic metric, which is bi-invariant, or on commutativity in a maximal torus.*

4. Full Isotropic Actuation

Let d denote the Riemannian distance on $\text{SU}(n)$ induced by (2). The next lemma is the geometric core of the fixed-time result.

Lemma 1 (Distance and minimum-norm logarithms). *For every $U \in \text{SU}(n)$,*

$$d(I, U) = \min_{\substack{X \in \mathfrak{su}(n) \\ e^X = U}} \|X\|. \quad (26)$$

Proof. The metric (2) is bi-invariant. For matrix Lie-algebra elements, the Levi–Civita connection is $\nabla_{X^L} Y^L = \frac{1}{2}[X, Y]^L$ on left-invariant fields and $\nabla_{X^R} Y^R = -\frac{1}{2}[X, Y]^R$ on right-invariant fields. In either convention, $\nabla_{X^{L/R}} X^{L/R} = 0$, so geodesics through the identity are one-parameter subgroups $t \mapsto e^{tX}$. Since $\text{SU}(n)$ is compact and connected, Hopf–Rinow gives a minimizing constant-speed geodesic from I to U , and the exponential map is surjective. If the minimizing geodesic is $t \mapsto e^{tX}$ on $[0, 1]$, then $e^X = U$ and its length is $\|X\|$. Conversely, every logarithm X gives the curve e^{tX} from I to U with length $\|X\|$. Taking the infimum over logarithms proves (26); compactness and the discrete branch structure ensure the infimum is attained. $\square$

Theorem 1 (Exact fixed-time isotropic optimum). *Let $U_f \in \text{SU}(n)$, let $T > 0$, and assume full actuation with isotropic cost (8). Then, the exact minimum over all L^2 controls is*

$$J_T^*(U_f) = \frac{d(I, U_f)^2}{2T} = \frac{1}{2T} \min_{\substack{X \in \mathfrak{su}(n) \\ e^X = U_f}} \|X\|^2. \quad (27)$$

If U_f has principal eigenphases as in (14), then

$$J_T^*(U_f) = \frac{1}{2T} \min_{\substack{m \in \mathbb{Z}^n \\ \sum_j (\vartheta_j + 2\pi m_j) = 0}} \sum_{j=1}^n (\vartheta_j + 2\pi m_j)^2. \quad (28)$$

For every minimizing branch m^ , an optimal control is constant:*

$$\Omega^*(t) = \frac{X_\star}{T}, \quad X_\star = V i \text{diag}(\vartheta_1 + 2\pi m_1^*, \dots, \vartheta_n + 2\pi m_n^*) V^\dagger. \quad (29)$$

Proof. For an absolutely continuous trajectory generated by $\Omega \in L^2$, define

$$L[U] = \int_0^T \|\Omega(t)\| dt. \quad (30)$$

Cauchy–Schwarz gives

$$J_T[\Omega] = \frac{1}{2} \int_0^T \|\Omega(t)\|^2 dt \geq \frac{1}{2T} \left(\int_0^T \|\Omega(t)\| dt \right)^2 = \frac{L[U]^2}{2T}. \quad (31)$$

Since $L[U] \geq d(I, U_f)$, every admissible control satisfies $J_T[\Omega] \geq d(I, U_f)^2/(2T)$. By Lemma 1, choose $X_\star \in \mathfrak{su}(n)$ with $e^{X_\star} = U_f$ and $\|X_\star\| = d(I, U_f)$. The curve $U_\star(t) = e^{tX_\star/T}$ has generator $X_\star/T$, reaches U_f , and attains the lower bound.

It remains to identify the logarithm branches. If $X \in \mathfrak{su}(n)$ and $e^X = U_f$, then X commutes with U_f . Both matrices are normal, so they are simultaneously unitarily diagonalizable with arbitrary basis choices inside degenerate eigenspaces. The eigenvalues of X are $i(\vartheta_j + 2\pi m_j)$ for integers m_j after relabeling within equal eigenspaces of U_f . The trace-zero condition is (15), and the unitary invariance of the Hilbert–Schmidt norm gives (28). $\square$

The integer minimization in (28) is finite in closed form.

Theorem 2 (Finite principal-phase unwrapping). Let $\vartheta_j \in (-\pi, \pi]$ be principal phases of $U \in \text{SU}(n)$ and set

$$q = \frac{1}{2\pi} \sum_{j=1}^n \vartheta_j \in \mathbb{Z}. \quad (32)$$

A minimizing branch in (28) is obtained by the following rule.

- (a) If $q = 0$, take $m_j^* = 0$ for all j .
- (b) If $q > 0$, choose q indices with the largest principal phases and set $m_j^* = -1$ on those indices and $m_j^* = 0$ elsewhere.
- (c) If $q < 0$, choose $-q$ indices with the smallest principal phases and set $m_j^* = +1$ on those indices and $m_j^* = 0$ elsewhere.

Every choice within a phase tie at the selected–unselected boundary gives a minimizing branch.

Proof. The constraint is $\sum_j m_j = -q$. Suppose first that $q > 0$. Define $f_j(k) = (\vartheta_j + 2\pi k)^2$. If an admissible vector m has some $m_a > 0$, then, because the total sum is negative, it also has some $m_b < 0$. Replacing (m_a, m_b) by $(m_a - 1, m_b + 1)$ preserves the total sum and does not increase the cost: the first coordinate strictly decreases because $\vartheta_a > -\pi$ and $m_a \geq 1$, while the second coordinate decreases or remains equal when it moves one step toward zero. Hence, no minimizer has a positive component.

Thus, a minimizer has $m_j = -r_j$ with $r_j \geq 0$ and $\sum_j r_j = q$. The marginal cost of the first downward shift in coordinate j is

$$f_j(-1) - f_j(0) = 4\pi(\pi - \vartheta_j), \quad (33)$$

whereas the marginal cost of the r th downward shift, $r \geq 2$, is

$$f_j(-r) - f_j(-(r-1)) = 4\pi((2r-1)\pi - \vartheta_j) \geq 8\pi^2. \quad (34)$$

Every unused first downward shift has cost strictly below $8\pi^2$, since principal phases lie in $(-\pi, \pi]$. Also $q < n$: indeed $\sum_j \vartheta_j \leq n\pi$ and at least one strict inequality holds unless n is even and all phases are π , in which case $q = n/2 < n$. Therefore, if some $r_j \geq 2$, there is an unused coordinate k with $r_k = 0$, and moving one downward shift from j to k lowers the cost. Hence, $r_j \in \{0, 1\}$ for all j . The q chosen coordinates must minimize (33), which is equivalent to choosing the largest ϑ_j .

The case $q < 0$ is the same with upward shifts. The first upward marginal cost is

$$f_j(1) - f_j(0) = 4\pi(\pi + \vartheta_j), \quad (35)$$

so the least-cost coordinates are those with the smallest principal phases. When $q = 0$, any nonzero feasible branch contains both positive and negative components; the two-coordinate reduction above lowers the cost until all components vanish. $\square$

Corollary 2 (Nonuniqueness). The value $J_T^*(U_f)$ is unique, but the optimal logarithm need not be. Nonuniqueness occurs when the finite unwrapping rule has a phase tie at its selected–unselected boundary or when a degenerate eigenspace of U_f can be decomposed into different minimizing branch subspaces. If the spectrum of U_f is simple and the relevant largest or smallest phases in Theorem 2 are separated by strict inequalities, then the minimizing logarithm is unique.

Cut Locus, Conjugate Points, and Branch Stability

For the bi-invariant metric, global nonuniqueness is governed by the Weyl alcove rather than by a failure of the local Euler–Arnold equation. The following statement gives

the concrete $SU(n)$ criterion; it is the matrix form of the standard cut- and conjugate-point description for compact symmetric spaces [5,32].

Proposition 2 (Alcove and conjugate-point criterion). *Let*

$$X_* = iV \operatorname{diag}(\alpha_1, \dots, \alpha_n) V^\dagger, \quad \alpha_1 \geq \dots \geq \alpha_n, \quad \sum_j \alpha_j = 0, \quad (36)$$

be a minimum-norm logarithm of U . Then,

$$\alpha_1 - \alpha_n \leq 2\pi. \quad (37)$$

If the inequality is strict, the minimum-norm logarithm is unique. If equality holds, U lies in the cut locus of the identity and the minimum-norm logarithms form a continuum. More precisely, the extremal phases exponentiate to the same eigenvalue of U . If their multiplicities are r and s , the choices of the r -dimensional branch subspace inside the combined $(r+s)$ -dimensional eigenspace are parametrized by the complex Grassmannian $\operatorname{Gr}(r, r+s)$; independent cut clusters give the corresponding product. Moreover, the differential of the exponential map at a logarithm $X = iV \operatorname{diag}(\alpha) V^\dagger$ is singular exactly when

$$\alpha_j - \alpha_k = 2\pi r \quad \text{for some } j \neq k, \quad r \in \mathbb{Z} \setminus \{0\}. \quad (38)$$

Thus, the first cut wall $\alpha_1 - \alpha_n = 2\pi$ is also a conjugate wall.

Proof. If a trace-zero logarithm has $\alpha_1 - \alpha_n > 2\pi$, replacing α_1 by $\alpha_1 - 2\pi$ and α_n by $\alpha_n + 2\pi$ leaves the exponential and trace unchanged, while the squared norm changes by

$$-4\pi(\alpha_1 - \alpha_n) + 8\pi^2 < 0. \quad (39)$$

Hence, a minimizer must satisfy (37). Under strict inequality, the logarithm is away from the affine cut wall. Every eigenvalue of U then has a unique real phase in the selected interval of width less than 2π ; within a repeated eigenspace, that phase is scalar, so a change of eigenbasis does not change X . Thus, the minimum logarithm is unique even if the ordered point lies on an ordinary Weyl-chamber wall $\alpha_j = \alpha_{j+1}$. At equality, the two-coordinate transfer leaves the norm unchanged. Because the extremal phases have the same exponential, rotating the branch subspace inside their combined eigenspace produces the Grassmannian family stated above. Finally, in left-trivialized form,

$$(d \exp)_X = L_{e^X} \circ \frac{1 - e^{-\operatorname{ad}_X}}{\operatorname{ad}_X}. \quad (40)$$

The eigenvalues of ad_X are $i(\alpha_j - \alpha_k)$, so the displayed operator is singular precisely under (38). $\square$

A quantitative separation margin is useful in floating-point implementations.

Proposition 3 (Branch-gap stability). *Suppose $q > 0$ in (32), and order the principal phases as $\vartheta_{(1)} \geq \dots \geq \vartheta_{(n)}$. Define*

$$\delta = \vartheta_{(q)} - \vartheta_{(q+1)}. \quad (41)$$

If $\delta > 0$, the selected index set in Theorem 2 is unique. The squared-norm gap between the optimum and the nearest branch obtained by exchanging one selected and one unselected phase is $4\pi\delta$, and the corresponding fixed-time action gap is

$$\Delta J = \frac{2\pi\delta}{T}. \quad (42)$$

If consistently matched phase perturbations satisfy $\max_j |\hat{\vartheta}_j - \vartheta_j| < \delta/2$ and do not change q , the selected branch is unchanged. For $q < 0$, the same statement holds after ordering the phases increasingly and using the gap between the $(-q)$ th and $(-q + 1)$ th smallest phases.

Proof. For $q > 0$, exchanging a shifted phase ϑ_a with an unshifted phase ϑ_b changes the squared norm by

$$(\vartheta_b - 2\pi)^2 + \vartheta_a^2 - (\vartheta_a - 2\pi)^2 - \vartheta_b^2 = 4\pi(\vartheta_a - \vartheta_b). \quad (43)$$

The smallest positive exchange uses the two phases adjacent to the selection boundary and gives $4\pi\delta$. Division by $2T$ gives (42). A perturbation smaller than $\delta/2$ in the maximum norm cannot reverse the ordering across that boundary. The case $q < 0$ is analogous. $\square$

Remark 2 (Value stability versus selector stability). *The optimal value is continuous at the cut locus, but no globally continuous single-valued choice of minimizing logarithm exists there. A robust implementation should therefore return the branch margin (41), enumerate all assignments in a declared numerical boundary cluster, and validate each reported candidate by both the trace condition and the endpoint residual. Near an eigenvalue on the principal logarithm branch cut, Schur-based matrix-function algorithms and explicit condition estimates are preferable to unguarded eigenvector formulas [28,30].*

Corollary 3 (Norm-limited minimum time). *Assume full actuation and the pointwise bound $\|\Omega(t)\| \leq \omega_{\max}$ almost everywhere with $\omega_{\max} > 0$. The minimum time to realize $U_f \in \text{SU}(n)$ is*

$$T_{\min}(U_f) = \frac{d(\mathbf{I}, U_f)}{\omega_{\max}} = \frac{1}{\omega_{\max}} \min_{\substack{X \in \text{su}(n) \\ e^X = U_f}} \|X\|. \quad (44)$$

If $U_f \neq \mathbf{I}$, an optimal control is the constant-speed geodesic generator

$$\Omega^*(t) = \omega_{\max} \frac{X_\star}{\|X_\star\|}, \quad (45)$$

where $X_\star$ is a minimum-norm logarithm.

Proof. For every admissible trajectory,

$$d(\mathbf{I}, U_f) \leq \int_0^T \|\Omega(t)\| dt \leq \omega_{\max} T. \quad (46)$$

The constant-speed geodesic generated by $X_\star$ attains equality. $\square$

Corollary 4 (Projective gates). *Let a physical gate be specified by a class $[Q] \in \text{PU}(n)$ with representative $Q \in \text{U}(n)$. Choose one determinant-one representative*

$$\hat{Q} = (\det Q)^{-1/n} Q \in \text{SU}(n), \quad (47)$$

where one branch of the n th root is fixed. The exact full-actuation isotropic fixed-time cost for implementing the projective gate is

$$J_{T,\text{PU}}^*([Q]) = \frac{1}{2T} \min_{k=0,\dots,n-1} d(\mathbf{I}, e^{2\pi i k/n} \hat{Q})^2. \quad (48)$$

The corresponding norm-limited time is the square root of the numerator in (48) divided by $\omega_{\max}$.

Proof. The determinant-one representatives of the projective class are exactly $e^{2\pi i k/n} \widehat{Q}$, $k = 0, \dots, n-1$. Apply Theorem 1 or Corollary 3 to each representative and choose the least value. $\square$

5. Drift-Adapted Full-Control Synthesis

The drifted system is

$$\dot{U}(t) = (\Omega_0 + \Omega_c(t))U(t), \quad U(0) = I. \quad (49)$$

The cost penalizes only Ω_c —not the fixed drift Ω_0 .

Theorem 3 (Interaction-picture isometry). *Define*

$$V(t) = e^{-t\Omega_0}U(t). \quad (50)$$

Then,

$$\dot{V}(t) = \tilde{\Omega}_c(t)V(t), \quad \tilde{\Omega}_c(t) = \text{Ad}_{e^{-t\Omega_0}}\Omega_c(t), \quad (51)$$

with endpoint

$$V(T) = e^{-T\Omega_0}U_f. \quad (52)$$

If the quadratic metric is invariant under the drift adjoint action,

$$\langle \text{Ad}_{e^{-t\Omega_0}}X, \mathcal{G} \text{Ad}_{e^{-t\Omega_0}}X \rangle = \langle X, \mathcal{G}X \rangle \quad \forall X \in \mathfrak{su}(n), \quad t \in [0, T], \quad (53)$$

then the drifted and driftless fixed-time problems are isometric:

$$J_T^*(U_f; \Omega_0, \mathcal{G}) = J_T^*(e^{-T\Omega_0}U_f; 0, \mathcal{G}). \quad (54)$$

In particular, for the isotropic metric,

$$J_T^*(U_f; \Omega_0) = \frac{1}{2T} \min_{\substack{X \in \mathfrak{su}(n) \\ e^X = e^{-T\Omega_0}U_f}} \|X\|^2, \quad (55)$$

and an optimal laboratory-frame control is

$$\Omega_c^*(t) = \text{Ad}_{e^{t\Omega_0}}\left(\frac{X_\star}{T}\right), \quad (56)$$

where $X_\star$ is a minimum-norm logarithm of $e^{-T\Omega_0}U_f$.

Proof. Differentiating (50) gives

$$\dot{V} = -\Omega_0 e^{-t\Omega_0}U + e^{-t\Omega_0}(\Omega_0 + \Omega_c)U \quad (57)$$

$$= e^{-t\Omega_0}\Omega_c e^{t\Omega_0}V = \text{Ad}_{e^{-t\Omega_0}}\Omega_c V. \quad (58)$$

This proves (51); (52) is immediate. The map $\Omega_c \mapsto \tilde{\Omega}_c$ is bijective with inverse $\tilde{\Omega}_c \mapsto \text{Ad}_{e^{t\Omega_0}}\tilde{\Omega}_c$. Under (53), it preserves the cost integrand and therefore preserves admissible fixed-time actions and endpoints as in (54). For $\mathcal{G} = \text{id}$, invariance follows from the Ad-invariance of (2). Applying Theorem 1 in the interaction picture gives (55), and transforming the constant interaction-picture generator back to the laboratory frame gives (56). $\square$

Remark 3 (Non-invariant anisotropy). The transformation (50) always removes the drift from the differential equation, but it is not generally cost-neutral. Without (53), the interaction-picture cost is

$$\frac{1}{2} \int_0^T \left\langle \tilde{\Omega}_c(t), (\text{Ad}_{e^{it\Omega_0}}^* \mathcal{G} \text{Ad}_{e^{it\Omega_0}}) \tilde{\Omega}_c(t) \right\rangle dt, \quad (59)$$

where $*$ denotes the adjoint with respect to (2). Thus, arbitrary drift is dynamically removable but generally converts a time-independent anisotropic metric into a time-dependent one.

Theorem 4 (Exact isotropic drifted speed limit). Under full actuation and $\|\Omega_c(t)\| \leq \omega_{\max}$ almost everywhere, with $\omega_{\max} > 0$, the minimum time in the presence of drift is

$$T_{\min}(U_f; \Omega_0, \omega_{\max}) = \min \left\{ T \geq 0 : d(I, e^{-T\Omega_0} U_f) \leq \omega_{\max} T \right\}. \quad (60)$$

For any minimizing $T > 0$, an optimal control is (56) with X_* a minimum-norm logarithm of $e^{-T\Omega_0} U_f$.

Proof. At a fixed time T , Theorem 3 and Corollary 3 show that the target is reachable under the bound if and only if

$$d(I, e^{-T\Omega_0} U_f) \leq \omega_{\max} T. \quad (61)$$

The left-hand side is continuous and bounded on $[0, \infty)$, while the right-hand side grows linearly. Hence, the feasible set is closed and nonempty, and its minimum is (60). The fixed-time construction at the minimizing time gives the control. If the minimizing time is positive, the feasibility inequality is in fact an equality unless the same target is reachable at time zero. $\square$

6. Weighted Maximal-Torus Synthesis

In the torus model, all generators commute. The time-ordered exponential reduces to an ordinary exponential, but compact phase topology remains.

6.1. Weighted Fixed-Time Action

For $u \in L^2([0, T], \mathbb{R}_0^n)$,

$$U(T) = \exp \left(i \text{diag} \left(\int_0^T u_1(t) dt, \dots, \int_0^T u_n(t) dt \right) \right), \quad (62)$$

so the endpoint depends only on the phase integral

$$x = \int_0^T u(t) dt \in \mathbb{R}_0^n. \quad (63)$$

Theorem 5 (Exact weighted torus action). Let $T > 0$, let $U_f = D(\theta)$ with $\theta \in \mathbb{R}_0^n$, and let G be symmetric positive definite on $\mathbb{R}_0^n$. The exact minimum of (10) over all torus controls on $[0, T]$ is

$$J_T^*(D(\theta)) = \frac{1}{2T} \min_{m \in L_n} (\theta + 2\pi m)^\top G (\theta + 2\pi m). \quad (64)$$

Every minimizing branch m^* gives the optimal constant control

$$u^*(t) = \frac{\theta + 2\pi m^*}{T}. \quad (65)$$

Proof. A control reaches $D(\theta)$ if and only if

$$x = \int_0^T u(t) dt = \theta + 2\pi m \quad \text{for some } m \in L_n. \quad (66)$$

For fixed x , Jensen's inequality in the Hilbert space defined by G gives

$$\int_0^T u(t)^\top G u(t) dt \geq \frac{1}{T} x^\top G x, \quad (67)$$

with equality for the constant control $u(t) = x/T$. Minimization over the lattice representatives in (66) proves (64). A minimizer exists because the positive quadratic form is proper on $\mathbb{R}_0^n$ and L_n is discrete. $\square$

Corollary 5 (Weighted Voronoi branch geometry). *Define the lifted value function*

$$\mathcal{J}_T(\theta) = \frac{1}{2T} \min_{m \in L_n} (\theta + 2\pi m)^\top G (\theta + 2\pi m), \quad \theta \in \mathbb{R}_0^n. \quad (68)$$

Then, $\mathcal{J}_T$ is $2\pi L_n$ -periodic, continuous, and piecewise quadratic. The branch cell for $m \in L_n$ is

$$\mathcal{V}_m^G = \left\{ \theta \in \mathbb{R}_0^n : (\theta + 2\pi m)^\top G (\theta + 2\pi m) \leq (\theta + 2\pi \ell)^\top G (\theta + 2\pi \ell) \quad \forall \ell \in L_n \right\}. \quad (69)$$

For distinct $m, \ell \in L_n$, the bisector between the corresponding quadratic branches is

$$(m - \ell)^\top G \theta = -\pi(m^\top G m - \ell^\top G \ell). \quad (70)$$

Inside the interior of a cell, the Euclidean gradient on $\mathbb{R}_0^n$ is

$$\nabla_{\mathbb{R}_0^n} \mathcal{J}_T(\theta) = \frac{1}{T} G(\theta + 2\pi m), \quad (71)$$

where G is understood as the self-adjoint operator on $\mathbb{R}_0^n$. If G is represented by an ambient matrix A , then (71) means $P_0 A(\theta + 2\pi m)/T$.

Proof. Periodicity follows by relabeling $m \mapsto m - k$ when θ is replaced by $\theta + 2\pi k$. Local finiteness of the relevant branches follows from the properness of the quadratic form on the discrete lattice. On a region where one branch is uniquely minimizing, (68) is a single quadratic polynomial whose gradient is (71). The equality of the m and ℓ branches gives

$$(\theta + 2\pi m)^\top G (\theta + 2\pi m) = (\theta + 2\pi \ell)^\top G (\theta + 2\pi \ell), \quad (72)$$

which reduces to (70). A locally finite minimum of continuous functions is continuous. $\square$

Lemma 2 (Finite certification of torus branches). *Let $\lambda_{\min}(G) > 0$ be the smallest eigenvalue of the self-adjoint operator G on $\mathbb{R}_0^n$. If $m_0 \in L_n$ is any candidate branch and*

$$R_0^2 = (\theta + 2\pi m_0)^\top G (\theta + 2\pi m_0), \quad (73)$$

then every minimizer of (64) lies in the finite set

$$\left\{ m \in L_n : \|\theta + 2\pi m\|_2 \leq \frac{R_0}{\sqrt{\lambda_{\min}(G)}} \right\}. \quad (74)$$

For the box-time problem below, if $S_0 = \|\theta + 2\pi m_0\|_{\mu, \infty}$, every minimizer lies in

$$\{m \in L_n : |\theta_j + 2\pi m_j| \leq \mu_j S_0 \text{ for all } j\}. \quad (75)$$

Proof. A weighted-action minimizer cannot have value larger than the value of m_0 , and

$$(\theta + 2\pi m)^\top G(\theta + 2\pi m) \geq \lambda_{\min}(G) \|\theta + 2\pi m\|_2^2. \quad (76)$$

This gives (74). The box statement follows directly from the definition of S_0 and the coordinate bounds. Both sets are finite because L_n is discrete. $\square$

6.2. Box-Constrained Minimum Time and Fixed-Time Action

Define the weighted coordinate norm on $\mathbb{R}_0^n$ by

$$\|x\|_{\mu, \infty} = \max_{1 \leq j \leq n} \frac{|x_j|}{\mu_j}. \quad (77)$$

Theorem 6 (Exact box-constrained torus time). *Under the box constraint (11), the reachable phase integrals at time T are exactly*

$$\mathcal{B}_\mu(T) = \{x \in \mathbb{R}_0^n : |x_j| \leq \mu_j T \text{ for } j = 1, \dots, n\}. \quad (78)$$

Consequently, for $D(\theta)$ with $\theta \in \mathbb{R}_0^n$, the exact minimum time is

$$T_{\min}^{\text{box}}(D(\theta)) = \min_{m \in L_n} \|\theta + 2\pi m\|_{\mu, \infty} = \min_{m \in L_n} \max_{1 \leq j \leq n} \frac{|\theta_j + 2\pi m_j|}{\mu_j}. \quad (79)$$

Every minimizing branch m^* gives a time-optimal constant control

$$u_j^*(t) = \frac{\theta_j + 2\pi m_j^*}{T_{\min}^{\text{box}}(D(\theta))}, \quad j = 1, \dots, n, \quad (80)$$

with the zero control if the target is the identity and the minimum time is zero.

Proof. If u satisfies (11), then its phase integral $x = \int_0^T u(t) dt$ satisfies

$$|x_j| \leq \int_0^T |u_j(t)| dt \leq \mu_j T, \quad (81)$$

so $x \in \mathcal{B}_\mu(T)$. Conversely, if $x \in \mathcal{B}_\mu(T)$ and $T > 0$, the constant control $u(t) = x/T$ lies in $\mathbb{R}_0^n$ and satisfies (11). Thus, (78) is exact. The target is reachable at time T if and only if some representative $\theta + 2\pi m$ lies in (78). For fixed m , the least such time is $\|\theta + 2\pi m\|_{\mu, \infty}$; minimizing over $m \in L_n$ proves (79). Existence follows from the properness of the norm on the discrete lattice coset, and (80) attains the value. $\square$

Corollary 6 (Fixed time with both quadratic cost and box bounds). *Let $T > 0$ and let*

$$\mathcal{F}_T(\theta) = \{m \in L_n : |\theta_j + 2\pi m_j| \leq \mu_j T \text{ for all } j\}. \quad (82)$$

Under the cost (10) and the box constraint (11), the target $D(\theta)$ is reachable at time T if and only if $\mathcal{F}_T(\theta)$ is nonempty. When it is nonempty,

$$J_{T, \mu}^*(D(\theta)) = \frac{1}{2T} \min_{m \in \mathcal{F}_T(\theta)} (\theta + 2\pi m)^\top G(\theta + 2\pi m), \quad (83)$$

and every minimizing branch is attained by the constant control $(\theta + 2\pi m)/T$.

Proof. The reachability statement is exactly Theorem 6 at fixed T . Conditional on a feasible branch $x = \theta + 2\pi m$, Jensen's inequality in the proof of Theorem 5 shows that the constant control minimizes the quadratic action among all controls with integral x . The constant control is feasible precisely because $m \in \mathcal{F}_T(\theta)$. $\square$

6.3. Closed Commuting Subsectors

The preceding formulas must be pulled back to the actually accessible commuting sector. Minimizing over the full maximal torus and then restricting the resulting pulse is generally invalid because the cheapest full-torus lift need not lie in a lower-dimensional control subspace. The next theorem supplies the required reduction for native-gate sectors.

Theorem 7 (Exact reduction on a closed commuting subtorus). *Let $\mathcal{B} : \mathbb{R}^r \rightarrow \mathfrak{su}(n)$ be an injective linear map whose image is Abelian, and assume*

$$\mathbb{T}_{\mathcal{B}} = \exp(\mathcal{B}(\mathbb{R}^r)) \quad (84)$$

is a closed r -dimensional torus. Its strict period lattice and projective period lattice are

$$\Lambda_{\mathcal{B}} = \{\lambda \in \mathbb{R}^r : \exp(\mathcal{B}\lambda) = \mathbf{I}\}, \quad (85)$$

$$\Lambda_{\mathcal{B}}^{\text{proj}} = \{\lambda \in \mathbb{R}^r : \exp(\mathcal{B}\lambda) \in Z(\text{SU}(n))\}, \quad (86)$$

where $Z(\text{SU}(n)) = \{e^{2\pi i k/n} \mathbf{I} : k = 0, \dots, n-1\}$. Both sets are full-rank lattices with $\Lambda_{\mathcal{B}} \subseteq \Lambda_{\mathcal{B}}^{\text{proj}}$ and finite quotient $\Lambda_{\mathcal{B}}^{\text{proj}} / \Lambda_{\mathcal{B}}$; equivalently, the projective lattice is a finite union of strict-lattice cosets.

Consider

$$\dot{U}(t) = \mathcal{B}(v(t))U(t), \quad J_T[v] = \frac{1}{2} \int_0^T v(t)^\top W v(t) dt, \quad (87)$$

where W is positive definite on $\mathbb{R}^r$. For $T > 0$ and $U_f = \exp(\mathcal{B}\beta)$,

$$J_{T,\mathcal{B}}^*(U_f) = \frac{1}{2T} \min_{\lambda \in \Lambda_{\mathcal{B}}} (\beta + \lambda)^\top W (\beta + \lambda). \quad (88)$$

Every minimizing lift is attained by the constant coefficient control $v(t) = (\beta + \lambda)/T$. If $|v_a(t)| \leq v_a$ with $v_a > 0$, then

$$T_{\min,\mathcal{B}}(U_f) = \min_{\lambda \in \Lambda_{\mathcal{B}}} \max_{1 \leq a \leq r} \frac{|\beta_a + \lambda_a|}{v_a}. \quad (89)$$

At fixed T , imposing both the quadratic cost and these box bounds restricts (88) to the finite feasible set satisfying $|\beta_a + \lambda_a| \leq v_a T$. For a projective target, the same three statements hold with $\Lambda_{\mathcal{B}}$ replaced by $\Lambda_{\mathcal{B}}^{\text{proj}}$.

Proof. Because the generators in $\mathcal{B}(\mathbb{R}^r)$ commute, the endpoint depends only on

$$y = \int_0^T v(t) dt, \quad U(T) = \exp(\mathcal{B}y). \quad (90)$$

The strict endpoint condition $U(T) = U_f$ is equivalent to $y - \beta \in \Lambda_{\mathcal{B}}$, while equality up to the global phase is equivalent to $y - \beta \in \Lambda_{\mathcal{B}}^{\text{proj}}$. The closedness of $\mathbb{T}_{\mathcal{B}}$ implies that the

kernel of the covering homomorphism $\mathbb{R}^r \rightarrow \mathbb{T}_B$ is a discrete cocompact, hence full-rank, lattice. For a fixed lift y , Jensen's inequality gives

$$\int_0^T v(t)^\top W v(t) dt \geq \frac{1}{T} y^\top W y, \quad (91)$$

with equality for $v(t) = y/T$. This proves (88). The reachable coefficient integrals under the box constraint are exactly $|y_a| \leq \nu_a T$, which are again attained by a constant control; minimizing the required T proves (89) and the fixed-time statement. Finally, the preimage of the finite subgroup $Z(\mathrm{SU}(n)) \cap \mathbb{T}_B$ is an additive subgroup containing Λ_B with finite quotient. It is therefore a full-rank lattice and, equivalently, a finite union of Λ_B -cosets. $\square$

Remark 4 (Two distinct lattice geometries). *The action problem uses the ellipsoidal norm induced by G , whereas the box-time problem uses the polyhedral norm (77). The same target can therefore select different logarithm branches for fluence minimization and for time minimization. The coincidence of the minimizer in a specific example is not a structural theorem.*

7. Analytic Examples and Benchmark Definitions

7.1. Drifted Qubit Phase Gate

Let $\sigma_x, \sigma_y, \sigma_z$ be the Pauli matrices and consider

$$\Omega_0 = -\frac{i\Delta}{2}\sigma_x, \quad U_f = e^{-i\phi\sigma_z/2}, \quad (92)$$

with full isotropic actuation. At a prescribed duration T , Theorem 3 reduces the problem to the logarithm of

$$W(T) = e^{-T\Omega_0} U_f = e^{i\Delta T\sigma_x/2} e^{-i\phi\sigma_z/2}. \quad (93)$$

Write

$$W(T) = \cos \frac{\chi(T)}{2} \mathbf{I} - i \sin \frac{\chi(T)}{2} \hat{n}(T) \cdot \sigma, \quad \chi(T) \in [0, 2\pi]. \quad (94)$$

Then,

$$\cos \frac{\chi(T)}{2} = \cos \frac{\Delta T}{2} \cos \frac{\phi}{2}. \quad (95)$$

For $0 < \chi(T) < 2\pi$, a minimum logarithm is

$$X_\star(T) = -\frac{i\chi(T)}{2} \hat{n}(T) \cdot \sigma. \quad (96)$$

Since $\| -i(\chi/2) \hat{n} \cdot \sigma \|^2 = \chi^2/2$,

$$J_T^\star = \frac{\chi(T)^2}{4T}. \quad (97)$$

The laboratory-frame optimum is

$$\Omega_c^\star(t) = \mathrm{Ad}_{e^{-i\Delta t\sigma_x/2}} \left(-\frac{i\chi(T)}{2T} \hat{n}(T) \cdot \sigma \right). \quad (98)$$

If $W(T) = \mathbf{I}$, the optimum is the zero interaction-picture control. If $W(T) = -\mathbf{I}$, equivalently $\chi(T) = 2\pi$, the axis $\hat{n}$ is not unique; every unit vector gives a minimum logarithm of norm squared $2\pi^2$, as expected at the cut locus of $\mathrm{SU}(2)$.

This example tests whether a numerical method handles drift in the cost-correct way. The interaction-picture optimum is constant, but the laboratory-frame optimum rotates by the drift adjoint action.

7.2. Weighted Qutrit Phase Synthesis and Branch Switching

Consider the diagonal qutrit target

$$U_f = \text{diag}(e^{7\pi i/5}, e^{\pi i/5}, e^{-8\pi i/5}). \quad (99)$$

The lift

$$\theta = \left(\frac{7\pi}{5}, \frac{\pi}{5}, -\frac{8\pi}{5}\right) \in \mathbb{R}_0^3 \quad (100)$$

is not the principal trace-zero lift. The principal trace-zero lift is obtained by the lattice vector $m = (-1, 0, 1)$:

$$\theta + 2\pi(-1, 0, 1) = \left(-\frac{3\pi}{5}, \frac{\pi}{5}, \frac{2\pi}{5}\right). \quad (101)$$

Let the quadratic form be the restriction of the ambient diagonal form

$$G = \text{diag}(1, 1, 4) \quad (102)$$

to $\mathbb{R}_0^3$, and let

$$\mu = (1, 1, 1/2). \quad (103)$$

For the branch $m = (0, 0, 0)$,

$$\theta^\top G \theta = \frac{306\pi^2}{25}. \quad (104)$$

For $m^* = (-1, 0, 1)$,

$$(\theta + 2\pi m^*)^\top G (\theta + 2\pi m^*) = \frac{26\pi^2}{25}. \quad (105)$$

Appendix B gives a finite certificate that this is the global lattice minimum. Hence,

$$J_T^*(U_f) = \frac{13\pi^2}{25T}, \quad u^*(t) = \frac{1}{T} \left(-\frac{3\pi}{5}, \frac{\pi}{5}, \frac{2\pi}{5}\right). \quad (106)$$

For the box constraint, Theorem 6 gives

$$T_{\min}^{\text{box}}(U_f) = \frac{4\pi}{5}, \quad u^*(t) = \left(-\frac{3}{4}, \frac{1}{4}, \frac{1}{2}\right), \quad (107)$$

again with $m^* = (-1, 0, 1)$. The third component saturates its amplitude bound. Representative branches are compared in Table 2.

Table 2. Representative branches for the qutrit target (99). The branch $(-1, 0, 1)$ is globally optimal for both the displayed weighted action and the displayed box-constrained time.

m	$\theta + 2\pi m$	$2TJ_T$	$T_{\min}^{\text{box}}$
(0,0,0)	$\left(\frac{7\pi}{5}, \frac{\pi}{5}, -\frac{8\pi}{5}\right)$	$\frac{306\pi^2}{25}$	$\frac{16\pi}{5}$
(-1,0,1)	$\left(-\frac{3\pi}{5}, \frac{\pi}{5}, \frac{2\pi}{5}\right)$	$\frac{26\pi^2}{25}$	$\frac{4\pi}{5}$
(0,-1,1)	$\left(\frac{7\pi}{5}, -\frac{9\pi}{5}, \frac{2\pi}{5}\right)$	$\frac{146\pi^2}{25}$	$\frac{9\pi}{5}$

Starting from the nonprincipal lift (100), a local method that never changes logarithm branch reports a weighted action larger by the factor 306/26.

7.3. Benchmark Definitions for Numerical Solvers

The formulas above define reproducible tests for gradient, sequential-update, random-basis, shooting, or collocation solvers. The numerical study in Section 9 instantiates these tests with fixed seeds.

- (a) *Full isotropic test.* Draw $U_f \in \text{SU}(n)$, compute principal eigenphases, unwrap them by Theorem 2, and compare the solver's fixed-time cost with (27). The residual should be measured against the globally minimizing branch rather than merely against the principal matrix logarithm.
- (b) *Drift test.* Draw $\Omega_0 \in \mathfrak{su}(n)$, replace the endpoint by $e^{-T\Omega_0}U_f$, and apply the same logarithmic certificate. A correct laboratory-frame solution is the adjoint transport (56) of a constant interaction-picture generator.
- (c) *Weighted torus test.* Choose G and a target lift $\theta \in \mathbb{R}_0^n$, solve the finite nearest-lattice problem certified by Lemma 2, and test whether the numerical method crosses the branch-switching hyperplanes (70) when required.
- (d) *Box-time test.* Choose amplitudes μ_j , solve the finite weighted lattice-Chebyshev problem (79), and compare the reported minimum time and saturated coordinates with the analytic branch.
- (e) *Fixed-time box-action test.* At a prescribed T , use (83). This test distinguishes algorithms that find a low-action logarithm branch from algorithms that also respect amplitude feasibility.

8. Practical Interpretation for Native Multiqubit Gates

The dimension n is the Hilbert-space dimension, so a register of k ideal qubits has $n = 2^k$. Full actuation on $\text{SU}(2^k)$ would require independent access to all $4^k - 1$ traceless Pauli directions and is therefore not a realistic hardware assumption except in very small effective subspaces. The exact full-control value is an ideal lower bound for a sparse closed plant only when that plant is an admissible subset of the full-control model under the same endpoint, duration, and cost normalization. Its direct physical use occurs after a native gate has been reduced to a fully actuated block or to a commuting effective sector. Table 3 summarizes the resulting division between direct use of the formulas and use as a benchmark for a larger numerical model.

Table 3. Connection between common native-gate effective models and the present exact problems. “Direct” means exact only for the displayed calibrated effective Hamiltonian; omitted effects must be restored in a device-level numerical plant.

Platform or Reduction	Commuting Effective Sector	Use and Principal Exclusions
Fixed-axis Mølmer-Sørensen gate	One generator S_ϕ^2 , or same-axis pair terms $\sigma_\phi^{(j)} \sigma_\phi^{(\ell)}$	Direct period-lattice calculation on the accessible closed subtorus after motional closure; excludes spectator-mode excitation, detuning, heating, and optical waveform constraints.
Echoed cross-resonance	Ideally isolated dominant Pauli generator, commonly ZX after frame and echo conventions	Direct one-parameter lower bound only after residual terms are negligible; otherwise, a seed and regression test for a sparse noncommuting model.
Tunable-coupler or fSim gate	Exchange $XX + YY$ together with a commuting conditional phase	Direct two-coordinate closed-subtorus calculation with distinct strict-SU(4) and projective period lattices; excludes leakage, transfer functions, bandwidth, and calibration drift.
Two-qubit Cartan core	Mutually commuting XX , YY , and ZZ	Exact nonlocal-torus calculation followed by a finite minimization over Weyl/local-equivalence representatives; local-gate costs must be specified separately.

8.1. Trapped-Ion Effective Interactions

For a fixed optical phase and a closed motional trajectory, a Mølmer–Sørensen entangler has the effective form

$$U_{\text{MS}}(\chi) = \exp(-i\chi S_\phi^2), \quad S_\phi = \frac{1}{2} \sum_{j=1}^k \sigma_\phi^{(j)}, \quad (108)$$

up to local phases and model-dependent corrections [21,22]. Apart from an identity contribution, S_ϕ^2 is a sum of same-axis pair couplings $\sigma_\phi^{(j)} \sigma_\phi^{(\ell)}$, and all such Pauli strings commute. With one collective coupling, (108) is a one-parameter subgroup; with independently calibrated commuting pair couplings, it is a higher-dimensional closed subtorus. Applying Theorem 7 to the experimentally accessible coefficient map—rather than minimizing over an inaccessible full Cartan algebra—determines the least-cost period-lattice lift and, with calibrated amplitude bounds, the ideal minimum time. The theorem does not replace the pulse shaping needed to close the motional trajectory, suppress spectator modes, or compensate optical detuning and heating. In those more complete models, the analytic gate is a nominal reference and a branch-aware warm start for numerical robust control; it is a rigorous lower bound only for a closed constrained model with the same endpoint and objective.

8.2. Superconducting Effective Interactions

Echoed cross-resonance gates are commonly described by a dominant entangling term together with residual single-qubit and two-qubit Pauli terms, whereas tunable-coupler devices can realize continuous families such as $f\text{Sim}$ gates [23–25]. If calibration and echo sequences isolate a single effective generator, Theorem 7 gives the ideal one-parameter action and speed limit. A concrete two-parameter example is the convention

$$U_{f\text{Sim}}(\theta, \varphi) = \exp\left[-\frac{i\theta}{2}(XX + YY)\right] \exp[-i\varphi |11\rangle\langle 11|]. \quad (109)$$

The two displayed generators commute. The standard matrix has determinant $e^{-i\varphi}$, so a determinant-one representative is

$$\tilde{U}_{f\text{Sim}}(\theta, \varphi) = e^{i\varphi/4} U_{f\text{Sim}}(\theta, \varphi) = \exp(\theta A_\theta + \varphi A_\varphi), \quad (110)$$

where

$$A_\theta = -\frac{i}{2}(XX + YY), \quad A_\varphi = -i(|11\rangle\langle 11| - I/4), \quad (111)$$

and $|11\rangle\langle 11| - I/4 = (-ZI - IZ + ZZ)/4$ still commutes with $XX + YY$.

Corollary 7 (Strict and projective $f\text{Sim}$ period lattices). *For the coefficient map $\mathcal{B}(a, b) = aA_\theta + bA_\varphi$,*

$$\Lambda_{f\text{Sim}} = (2\pi\mathbb{Z}) \times (8\pi\mathbb{Z}), \quad \Lambda_{f\text{Sim}}^{\text{proj}} = (2\pi\mathbb{Z}) \times (2\pi\mathbb{Z}). \quad (112)$$

Consequently, if W is the calibrated positive coefficient metric, the exact strict-SU(4) fixed-time action is

$$J_{T,f\text{Sim}}^* = \frac{1}{2T} \min_{k, \ell \in \mathbb{Z}} \begin{pmatrix} \theta + 2\pi k \\ \varphi + 8\pi \ell \end{pmatrix}^\top W \begin{pmatrix} \theta + 2\pi k \\ \varphi + 8\pi \ell \end{pmatrix}. \quad (113)$$

For the physical projective gate, replace $8\pi\ell$ by $2\pi\ell$. Under coefficient bounds $|v_\theta| \leq \nu_\theta$ and $|v_\varphi| \leq \nu_\varphi$, the corresponding exact minimum-time formula is the minimum over the same strict or projective lattice of

$$\max\left\{\frac{|\theta + 2\pi k|}{\nu_\theta}, \frac{|\varphi + 8\pi\ell|}{\nu_\varphi}\right\}, \quad (114)$$

again replacing $8\pi\ell$ by $2\pi\ell$ projectively.

Proof. In the joint eigenbasis $|00\rangle, (|01\rangle + |10\rangle)/\sqrt{2}, (|01\rangle - |10\rangle)/\sqrt{2}, |11\rangle$, the four eigenphases of (110) are

$$\frac{\varphi}{4}, \quad \frac{\varphi}{4} - \theta, \quad \frac{\varphi}{4} + \theta, \quad -\frac{3\varphi}{4}. \quad (115)$$

Equality of all four exponentials up to a common phase requires $\theta \in 2\pi\mathbb{Z}$ and $\varphi \in 2\pi\mathbb{Z}$, proving the projective lattice. Strict equality to the identity additionally requires $e^{i\varphi/4} = 1$; hence, $\varphi \in 8\pi\mathbb{Z}$. The action and time formulas are Theorem 7 for these two lattices. $\square$

Thus, the ideal $f\text{Sim}$ family is an exact two-dimensional closed-subtorus problem once the coefficient metric and bounds have been calibrated. Noncommuting residuals, leakage levels, transfer functions, and bandwidth limits break this exact reduction and must remain in the numerical plant model.

8.3. Two-Qubit Cartan Reduction

Any ideal two-qubit gate admits a local–nonlocal decomposition of the form

$$Q = (k_1 \otimes k_2) \exp\left[-\frac{i}{2}(c_1 XX + c_2 YY + c_3 ZZ)\right] (k_3 \otimes k_4), \quad (116)$$

where the three nonlocal generators commute [6,26]. If local gates are treated as fast or assigned a separate cost, the nonlocal core lies in a conjugate maximal torus, and the weighted nearest-lattice construction can be applied to (c_1, c_2, c_3) with hardware-derived weights and bounds. Weyl symmetries and local equivalence add a finite outer minimization over canonical representatives. This provides a concrete route from the present phase-lattice geometry to native two-qubit compilation while preserving the distinction between an exact effective model and the waveform-level realization.

9. Numerical Validation and Sensitivity Analysis

Complex Gaussian QR factorization gives a Haar matrix $Q \in \text{U}(n)$ [33]; the code then sets $F(Q) = e^{-i \text{Arg det}(Q)/n} Q \in \text{SU}(n)$. This determinant correction preserves Haar sampling because $F(SQ) = SF(Q)$ for every $S \in \text{SU}(n)$, so the push-forward law is left-invariant and hence Haar on $\text{SU}(n)$. Principal phases were extracted with a complex Schur decomposition, canonicalized to $(-\pi, \pi]$, and processed by Theorem 2. The minimum-cost representative is selected from the actual floating-point phase ordering; the declared 10^{-10} tolerance is used only to describe numerical ambiguity at the selected–unselected boundary and never to replace that representative. Unless otherwise stated, timing excludes random-target generation and all independent validation passes, and it includes the Schur decomposition, branch selection, and endpoint reconstruction. After one unreported warm-up call per dimension, each target was evaluated five times, and its median was used before aggregation. The fixed random seed was 4,341,260.

For pulse-optimization comparisons, the endpoint error was

$$\mathcal{E}(U) = 1 - \frac{1}{n} \Re \text{Tr}(U_f^\dagger U) = \frac{\|U - U_f\|_F^2}{2n}, \quad (117)$$

which is phase-sensitive and therefore tests the stated $SU(n)$ endpoint rather than only its projective class. The compact reference implementations used a common piecewise-constant model and two-step penalty continuation. “GRAPE–L-BFGS” denotes concurrent exact NumPy/SciPy gradients—matrix-exponential Fréchet derivatives for the full system and analytic phase gradients for the commuting qutrit—followed by L-BFGS-B. “Krotov-type” denotes a first-order sequential discretize-then-optimize construction: old-field costates propagate backward, updated-field states propagate forward, and each time slice approximately minimizes the discrete Krotov auxiliary function to the stated local L-BFGS-B budget using exact matrix-exponential Fréchet derivatives, the same physical fluence, and a positive proximal metric. Here, “exact discrete” refers to the exact propagation and differentiation of the chosen piecewise-constant model rather than to the symbolic or unlimited-accuracy minimization of each local subproblem. A full sweep is retained only if the physical fluence-regularized objective is non-increasing within the declared tolerance. “CRAB–Powell” denotes a constant plus one sine/cosine pair per generator optimized by derivative-free Powell search [14,17,19,20]. We report both a randomized harmonic and, for the drifted qubit, a drift-adapted frequency $\Delta = 1.1$; the latter tests the effect of matching the basis to the known rotating-frame structure.

9.1. Random-Target Verification and Higher-Dimensional Scaling

Table 4 summarizes 1506 independent Haar-random targets. Every target also underwent a complete cost-bounded Bellman cross-check. For an incumbent action coefficient C , every branch capable of improving or tying it must obey

$$|\vartheta_j + 2\pi m_j| \leq \sqrt{2C} \quad (j = 1, \dots, n), \quad (118)$$

which gives finite coordinate-wise integer ranges without imposing an arbitrary branch radius. After outward rounding of the real-to-integer range boundaries, the dynamic program uses the cumulative integer branch sum as its state and minimizes over all combinations represented by those ranges without invoking Theorem 2. It agreed with the analytic selector for all 1506 targets; the largest absolute action-coefficient difference was 2.84×10^{-14} . The first 100 targets in each of $n = 2, 3, 4$ were additionally subjected to direct Cartesian enumeration of every trace-compatible vector in the same cost box. All 300 direct searches agreed exactly at the recorded precision, and the largest DP–Cartesian difference was 1.78×10^{-15} ; the largest Cartesian range product was eight and the largest number of trace-compatible candidates was three. These are complete finite floating-point checks under the stated outward-rounded incumbent bounds rather than interval-arithmetic or formal proofs. Across all dimensions, the largest normalized endpoint reconstruction residual was 6.21×10^{-15} , and the largest trace residual was 1.60×10^{-14} in the archived run.

Table 4. Random-target validation and timing. The column $q \neq 0$ reports the fraction of targets whose principal eigenphases require a trace-restoring branch shift. After one unreported warm-up call per dimension, the reported input for each target is the median of five repetitions. Total timing includes Schur decomposition, tolerance-aware branch selection, and endpoint reconstruction; branch timing isolates phase canonicalization, sorting, clustering, and index shifts. Residuals are normalized Frobenius endpoint errors.

n	Samples	$q \neq 0$ (%)	Max. Residual	Median Total (ms)	90th Pct. (ms)	Median Branch (μ s)
2	500	0.00	7.95×10^{-16}	0.072	0.109	44.4
3	400	2.50	1.68×10^{-15}	0.072	0.098	43.5
4	300	4.67	1.71×10^{-15}	0.078	0.107	44.9

Table 4. Cont.

n	Samples	$q \neq 0$ (%)	Max. Residual	Median Total (ms)	90th Pct. (ms)	Median Branch (μ s)
8	160	13.75	1.99×10^{-15}	0.096	0.126	44.5
16	80	13.75	2.36×10^{-15}	0.180	0.214	49.6
32	40	17.50	3.14×10^{-15}	0.549	0.620	55.3
64	16	43.75	4.32×10^{-15}	2.917	3.242	98.3
128	10	40.00	6.21×10^{-15}	114.619	127.274	244.5

9.1.1. Independent Verification Hierarchy

The random-target test was supplemented by 120 random conjugations of analytically prescribed spectra at every displayed dimension: 96 distinct spectra strictly inside the affine alcove, 16 near-cut spectra with margins 2×10^{-6} or 2×10^{-10} , and eight exact-cut spectra. For the interior cases, the maximum absolute cost error, sorted-logarithm-spectrum error, and normalized Frobenius error of the unique logarithm were, respectively, 2.13×10^{-14} , 4.11×10^{-15} , and 6.33×10^{-15} . In the near-cut cases, the cost and spectrum remained accurate to the displayed precision, whereas the matrix-valued logarithm error reached 6.64×10^{-6} at a 2×10^{-10} affine-wall margin. This is the expected eigenprojector sensitivity near a coalescing cut eigenspace rather than a value-function error. At the exact cut, only cost, spectrum, endpoint, and set-valued branch diagnostics are compared because there is no distinguished single logarithm.

A separate selector regression uses $n = 20$, $q = \pm 5$, and phases separated by less than the declared cluster tolerance. The numerical cluster then contains

$$\binom{20}{5} = 15504 \quad (119)$$

finite assignments, which is above the 10^4 explicit-enumeration cap. The implementation nevertheless selects the exact floating-point minimum from the actual phase order independently of the reporting tolerance; a direct enumeration of all 15,504 fixed-cardinality assignments and the cost-bounded dynamic program agree with that representative exactly at the recorded precision. Thus, the cap can limit only the list of reported near-tied alternatives—not the computed optimum. The complete numerical verification hierarchy is summarized in Table 5.

Table 5. Independent numerical verification hierarchy. Cost discrepancies are absolute differences in the action coefficient $C = \frac{1}{2} \sum_j \alpha_j^2$. All values are from the archived deterministic run.

Validation Layer	Cases	Max. Cost Discrepancy	Independent Diagnostic
Haar targets, Bellman cross-check	1506	2.84×10^{-14}	Max. endpoint residual 6.21×10^{-15} ; at most 711 branch-sum states per layer
Low-dimensional Cartesian search	300	0	Max. DP–Cartesian discrepancy 1.78×10^{-15}
Prescribed interior spectra	96	2.13×10^{-14}	Max. sorted-spectrum and logarithm errors $4.11 \times 10^{-15}/6.33 \times 10^{-15}$
Prescribed near/exact-cut spectra	24	0	Correct cost and spectrum; near-cut logarithm sensitivity and exact-cut set-valued output
Large-cluster selector regression	12	0	Maximum cluster count 15,504; exact representative unchanged by enumeration cap

For a dense target, the dominant production operation is the complex Schur decomposition with standard $O(n^3)$ arithmetic and $O(n^2)$ storage. Away from a numerical boundary cluster, the finite selector uses one sort and index shifts, hence $O(n \log n)$ time and $O(n)$ additional storage; order-statistic selection could reduce the generic selection step to expected linear time. At a high-multiplicity boundary, explicitly listing every cluster assignment is necessarily combinatorial, so the implementation caps explicit enumeration and returns the exact representative, the combinatorial cluster count, and a compact cluster/projector description. The Bellman cross-check is algorithmically separate from the production selector. In the archived random study, the largest raw cost-box Cartesian product was about $10^{104.29}$ branches, yet branch-sum compression required at most 711 states per layer and 297,703 transitions for any target. Even at $n = 128$, production branch selection remained negligible compared with the complete dense calculation in the recorded environment. The smallest observed signed affine-wall margin $2\pi - (\max_j \alpha_j - \min_j \alpha_j)$ over all 1506 random targets was $2.34 \times 10^{-2} > 0$, so none of these random samples was numerically on a cut wall. The high-dimensional samples are modest (16 at $n = 64$ and 10 at $n = 128$); the timings are descriptive machine-specific data rather than a fitted asymptotic law.

A separately archived stress test extends the same production Schur/unwrapping pipeline to one Haar-random $SU(256)$ target. The normalized endpoint residual was 1.05×10^{-14} , the trace residual and selector–Bellman cost discrepancy were zero at the recorded precision, and the signed affine-wall margin was 2.85×10^{-2} . The incumbent-derived raw Cartesian branch-box product had base-10 logarithm 246.94, whereas the independent Bellman recursion used at most 2107 states and 2,480,838 transitions. This single case is deliberately reported as a stress target rather than as a statistical scaling sample or a replacement for the multi-target study through $n = 128$.

9.1.2. Random Weighted-Torus Validation

The anisotropic and box-constrained formulas were tested separately from the full-isotropic selector. For each $n \in \{3, 4, 5, 6\}$, the archive contains 60 seeded trace-zero phase targets, which were split equally among random positive quadratic forms with restricted condition numbers 1, 10, and 100; the coordinate bounds μ_j were also randomized. Weighted action was computed by a Cholesky sphere decoder in integer coordinates and independently by an exhaustive Cartesian search over the complete incumbent-certified box implied by

$$(\theta + 2\pi m)^T G(\theta + 2\pi m) \geq \lambda_{\min}(G|_{\mathbb{R}_0^n}) \|\theta + 2\pi m\|_2^2. \quad (120)$$

The box-constrained minimum time was computed by finite Cartesian enumeration and independently by bisection on an integer-interval feasibility predicate: at a trial time T , each m_j belongs to a contiguous integer interval, and a trace-zero branch exists exactly when zero lies between the sums of the interval lower and upper endpoints. The paired algorithms agreed in all 240 cases. Fifty-seven action minima and 20 box-time minima used nonzero lattice branches, so the suite does not merely retest the zero lift. The maximum action discrepancy was zero at the recorded precision, the maximum box-time discrepancy was 1.78×10^{-15} , and the maximum endpoint residual was 8.95×10^{-16} .

The experiment also exposes a common invalid shortcut. If the principal logarithm has trace $2\pi i q$, subtracting its scalar trace component produces a traceless matrix but changes the endpoint by the central phase $e^{-2\pi i q/n}$. The normalized residual is exactly

$$\left| e^{-2\pi i q/n} - 1 \right|, \quad (121)$$

which was $\sqrt{3}$ for every nonzero- q sample at $n = 3$ and $\sqrt{2}$ at $n = 4$. The finite branch rule restores trace without changing the endpoint. The randomized weighted-torus results are summarized in Table 6.

Table 6. Independent randomized validation of weighted action and box-constrained minimum time. Each dimension contains 20 cases at each restricted metric condition number 1, 10, and 100. “Sphere nodes” counts the median visited nodes of the exact Cholesky decoder; the final column is the largest incumbent-certified Cartesian product for the action check.

n	Cases	Condition Numbers	Max. Action Diff.	Max. Box-Time Diff.	Median Sphere Nodes	Max. Action Product
3	60	1, 10, 100	0	1.33×10^{-15}	3	144
4	60	1, 10, 100	0	1.78×10^{-15}	4	125
5	60	1, 10, 100	0	1.78×10^{-15}	5	1080
6	60	1, 10, 100	0	1.33×10^{-15}	7	16,807

9.2. Comparison with Numerical Pulse Optimizers

The drifted-qubit benchmark used $\Delta = 1.1$, $\phi = 1.4$, $T = 1$, and eight equal time slices. The continuously rotating analytic control (98) was midpoint sampled to quantify discretization alone. Its action is exactly J^* at every sampling density because the interaction-picture norm is constant; the endpoint error decreased from 1.60×10^{-5} at four slices to 2.42×10^{-10} at 64 slices. Each numerical method was run from ten independently perturbed initial pulses. Table 7 reports medians and the fraction reaching $\mathcal{E} \leq 10^{-6}$.

Table 7. Drifted-qubit optimizer benchmark. J/J^* is the action relative to the exact continuous-time optimum. Timings are medians over ten seeded restarts except for the sampled analytic control. The randomized and drift-adapted CRAB rows use the same nine-parameter basis size but different frequencies.

Method	Runs	Success (%)	Median $\mathcal{E}$	Median J/J^*	Median Time (s)
Sampled analytic control	1	100	9.93×10^{-7}	1.0000	0
GRAPE-L-BFGS	10	100	2.71×10^{-10}	1.00108	0.120
First-order Krotov-type	10	100	2.67×10^{-10}	1.02769	10.890
CRAB-Powell, randomized mode	10	100	5.87×10^{-8}	2.71233	0.355
CRAB-Powell, drift-adapted mode	10	100	3.20×10^{-10}	1.00912	0.350

In this configuration, the concurrent-gradient method recovered the analytic action to about 1.1×10^{-3} relative accuracy. The first-order Krotov-type method reached the endpoint in every restart, had no rejected monotonicity sweep, and retained 2.77% median excess action under its fixed sweep/proximal budget. The deliberately compressed randomized one-mode CRAB parametrization used nine coefficients rather than the 24 independent slice coefficients and reached the endpoint in every restart, but its median action was $2.71J^*$ (90th percentile $4.20J^*$). Its randomized frequency lies near the first Fourier harmonic and excludes the known laboratory-frame rotation at $\Delta = 1.1$. Keeping nine coefficients and replacing that frequency by Δ reduces the median action to $1.00912J^*$. Thus, the large randomized-row deficit measures this specific basis mismatch, while the adapted row demonstrates that one-mode CRAB can approach the exact action when supplied with the correct physical frequency.

The qutrit benchmark isolates global branch selection. “Branch-blind” initial pulses were centered on the nonprincipal lift (100); “branch-aware” pulses were centered on the exact lift selected by the lattice calculation. Every algorithm reached high endpoint fidelity in both cases. The branch-blind runs remained near the local action ratio $306/26 = 11.769$,

whereas every branch-aware run recovered the globally correct branch; GRAPE–L-BFGS and CRAB–Powell approached J^* to the displayed tolerance, while the Krotov-type method retained 0.233% median excess action under its fixed budget. The detailed results are reported in Table 8.

Table 8. Outrit branch-selection benchmark over ten seeded perturbations. All success rates at $\mathcal{E} \leq 10^{-6}$ were 100%. Ratios slightly below one in branch-aware runs arise from the small nonzero endpoint tolerance in the penalized numerical objective.

Initialization	Method	Median $\mathcal{E}$	Median J/J^*	Median Time (s)
Branch-blind	GRAPE–L-BFGS	1.98×10^{-9}	11.76883	0.002
Branch-blind	First-order Krotov-type	2.06×10^{-9}	11.77017	3.111
Branch-blind	CRAB–Powell, one mode	1.98×10^{-9}	11.76889	0.037
Branch-aware	GRAPE–L-BFGS	1.39×10^{-10}	0.999972	0.002
Branch-aware	First-order Krotov-type	1.54×10^{-10}	1.00233	2.758
Branch-aware	CRAB–Powell, one mode	1.45×10^{-10}	0.999996	0.039

This is the central comparative advantage of the exact lattice solution: it is not merely another local optimizer. It supplies the missing discrete branch decision, after which standard numerical algorithms can refine a waveform within the correct global sector. Conversely, the numerical methods can incorporate bandwidth, leakage, robustness, and sparse controls that are absent from the exact model.

9.3. Cut-Locus Stress Tests

For $SU(3)$, consider the principal-phase family

$$\vartheta(\varepsilon) = \left(\frac{2\pi}{3} + \varepsilon, \frac{2\pi}{3} - \varepsilon, \frac{2\pi}{3} \right), \quad q = 1. \quad (122)$$

At $\varepsilon = 0$, the target is the central element $e^{2\pi i/3}I$. In a fixed diagonal basis, there are three index assignments, but the full- $SU(3)$ minimum logarithms form the continuum

$$X_P = i \left(\frac{2\pi}{3} I - 2\pi P \right), \quad P^2 = P = P^\dagger, \quad \text{Tr } P = 1. \quad (123)$$

For $\varepsilon \neq 0$, the nearest single-exchange action differs by exactly $2\pi|\varepsilon|$ at $T = 1$, which is in agreement with Proposition 3. The numerical test constructs a Fourier-conjugated target matrix, invokes the same Schur extraction and tolerance-aware unwrapping routine used for the random targets, and computes the observed branch costs from the extracted phases. It tracks the law down to $|\varepsilon| = 10^{-12}$ with maximum relative deviation below 8.2×10^{-4} ; at $\varepsilon = 0$, the code reports three coordinate representatives and flags the continuum projector freedom. Figure 1 shows the result.

For $SU(2)$, approach the cut point $-I$ from two directions,

$$U_z(\varepsilon) = e^{-i(\pi-\varepsilon)\sigma_z}, \quad U_x(\varepsilon) = e^{-i(\pi-\varepsilon)\sigma_x}. \quad (124)$$

The targets coalesce as $\varepsilon \rightarrow 0$, but their unique minimizing logarithms for $\varepsilon > 0$ approach different elements of the continuum at $-I$. Direct calculation gives

$$\frac{\|X_z(\varepsilon) - X_x(\varepsilon)\|_F}{\|U_z(\varepsilon) - U_x(\varepsilon)\|_F} = \frac{\pi - \varepsilon}{\sin \varepsilon} \sim \frac{\pi}{\varepsilon}. \quad (125)$$

Thus, the optimal value remains stable, but any single-valued selector of the optimal logarithm is ill-conditioned. The test constructs both matrices, passes each through the implemented Schur/minimum-logarithm routine, and forms the displayed norm ratio from the returned logarithms. It reproduces the inverse-distance divergence down to $\varepsilon = 10^{-12}$ with maximum relative deviation below 4.3×10^{-4} . The directional conditioning is shown in Figure 2.




Figure 1. Schur-selector and analytic action gaps between the best and second-best coordinate branches for the SU(3) family (122). The linear collapse explains why finite-precision branch identification becomes tolerance-dependent at the cut locus.




Figure 2. Directional conditioning of a single-valued minimum-logarithm selector near the SU(2) cut point $-I$. The value function is continuous, whereas the selected generator changes by an $O(1)$ amount under an $O(\varepsilon)$ change of target.

These experiments motivate three implementation rules: retain a branch-gap diagnostic; cluster phases within a declared tolerance and return all finite assignments in that numerical cluster when manageable—otherwise the exact floating-point representative, the combinatorial cluster count, and a compact cluster description; and distinguish those coordinate representatives from the Grassmannian projector freedom caused by exact spectral degeneracy. The optimal-value conditioning must likewise be distinguished from the conditioning of a selected logarithm. A black-box principal logarithm alone cannot make these distinctions.

10. Parameter Uncertainty, Decoherence, and Benchmark Use

The exact results describe a closed, finite-dimensional, perfectly calibrated system. They remain informative under model error, but the sense in which they remain valid must be stated precisely.

10.1. Hamiltonian Parameter Errors

Let U_0 solve $\dot{U}_0 = \Omega U_0$ and let U_δ solve $\dot{U}_\delta = (\Omega + \Delta\Omega)U_\delta$ with both generators skew-Hermitian. The variation of constants in the interaction frame gives the nonasymptotic bound

$$\|U_\delta(T) - U_0(T)\|_2 \leq \int_0^T \|\Delta\Omega(t)\|_2 dt. \quad (126)$$

The first-order error is governed by the toggling-frame integral

$$U_0(T)^\dagger U_\delta(T) - I = \int_0^T U_0(t)^\dagger \Delta\Omega(t) U_0(t) dt + O(\|\Delta\Omega\|_{L^1}^2). \quad (127)$$

Consequently, the exact pulse gives a nominal optimum and a computable sensitivity baseline, while robust pulse design seeks the cancellation of (127) over an uncertainty ensemble. Sequential convex programming, ensemble GRAPE, and related approaches can then report the additional action required relative to the exact nominal certificate [18,34].

10.2. Markovian Decoherence

For an open system with a Gorini–Kossakowski–Sudarshan–Lindblad generator, write the superoperator as an ideal Hamiltonian part plus a dissipator $\mathcal{D}_t$ [35,36]. Let $\mathcal{U}_{t:0}$ be the ideal unitary channel and use a submultiplicative superoperator norm invariant under unitary pre- and postcomposition, such as the diamond norm. If the interaction-picture dissipator obeys

$$\|\mathcal{U}_{t:0}^{-1} \circ \mathcal{D}_t \circ \mathcal{U}_{t:0}\| \leq \eta(t), \quad (128)$$

then a variation of constants and Grönwall's inequality yield

$$\|\mathcal{E}_T - \mathcal{U}_T\| \leq \exp\left(\int_0^T \eta(t) dt\right) - 1, \quad (129)$$

where $\mathcal{U}_T$ is the ideal unitary channel. The closed-system pulse is therefore no longer certified by the closed-system theorem and generally need not remain globally optimal for a fidelity-plus-decoherence objective, although it may remain optimal in special noise models. For a fixed driftless endpoint (or a fixed interaction-picture endpoint), increasing T lowers the ideal quadratic action as $1/T$ but generally increases the decoherence exposure in (129). With a laboratory-frame drift, the interaction-picture endpoint itself depends on T , so even the ideal action need not follow a pure $1/T$ law. In either case, the combined noisy optimum is a new problem.

The exact formulas nevertheless remain useful in four ways. First, the norm-limited time is an ideal-model speed reference; it is a lower bound for a more constrained closed model only under the same endpoint and norm-bound objective. Second, the analytic pulse is a branch-aware initialization for open-system or robust optimization. Third, controlled ablations can compare the exact closed benchmark, the same discretized closed plant, and the enlarged noisy plant; this helps diagnose numerical and model-induced contributions, although a gap relative to the ideal value alone cannot identify them without independent constrained-solver convergence checks or stronger bounds. Fourth, the branch and affine-wall margins quantify when small calibration errors can change the selected logarithm sector even if the optimal value itself remains continuous. Leakage, non-Markovian

noise, waveform filtering, and calibration drift require correspondingly enlarged models and are not covered by the theorems in this paper.

A device-facing use protocol therefore includes the following: identify a calibrated effective block or commuting sector; infer the quadratic weights and amplitude bounds from the calibration model; compute the exact branch, ideal-model action/time reference, and branch margin; use the corresponding analytic pulse as the seed or reference for an ensemble, leakage-aware, or open-system optimizer; and report the resulting action or time relative to that nominal value. For a more constrained closed plant with the same objective, this reference is also a lower bound; for an open or otherwise changed objective, it is not. This protocol does not transfer the closed-system optimality proof to the enlarged plant, but it makes the origin of any performance gap auditable.

11. Discussion

The analytic and numerical results separate four distinct sources of difficulty that are often conflated in quantum gate synthesis.

First, full isotropic actuation is a bi-invariant Riemannian problem. The global result is not a consequence of the local Euler–Arnold equation alone; it also uses compactness, the exact geodesic-distance formula, and the lattice of logarithm branches. For a dense target, spectrum extraction is the $O(n^3)$ bottleneck, whereas the implemented finite unwrapping step is an $O(n \log n)$ sort followed by index shifts. The separate timing data confirm that branch selection is negligible compared with dense matrix decomposition.

Second, drift adaptation is exact but conditional. The interaction picture always removes the drift from the state equation, but it preserves the cost only under the drift-adjoint invariance condition (53). The isotropic Hilbert–Schmidt metric has this invariance. A generic anisotropic metric becomes time dependent in the interaction picture, so the constant-generator global certificate is lost; special symmetric or commuting instances may still admit a constant optimum.

Third, even an Abelian control sector has a nonsmooth global landscape. The torus value function is continuous and piecewise quadratic, but its gradient and minimizing branch change across weighted Voronoi bisectors.

The qutrit experiment shows that endpoint fidelity alone does not reveal a branch error: all three numerical optimizers reached the target while retaining an action nearly twelve times too large when initialized in the wrong lift. The exact lattice calculation supplies the discrete global information that a local waveform optimizer lacks.

The exact reductions have different high-dimensional algorithmic implications. The generic unique isotropic full-control branch is selected by the sorting rule of Theorem 2. By contrast, a general anisotropic torus or embedded-subtorus problem is a closest-vector problem in its period lattice; exact enumeration, sphere decoding, or branch-and-bound may grow rapidly with dimension and conditioning. Explicit output of every assignment in a high-multiplicity numerical boundary cluster can also be combinatorial, where a compact projector description is preferable. The high-dimensional experiment therefore validates the generic full-isotropic solver rather than a dimension-independent polynomial-time algorithm for every weighted lattice instance.

Fourth, the optimal value and an optimal control need not have the same numerical conditioning. At a full- $SU(n)$ cut wall, the value remains continuous while the minimizing logarithms form a Grassmannian continuum; only after fixing a diagonal torus or eigenbasis does one obtain finitely many coordinate branch assignments. The branch margin (41), projector family (123), and directional experiment (125) make this distinction quantitative. Numerical code should therefore report both the finite boundary cluster and

any degenerate-eigenspace projector freedom rather than manufacture a spurious unique answer from an arbitrary eigenbasis.

The practical scope is correspondingly precise. For an exactly isolated one-parameter or commuting native interaction, the formulas can be applied directly. For two-qubit Cartan synthesis, they solve the ideal nonlocal torus once local equivalence and weights have been specified. For a sparse or bandwidth-limited closed plant with the same endpoint and objective, the full-control result is a lower bound because its admissible set is larger. For open, leakage-prone, or otherwise changed objectives, the formulas are nominal baselines, seeds, and regression tests rather than certified bounds. They cannot certify optimality for a plant outside the assumptions.

This paper, which is a numerical study, has limitations of its own. Dense Haar-random matrices probe algebraic correctness and scaling but do not reproduce a specific laboratory Hamiltonian. GRAPE–L-BFGS, a first-order sequential Krotov-type method, and CRAB–Powell are compact reference implementations with different update units and fixed method-specific budgets, so runtime and evaluation counts are not directly interchangeable and the ordering is configuration dependent. The randomized one-mode CRAB frequency is near the first Fourier harmonic and does not contain the known drift rotation at $\Delta = 1.1$; the added drift-adapted CRAB row quantifies that basis mismatch directly. More extensive package-level and mode-count comparisons, sparse many-body models, experimental transfer functions, and open-system benchmarks remain natural next steps.

12. Conclusions

This paper provides a family of exactly solved quantum-control benchmarks on $SU(n)$ under explicit hypotheses. The smooth stationary curves of a quadratic full-control problem satisfy the right-trivialized Euler–Arnold equation. In the fully actuated isotropic case, the global fixed-time optimum is generated by a minimum-norm skew-Hermitian logarithm and is computable by a finite principal-phase unwrapping theorem. At the affine cut wall of spectral width 2π , the exponential is conjugate, the minimizers form Grassmannian families, and a single-valued selector is ill-conditioned even though the value remains continuous. Drift is absorbed into a moving endpoint whenever the metric is drift-adjoint invariant. On maximal tori and accessible closed commuting subtori, weighted action and box-constrained time reduce to exact period-lattice problems with separate strict and projective lattices when the global phase is irrelevant.

The computational study tests these claims on 1506 random targets with a complete cost-bounded Bellman cross-check for every target, 300 direct low-dimensional Cartesian searches, 120 analytically prescribed random-conjugation cases, 240 independently cross-checked anisotropic torus cases, a separate $SU(256)$ stress target, large-cluster selector regressions, ten-restart optimizer comparisons, and solver-level perturbations down to 10^{-12} from a cut locus. It also clarifies the division of labor between analytic and numerical control: the exact solution supplies a global branch decision within its model, an ideal-model reference value, a margin diagnostic, and a physically interpretable seed; numerical optimization supplies device-specific constraints and robustness. That value is a rigorous lower bound only for a more constrained closed problem with the same endpoint and objective. Controlled ablations and independent convergence checks are still required to attribute a performance gap uniquely to topology, discretization, model error, or numerical suboptimality.

Funding: This research received no external funding.

Data Availability Statement: No empirical data are generated or analyzed in this theoretical study.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Variation Identity

For completeness, we derive (19). Let $U_\varepsilon(t)$ be a smooth variation and define

$$\eta(t) = \left. \frac{\partial U_\varepsilon(t)}{\partial \varepsilon} \right|_{\varepsilon=0} U(t)^{-1}. \quad (\text{A1})$$

Then, $\delta U = \eta U$. Since $\Omega = \dot{U}U^{-1}$,

$$\delta \Omega = \delta(\dot{U}U^{-1}) \quad (\text{A2})$$

$$= \delta \dot{U} U^{-1} + \dot{U} \delta(U^{-1}) \quad (\text{A3})$$

$$= (\dot{\eta} U + \eta \dot{U}) U^{-1} - \dot{U} U^{-1} \eta \quad (\text{A4})$$

$$= \dot{\eta} + \eta \Omega - \Omega \eta = \dot{\eta} + [\eta, \Omega]. \quad (\text{A5})$$

Appendix B. Finite Certificate for the Qutrit Lattice Minimum

Write $m = (a, b, -a - b) \in L_3$. For the qutrit lift (100), define dimensionless variables

$$y(a, b) = \frac{1}{\pi}(\theta + 2\pi m) = \left(\frac{7}{5} + 2a, \frac{1}{5} + 2b, -\frac{8}{5} - 2a - 2b \right). \quad (\text{A6})$$

The weighted coefficient divided by π^2 is

$$F(a, b) = \left(\frac{7}{5} + 2a \right)^2 + \left(\frac{1}{5} + 2b \right)^2 + 4 \left(-\frac{8}{5} - 2a - 2b \right)^2. \quad (\text{A7})$$

The continuous minimizer of this strictly convex quadratic is

$$(a_0, b_0) = \left(-\frac{7}{10}, -\frac{1}{10} \right), \quad (\text{A8})$$

and $F(a_0, b_0) = 0$. The Hessian is

$$H = \begin{pmatrix} 40 & 32 \\ 32 & 40 \end{pmatrix}, \quad (\text{A9})$$

with eigenvalues 8 and 72. Hence,

$$F(a, b) = \frac{1}{2} \begin{pmatrix} a - a_0 & b - b_0 \end{pmatrix} H \begin{pmatrix} a - a_0 \\ b - b_0 \end{pmatrix}. \quad (\text{A10})$$

Since $\frac{1}{2} \lambda_{\min}(H) \|(a, b) - (a_0, b_0)\|_2^2 = 4 \|(a, b) - (a_0, b_0)\|_2^2$, any integer point with Euclidean distance greater than $\sqrt{26}/10$ from (a_0, b_0) has $F > 26/25$. The only integer point within this radius is $(-1, 0)$, for which

$$F(-1, 0) = \frac{26}{25}. \quad (\text{A11})$$

All other integer points have larger weighted action. Thus, $m = (-1, 0, 1)$ is the unique weighted-action minimizer.

For the box-time problem, the dimensionless objective is

$$B(a, b) = \max \left\{ \left| \frac{7}{5} + 2a \right|, \left| \frac{1}{5} + 2b \right|, \left| 2 - \frac{8}{5} - 2a - 2b \right| \right\}. \quad (\text{A12})$$

At $(-1, 0)$, $B(-1, 0) = 4/5$. If $B(a, b) < 4/5$, then

$$\left| \frac{7}{5} + 2a \right| < \frac{4}{5} \Rightarrow a = -1, \quad (\text{A13})$$

and

$$\left| \frac{1}{5} + 2b \right| < \frac{4}{5} \Rightarrow b = 0. \quad (\text{A14})$$

But at the only possible pair $(-1, 0)$, the third weighted coordinate equals $4/5$; hence, it cannot be strictly less than $4/5$. Therefore, no branch has box objective below $4/5$, and $m = (-1, 0, 1)$ is a global box-time minimizer.

Appendix C. Numerical Safeguards

The exact solver used the following steps for every target $U \in \text{SU}(n)$.

- (1) Compute a complex Schur decomposition $U = VTV^\dagger$. For a normal unitary matrix, T is diagonal up to roundoff.
- (2) Set $\vartheta_j = \text{Arg}(T_{jj})$, map numerical values at $-\pi$ to the theorem's convention $+\pi$, and set $q = \text{round}(\sum_j \vartheta_j / (2\pi))$ after checking its distance from the nearest integer.
- (3) Apply Theorem 2. Select the representative from the actual stable ordering of the floating-point phases. Use the declared phase-cluster tolerance 10^{-10} only to describe numerical ambiguity at the selected-unselected boundary. Enumerate every boundary subset when the cluster count is at most 10^4 ; otherwise, retain the exact floating-point representative together with the combinatorial cluster count and a compact description. A machine-level degeneracy flag indicates when continuous Grassmannian projector freedom may be present; a floating-point flag is not itself a proof of exact spectral equality.
- (4) Form $X = Vi \text{diag}(\alpha) V^\dagger$ and verify

$$\frac{\|e^X - U\|_F}{\sqrt{n}}, \quad |\text{Tr } X|, \quad \min_{\text{candidate branches}} \frac{\|X\|_F^2}{2T}. \quad (\text{A15})$$

The supplied implementation reconstructs e^X spectrally rather than calling a second dense exponential. For every random and prescribed-spectrum target, the analytically obtained incumbent cost C is used only as an upper bound: every improving or tying branch must satisfy (118). The code derives the corresponding integer ranges coordinate by coordinate, uses `nextafter` to round each real boundary outward before converting it to integers, and applies a Bellman recursion whose state is the cumulative branch sum. This independently minimizes over all trace-compatible branches represented by the cost box without using the finite-shift conclusion being tested. For 300 low-dimensional targets, a second routine directly enumerates the complete Cartesian product and agrees with both the selector and the Bellman result. These are exhaustive finite floating-point checks under the stated bound.

The optimizer scripts use real coefficients in skew-Hermitian bases and the induced coefficient metric. Exact matrix exponentials are used on every piecewise-constant interval, and the same phase-sensitive endpoint error (117) is used for every method. The first-order sequential Krotov-type routine uses terminal costate $P_N = -U_f/n$, old-field backward propagation, new-field forward propagation, and exact `expm_frechet` derivatives in each slice minimization. The term “exact discrete” refers only to the propagator and derivative of the chosen finite-dimensional piecewise-constant model; each local minimization is numerical with the fixed tolerance and iteration budget recorded in the configuration file.

For completeness, the strong-convexity constants quoted for the displayed local sub-problems follow directly from an operator-norm Hessian bound. Write the endpoint part of one local objective in a direction v as

$$h(v) = \Re \operatorname{Tr} \left(P_{k+1}^\dagger e^{\Delta t \Omega(v)} U_k \right). \quad (\text{A16})$$

The terminal condition and unitary backward propagation give $\|P_{k+1}\|_* = 1$, while $\|U_k\|_2 = 1$. The integral formula for the second Fréchet derivative of the exponential therefore yields

$$D^2 h(v)[w, w] \geq -\Delta t^2 \|G(w)\|_2^2, \quad (\text{A17})$$

where $G(w)$ is the skew-Hermitian control generator in direction w . For the Hilbert–Schmidt-orthonormal qubit basis, $\|G(w)\|_2^2 = \|w\|_2^2/2$ and $\Delta t = 1/8$. The unit-weight proximal term has Hessian $\Delta t I$, so the local Hessian is bounded below by $(1/8 - 1/128)I = (15/128)I$. For the qutrit coordinates, $G(w) = i \operatorname{diag}(w_1, w_2, -w_1 - w_2)$, hence $\|G(w)\|_2^2 \leq 2\|w\|_2^2$; moreover, $\Delta t = 1/5$ and the induced metric $M = \begin{pmatrix} 5 & 4 \\ 4 & 5 \end{pmatrix}$ has $\lambda_{\min}(M) = 1$. Thus, the lower bound is $(1/5 - 2/25)I = (3/25)I$. The physical fluence term is positive semidefinite and can only increase these constants. Finite-difference checks of the local Fréchet gradients have maximum absolute discrepancies below 5.3×10^{-11} , and every recorded Krotov-type sweep is checked for the monotonicity of the physical fluence-regularized objective.

The concurrent GRAPE gradient propagates exact forward and suffix products and inserts `expm_frechet` derivatives at every full-system slice; the commuting qutrit uses the corresponding analytic phase gradient. Seeded central-difference checks give maximum absolute discrepancies 3.94×10^{-11} for the full drifted-qubit problem and 2.81×10^{-11} for the qutrit problem with objective-consistency errors at machine precision.

References

1. Jurdjevic, V.; Sussmann, H.J. Control systems on Lie groups. *J. Differ. Equ.* **1972**, *12*, 313–329. [https://doi.org/10.1016/0022-0396\(72\)90035-6](https://doi.org/10.1016/0022-0396(72)90035-6).
2. Jurdjevic, V. *Geometric Control Theory*; Cambridge Studies in Advanced Mathematics; Cambridge University Press: Cambridge, UK, 1996; Volume 52. <https://doi.org/10.1017/CBO9780511530036>.
3. Agrachev, A.A.; Sachkov, Y.L. *Control Theory from the Geometric Viewpoint*; Springer: Berlin, Heidelberg, Germany, 2004. <https://doi.org/10.1007/978-3-662-06404-7>.
4. D'Alessandro, D. *Introduction to Quantum Control and Dynamics*; Chapman & Hall/CRC: Boca Raton, FL, USA, 2007. <https://doi.org/10.1201/9781584888833>.
5. Hall, B.C. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, 2nd ed.; Springer: Cham, Switzerland, 2015. <https://doi.org/10.1007/978-3-319-13467-3>.
6. Khaneja, N.; Brockett, R.; Glaser, S.J. Time optimal control in spin systems. *Phys. Rev. A* **2001**, *63*, 032308. <https://doi.org/10.1103/PhysRevA.63.032308>.
7. Boscain, U.; Mason, P. Time minimal trajectories for a spin 1/2 particle in a magnetic field. *J. Math. Phys.* **2006**, *47*, 062101. <https://doi.org/10.1063/1.2203236>.
8. Carlini, A.; Hosoya, A.; Koike, T.; Okudaira, Y. Time-optimal quantum evolution. *Phys. Rev. Lett.* **2006**, *96*, 060503. <https://doi.org/10.1103/PhysRevLett.96.060503>.
9. Carlini, A.; Hosoya, A.; Koike, T.; Okudaira, Y. Time-optimal unitary operations. *Phys. Rev. A* **2007**, *75*, 042308. <https://doi.org/10.1103/PhysRevA.75.042308>.
10. Nielsen, M.A.; Dowling, M.R.; Gu, M.; Doherty, A.C. Quantum computation as geometry. *Science* **2006**, *311*, 1133–1135. <https://doi.org/10.1126/science.1121541>.
11. Dowling, M.R.; Nielsen, M.A. The geometry of quantum computation. *Quantum Inf. Comput.* **2008**, *8*, 861–899. <https://doi.org/10.26421/QIC8.10-1>.
12. Peirce, A.P.; Dahleh, M.A.; Rabitz, H. Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications. *Phys. Rev. A* **1988**, *37*, 4950–4964. <https://doi.org/10.1103/PhysRevA.37.4950>.

13. Palao, J.P.; Kosloff, R. Optimal control theory for unitary transformations. *Phys. Rev. A* **2003**, *68*, 062308. <https://doi.org/10.1103/PhysRevA.68.062308>.
14. Khaneja, N.; Reiss, T.; Kehlet, C.; Schulte-Herbrüggen, T.; Glaser, S.J. Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms. *J. Magn. Reson.* **2005**, *172*, 296–305. <https://doi.org/10.1016/j.jmr.2004.11.004>.
15. Werschnik, J.; Gross, E.K.U. Quantum optimal control theory. *J. Phys. B At. Mol. Opt. Phys.* **2007**, *40*, R175–R211. <https://doi.org/10.1088/0953-4075/40/18/R01>.
16. Brif, C.; Chakrabarti, R.; Rabitz, H. Control of quantum phenomena: Past, present and future. *New J. Phys.* **2010**, *12*, 075008. <https://doi.org/10.1088/1367-2630/12/7/075008>.
17. Goerz, M.H.; Basilewitsch, D.; Gago-Encinas, F.; Krauss, M.G.; Horn, K.P.; Reich, D.M.; Koch, C.P. Krotov: A Python implementation of Krotov's method for quantum optimal control. *SciPost Phys.* **2019**, *7*, 080. <https://doi.org/10.21468/SciPostPhys.7.6.080>.
18. Koch, C.P.; Boscain, U.; Calarco, T.; Dirr, G.; Filipp, S.; Glaser, S.J.; Kosloff, R.; Montangero, S.; Schulte-Herbrüggen, T.; Sugny, D.; Wilhelm, F.K. Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe. *EPJ Quantum Technol.* **2022**, *9*, 19. <https://doi.org/10.1140/epjqt/s40507-022-00138-x>.
19. Caneva, T.; Calarco, T.; Montangero, S. Chopped random-basis quantum optimization. *Phys. Rev. A* **2011**, *84*, 022326. <https://doi.org/10.1103/PhysRevA.84.022326>.
20. Machnes, S.; Sander, U.; Glaser, S.J.; de Fouquières, P.; Gruslys, A.; Schirmer, S.; Schulte-Herbrüggen, T. Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework. *Phys. Rev. A* **2011**, *84*, 022305. <https://doi.org/10.1103/PhysRevA.84.022305>.
21. Sørensen, A.; Mølmer, K. Quantum computation with ions in thermal motion. *Phys. Rev. Lett.* **1999**, *82*, 1971–1974. <https://doi.org/10.1103/PhysRevLett.82.1971>.
22. Bruzewicz, C.D.; Chiaverini, J.; McConnell, R.; Sage, J.M. Trapped-ion quantum computing: Progress and challenges. *Appl. Phys. Rev.* **2019**, *6*, 021314. <https://doi.org/10.1063/1.5088164>.
23. Krantz, P.; Kjaergaard, M.; Yan, F.; Orlando, T.P.; Gustavsson, S.; Oliver, W.D. A quantum engineer's guide to superconducting qubits. *Appl. Phys. Rev.* **2019**, *6*, 021318. <https://doi.org/10.1063/1.5089550>.
24. Chow, J.M.; Córcoles, A.D.; Gambetta, J.M.; Rigetti, C.; Johnson, B.R.; Smolin, J.A.; Rozen, J.R.; Keefe, G.A.; Rothwell, M.B.; Ketchen, M.B.; Steffen, M. Simple all-microwave entangling gate for fixed-frequency superconducting qubits. *Phys. Rev. Lett.* **2011**, *107*, 080502. <https://doi.org/10.1103/PhysRevLett.107.080502>.
25. Foxen, B.; Neill, C.; Dunsworth, A.; Roushan, P.; Chiaro, B.; Megrant, A.; Kelly, J.; Chen, Z.; Satzinger, K.; Barends, R.; et al. Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms. *Phys. Rev. Lett.* **2020**, *125*, 120504. <https://doi.org/10.1103/PhysRevLett.125.120504>.
26. Zhang, J.; Vala, J.; Sastry, S.; Whaley, K.B. Geometric theory of nonlocal two-qubit operations. *Phys. Rev. A* **2003**, *67*, 042313. <https://doi.org/10.1103/PhysRevA.67.042313>.
27. Bhatia, R. *Matrix Analysis*; Springer: New York, NY, USA, 1997. <https://doi.org/10.1007/978-1-4612-0653-8>.
28. Higham, N.J. *Functions of Matrices: Theory and Computation*; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 2008. <https://doi.org/10.1137/1.9780898717778>.
29. Conway, J.H.; Sloane, N.J.A. *Sphere Packings, Lattices and Groups*, 3rd ed.; Springer: New York, NY, USA, 1999. <https://doi.org/10.1007/978-1-4757-6568-7>.
30. Al-Mohy, A.H.; Higham, N.J. Improved inverse scaling and squaring algorithms for the matrix logarithm. *SIAM J. Sci. Comput.* **2012**, *34*, C153–C169. <https://doi.org/10.1137/110852553>.
31. Arnold, V.I. Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. *Ann. Inst. Fourier* **1966**, *16*, 319–361. <https://doi.org/10.5802/aif.233>.
32. Crittenden, R.J. Minimum and conjugate points in symmetric spaces. *Can. J. Math.* **1962**, *14*, 320–328. <https://doi.org/10.4153/CJM-1962-024-8>.
33. Mezzadri, F. How to generate random matrices from the classical compact groups. *Not. Am. Math. Soc.* **2007**, *54*, 592–604. Available online: <https://www.ams.org/notices/200705/fea-mezzadri-web.pdf> (accessed on 11 July 2026).
34. Kosut, R.L.; Grace, M.D.; Brif, C. Robust control of quantum gates via sequential convex programming. *Phys. Rev. A* **2013**, *88*, 052326. <https://doi.org/10.1103/PhysRevA.88.052326>.
35. Lindblad, G. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.* **1976**, *48*, 119–130. <https://doi.org/10.1007/BF01608499>.
36. Gorini, V.; Kossakowski, A.; Sudarshan, E.C.G. Completely positive dynamical semigroups of N -level systems. *J. Math. Phys.* **1976**, *17*, 821–825. <https://doi.org/10.1063/1.522979>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.