Quantum speed limits from symmetries: A numerical analysis

Santana Lujan,[1,2,*] Marco Wiedmann,[2] and Daniel Burgarth[2]

[1] *Institute for Space Operations and Astronaut Training, German Aerospace Center (DLR), 82234 Weßling, Germany*

[2] *Department of Physics, Friedrich-Alexander-Universität Erlangen-Nürnberg, 91058 Erlangen, Germany*

Target-specific quantum speed limits (QSLs) bound the minimal time required to realize a given operation. Recent theoretical work derives these limits directly from the symmetries of the dynamical Lie algebra of the control system and a symmetry-restoring perturbation, bypassing the need to solve the equations of motion. In practice, however, identifying a suitable symmetry and a perturbation that restores it demands strong physical intuition and offers no guarantee of tightness. We show that, for any given symmetry, the minimal restoring perturbation is the solution of a convex semidefinite program, which can be solved to global optimality and thus yields the tightest achievable speed limit for that symmetry. This semidefinite program is one building block of our numerical method, alongside an explicit computation of the space of admissible symmetries and a search over this space. Since this search is non-convex, it is addressed with gradient-based optimization, black-box heuristics, or the injection of physically motivated candidate symmetries. We apply the method to three examples: a two-qubit CNOT gate, a SWAP on a spinless-particle chain, and a Dicke state preparation on a qubit ring. Our results improve on previously known symmetry-based bounds, most notably for the SWAP, where solving the semidefinite program for the known analytical symmetry tightens the bound by a factor of approximately 2.5. For the Dicke state preparation, however, the bounds decrease with system size, which we trace back to a symmetry-independent perturbation ceiling. This work provides a reproducible numerical procedure for constructing target-specific quantum speed limits directly from the system Hamiltonians and the target operation.

arXiv:2611.22630v1 [quant-ph] 17 Nov 2026

I. INTRODUCTION

As quantum processors transition from the noisy intermediate-scale quantum computing (NISQ) era toward fault tolerance, executing operations as rapidly as possible is essential to maximize gate throughput and mitigate decoherence [1]. In quantum control, whenever control fields alone cannot implement a desired operation, the finite interaction strength of the drift Hamiltonian imposes an intrinsic minimal duration, known as the quantum speed limit (QSL) or controllability time [2]. Historically rooted in time-energy uncertainty relations, fundamental speed limits were formalized by Mandelstam–Tamm and Margolus–Levitin bounds [3, 4]. While such QSLs provide operation-independent, worst-case bounds, practical interest typically centers on specific operations, such as multiqubit entangling gates or state preparations, where worst-case limits are often loose.

Target-specific operation times are commonly benchmarked from above via numerical pulse optimization such as GRAPE [5] or CRAB [6], which search for time-optimal control fields, or from the quantum brachistochrone equations for optimal Hamiltonian trajectories [7, 8]. However, these approaches require explicitly simulating the time evolution of the system, scaling unfavorably with the Hilbert space dimension. Lower bounds can be obtained from the geometric frameworks of *distance to uncontrollability* [9] and its target-specific generalization, the *distance to unreachability* [10]. These frameworks establish QSLs directly from the Lie-algebraic generators of the system, without solving the Schrödinger equation. Concretely, given a target operation that is not reachable by the controls alone, a symmetry S of the control-only algebra which is broken only by the drift term is identified. The obtained speed limit depends inversely on the spectral norm of the minimal perturbation that restores the symmetry.

In practice, identifying valid symmetries and minimal restoring perturbations remains non-trivial: while Ref. [9] briefly noted that symmetries can be computed by solving a linear system, previous applications relied on system-specific analytical constructions [10], offering no guarantee on the minimality of the restoring perturbation. Since the symmetry-restoring condition is linear in the perturbation, minimizing the spectral norm of the perturbation is a convex problem, which we cast as a semidefinite program (SDP). This guarantees that, for any candidate symmetry, the computed perturbation is globally minimal, so that the resulting bound can only be improved by choosing a different symmetry. Solving the SDP serves as the inner optimization of our numerical method, which first computes an explicit basis for the space of admissible symmetries via matrix factorizations and then searches this space for the symmetry yielding the tightest bound. Although optimizing over the symmetry space is non-convex, having an explicit basis enables structured automated searches while also accommodating physically motivated candidate symmetries. Altogether, this transforms the calculation of target-specific QSLs into a systematic, reproducible procedure requiring only the system Hamiltonians and the target operation as input.

We demonstrate the method on three physical control systems. First, for a CNOT gate on locally controlled

* santana.lujan@dlr.de

qubits, automated search reproduces and marginally improves upon the known symmetry-based analytical bound [10], validating the numerical implementation. Second, for the SWAP operation across a spinless-particle chain, solving the SDP for the hand-derived symmetry tightens the previous bound by a constant factor. However, automated heuristic searches over the computed symmetry space struggle to locate this narrow optimum, demonstrating the importance of physically motivated symmetry candidates. Finally, for a Dicke state preparation on a globally controlled qubit ring [11], the computed speed limit decreases with system size across all sampled symmetries. We show that the optimization encounters a symmetry-independent perturbation ceiling governed by the drift Hamiltonian, while leaving open whether symmetries that yield tighter bounds exist in the admissible symmetry space.

The remainder of the paper is organized as follows: Sec. II reviews the theoretical framework of symmetry-based speed limits, while Sec. III details the numerical method and its individual stages. Section IV presents the application to the three example systems, followed by a discussion of scaling trends and computational bottlenecks in Sec. V, and concluding remarks in Sec. VI.

II. THEORY

A bilinear quantum control system on a Hilbert space $\mathcal{H}$ of dimension d ($d < \infty$) is defined by a constant, always-on drift term H_d and variable controls H_j, such that the unitary propagator $U(t)$ evolves according to the Schrödinger equation

$$\frac{\mathrm{d}}{\mathrm{d}t}U(t) = -\frac{\mathrm{i}}{\hbar}H(t)U(t),\ U(0) = \mathbb{1}, \tag{1}$$

with the time-dependent Hamiltonian

$$H(t) = H_d + \sum_{j=1}^{r} f_j(t)H_j. \tag{2}$$

We allow for the control pulses $f_j : \mathbb{R} \to \mathbb{R}$ to be chosen from any reasonable set of functions (e.g., piecewise constant). In particular, we impose no *a priori* bounds on the control pulses, meaning that the speed limits we obtain are fundamental and not only arise from experimental constraints.

For any given bilinear quantum control system, one can determine the set of all reachable unitaries

$$\mathcal{R} = \{V \in \mathrm{U}(d) \mid \exists T \geq 0, \exists f_j : \mathbb{R} \to \mathbb{R} \text{ s.t. } U(T) = V\} \tag{3}$$

that can be implemented on that control system. To do that, one computes the so-called dynamical Lie algebra (DLA) of the control system

$$\mathfrak{g} = \langle iH_d, iH_1, iH_2, ...iH_r\rangle_{\mathrm{Lie}}, \tag{4}$$

i.e., the smallest Lie algebra that contains the drift and all controls. The reachable set is then dense in the associated Lie group (Theorem 3.2.1 in Ref. [12])

$$\bar{\mathcal{R}} = \exp(\mathfrak{g}). \tag{5}$$

If the DLA is compact, the closure is not needed and we get exactly $\mathcal{R} = \exp(\mathfrak{g})$.

Since we do not impose bounds on the control functions, every target that can be reached purely by the controls can also be reached arbitrarily fast. Formally, we introduce the control Lie algebra

$$\mathfrak{h} = \langle iH_1, ..., iH_r\rangle_{\mathrm{Lie}} \tag{6}$$

and the set of unitaries that is reachable purely from the controls $\mathcal{R}_{\mathrm{ctrl}} = \exp(\mathfrak{h})$. QSLs only occur in this setting if $\mathfrak{h}$ is strictly smaller than $\mathfrak{g}$ and therefore if it is in particular not universal. Proper subalgebras of compact Lie algebras are characterized by the enlargement of the set of so-called quadratic symmetries [13]. These are Hermitian matrices $S^{(2)} \in \mathbb{C}^{d^2\times d^2}$ that commute with the symmetric embeddings of the drift and control Hamiltonians onto $\mathcal{H} \otimes \mathcal{H}$

$$\forall j = 1, ..., r\text{: } [S^{(2)}, H_j \otimes \mathbb{1} + \mathbb{1} \otimes H_j] = 0. \tag{7}$$

The set of all quadratic symmetries is also called the second-order commutant $\mathfrak{h}^{(2)}$ of $\mathfrak{h}$. This set is higher dimensional than the commutant $\mathfrak{h}^{(1)}$ of $\mathfrak{h}$ (i.e., all Hermitian matrices $S^{(1)} \in \mathbb{C}^{d\times d}$ for which $[S^{(1)}, H_j] = 0$, $j = 1, ..., r$), and every element in the commutant has a corresponding quadratic symmetry obtained by the same symmetric embedding. A Lie algebra $\mathfrak{h}$ is a proper subalgebra of the compact Lie algebra $\mathfrak{g}$ if and only if $\dim \mathfrak{h}^{(2)} > \dim \mathfrak{g}^{(2)}$ [13]. We always have trivially $\dim \mathfrak{g}^{(2)} \geq 2$ as the identity and the permutation of the tensor components are always trivial quadratic symmetries.

For any target operation $U \in \mathcal{R} \setminus \mathcal{R}_{\mathrm{ctrl}}$ that is not reachable with the control Hamiltonians alone, the constant strength of the drift term imposes a QSL, i.e., a minimal time necessary to reach that target. Since by assumption $\mathcal{R}$ is strictly larger than $\mathcal{R}_{\mathrm{ctrl}}$, we infer that there exists a quadratic symmetry $S^{(2)} \in \mathfrak{h}^{(2)}$ that does not commute with $U \otimes U$, i.e., $[U \otimes U, S^{(2)}] \neq 0$. The guaranteed existence of such a quadratic symmetry is leveraged in Ref. [10] to derive a lower bound on the QSL on implementing U

$$T \geq \frac{\|[U^{\otimes 2}, S^{(2)}]\|_F}{4\|S^{(2)}\|_F\|\Delta H\|_\infty}. \tag{8}$$

Here, $\|A\|_F = \sqrt{\mathrm{Tr}(A^\dagger A)}$ is the Frobenius norm, $\|\cdot\|_\infty$ denotes the spectral norm and ΔH is the minimal perturbation to the drift that restores the quadratic symmetry $S^{(2)}$ to it, i.e.,

$$[S^{(2)}, (H_d + \Delta H) \otimes \mathbb{1} + \mathbb{1} \otimes (H_d + \Delta H)] = 0. \tag{9}$$

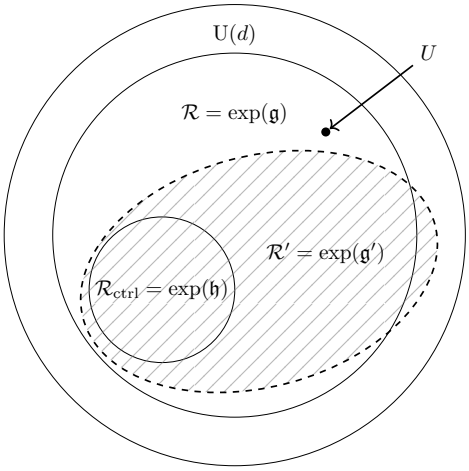

Figure 1. Schematic of the reachable-unitary hierarchy for a controlled system with Hamiltonian $H = H_d + \sum_j f_j(t)H_j$. The system's DLA generates a reachable set $\mathcal{R} \subseteq \mathrm{U}(d)$ containing the target U, while the controls alone reach a subset $\mathcal{R}_{\text{ctrl}} \subset \mathcal{R}$. Perturbing the drift restores a symmetry of the control algebra $\mathfrak{h}$ to the perturbed algebra $\mathfrak{g}'$, which generates $\mathcal{R}'$ (dashed ellipse). This symmetry forces $U \notin \mathcal{R}'$, underlying the speed limit for realizing U of Ref. [10].

In other words, the perturbation ΔH should define a (fictional) perturbed control system

$$\mathfrak{g}' = \langle i(H_d + \Delta H), iH_1, ..., iH_r \rangle_{\text{Lie}}, \tag{10}$$

such that $S^{(2)} \in \mathfrak{g}'^{(2)}$. That the QSL depends inversely on the spectral norm of that perturbation was originally shown in Ref. [9]. For the relation between all the different Lie algebras and their reachable sets that we introduced in this section, see also Fig. 1.

If there already exists a linear symmetry $S^{(1)} \in \mathfrak{h}^{(1)} \setminus \mathfrak{g}^{(1)}$, a tighter bound on the QSL can be derived [10]

$$T \geq \frac{\|[U, S^{(1)}]\|_F}{2\|S^{(1)}\|_F \|\Delta H\|_\infty}, \tag{11}$$

where the perturbation ΔH should now restore the linear symmetry instead, i.e.,

$$[H_d + \Delta H, S^{(1)}] = 0. \tag{12}$$

Since Eqs. (11) and (8) share the same structure, we combine them into a single bound. For a symmetry $S^{(p)}$ of order p, the speed limit is then

$$T \geq \frac{\|[U^{\otimes p}, S^{(p)}]\|_F}{2p\|S^{(p)}\|_F \|\Delta H\|_\infty}. \tag{13}$$

Here, $p = 1$ denotes linear and $p = 2$ quadratic symmetries, and ΔH restores $S^{(p)}$ according to Eq. (12) or Eq. (9), respectively. From here on, we write S without the superscript whenever a statement holds for both orders.

Finally, notice that any components of $S^{(2)}$ or $S^{(1)}$ along $\mathfrak{g}^{(2)}$ or $\mathfrak{g}^{(1)}$ do not contribute to the numerators in Eqs. (8) and (11), but increase the Frobenius norm of the symmetry which appears in the denominator. Therefore, to get the tightest possible bounds, we should always discard these components, i.e., we should choose the symmetry from the quotient space $\mathfrak{h}^{(2)}/\mathfrak{g}^{(2)}$ or $\mathfrak{h}^{(1)}/\mathfrak{g}^{(1)}$ correspondingly.

III. NUMERICAL OPTIMIZATION OF THE SPEED LIMIT

This section showcases the mapping of the algebraic framework established in Sec. II into a systematic numerical procedure, as depicted in Fig. 2. The required input consists solely of a generating set for the Lie algebras $\mathfrak{h}, \mathfrak{g}$, alongside the target unitary operator U. Since the speed limit T represents a lower bound on evolution time, our objective is to maximize T to obtain the tightest bound.

Following Eq. (13), we formulate this maximization as a nested optimization problem: we seek a symmetry operator S from the admissible symmetry space that maximizes the normalized commutator ratio $\|[U^{\otimes p}, S]\|_F / \|S\|_F$, while minimizing the spectral norm of the symmetry-restoring perturbation $\|\Delta H\|_\infty$. The method is implemented in Python 3.13 and is structured into three stages: (1) *Symmetry basis construction*: computing an orthonormal basis for the vector space of symmetries; (2) *Outer optimization*: determining the optimal symmetry S within this space; and (3) *Inner optimization*: solving a convex SDP to find the minimal-norm perturbation ΔH that restores the symmetry S.

A. Calculating a basis for the vector space of symmetries

For a DLA $\mathfrak{m}$, we seek all symmetries $S \in \mathbb{C}^{n\times n}$ with $n = d^p$ that commute with every $G \in \mathfrak{m}$. In the following, $\mathrm{vec}(\cdot)\colon \mathbb{C}^{n\times n} \to \mathbb{C}^{n^2}$ denotes row-wise vectorization, mapping a matrix to a vector by stacking its rows. For quadratic symmetries, we replace each G by the embedding $G \otimes \mathbb{1} + \mathbb{1} \otimes G$. Using Roth's lemma $\mathrm{vec}(ABC) = (A \otimes C^T)\,\mathrm{vec}(B)$, the condition $[G, S] = 0$ is mapped to a linear system

$$\mathrm{vec}([G, S]) = \left(G \otimes \mathbb{1} - \mathbb{1} \otimes G^T\right) \mathrm{vec}(S) = 0\,. \tag{14}$$

By the Jacobi identity, if an operator S commutes with two generators G_1 and G_2, it must also commute with their commutator $[G_1, G_2]$, since $[[G_1, G_2], S] = [[G_1, S], G_2] + [G_1, [G_2, S]] = 0$. Because the DLA is constructed as the Lie closure of a generating set $\{G_1, \ldots, G_g\}$, enforcing Eq. (14) solely on these generators is sufficient to guarantee that S commutes with the entire DLA. Defining the constraint block for each generator as $M_i \coloneqq G_i \otimes \mathbb{1} - \mathbb{1} \otimes G_i^T \in \mathbb{C}^{m\times m}$ with vectorized dimension $m \coloneqq n^2$, we consolidate all individual

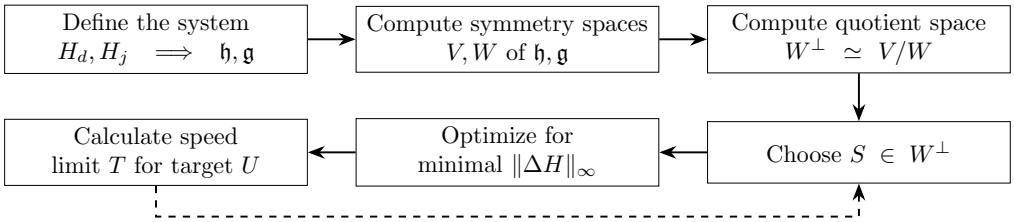

Figure 2. Overview of the numerical procedure for calculating the speed limit T. The drift and control Hamiltonians define the dynamical Lie algebras (DLAs), whose symmetry spaces are computed. We calculate a basis for the quotient space via the orthogonal complement, which serves as a pool of candidate symmetries. After selecting a symmetry S, solving a semidefinite program (SDP) yields the minimal-norm perturbation that restores S. Together with the target unitary U, the speed limit is calculated via the corresponding linear or quadratic bound. The dashed arrow denotes an optional iterative search over S, e.g., using a heuristic or gradient ascent scheme.

symmetry requirements into a single joint linear system:

$$\mathcal{M}\,\mathrm{vec}(S) = 0, \quad \text{where} \quad \mathcal{M} \coloneqq \begin{pmatrix} M_1 \\ M_2 \\ \vdots \\ M_g \end{pmatrix} \in \mathbb{C}^{gm \times m}. \tag{15}$$

Here, g is the number of generators of the DLA, and the matrix $\mathcal{M}$ is assembled by vertically stacking the individual blocks M_i for each generator G_i.

Finding the symmetry space of the DLA is thus equivalent to computing $\ker(\mathcal{M})$. This is a standard null-space problem, which we solve numerically using either the singular value decomposition (SVD) [14, p. 78] or a rank-revealing QR decomposition [14, pp. 276–278]. We discuss the computational complexity of this calculation in Sec. V B. The output of this stage is an explicit orthonormal basis spanning the symmetry space of the DLA.

Next, we isolate symmetries of the control algebra $\mathfrak{h}$ not present in the full system algebra $\mathfrak{g}$. Let $V \coloneqq \mathfrak{h}^{(p)}$ and $W \coloneqq \mathfrak{g}^{(p)}$ denote the vector spaces of (linear or quadratic) symmetries for the algebras $\mathfrak{h}$ and $\mathfrak{g}$, respectively. Since $\mathfrak{h} \subset \mathfrak{g}$ by construction, it follows that $W \subset V$, which allows us to construct the quotient space V/W. To compute this quotient space numerically, we exploit the fact that in finite-dimensional inner product spaces, $V/W \simeq W^\perp$ [15] for a closed vector subspace W in V. The task thus reduces to calculating an orthonormal basis for this orthogonal complement.

To achieve this, let $\mathbf{V}$ and $\mathbf{W}$ be matrices whose columns consist of the vectorized orthonormal basis elements of V and W, respectively. We perform a thin QR decomposition [14, p. 248] without column pivoting, which preserves the block ordering, of the horizontally concatenated matrix

$$\left[\,\mathbf{W}\,\middle|\,\mathbf{V}\,\right] \xrightarrow{\mathrm{QR}} \underbrace{\left[\,\mathbf{W}\,\middle|\,\mathbf{W}^\perp\,\middle|\,*\,\right]}_{Q} R, \tag{16}$$

where $Q \in \mathbb{C}^{m \times (\dim(W)+\dim(V))}$ has orthonormal columns, and R is square and upper triangular. Because the columns of the input matrix are processed sequentially from left to right, the first block of Q reproduces the basis of W, while the second block $\mathbf{W}^\perp$ contains exactly the $k \coloneqq \dim(V) - \dim(W)$ columns that span the orthogonal complement $W^\perp$. The remaining $\dim(W)$ columns originate from linearly dependent input columns and are discarded. Extracting the columns of $\mathbf{W}^\perp$ yields the vectorized basis elements s_j. By inverting the vectorization, $S_j \coloneqq \mathrm{vec}^{-1}(s_j)$, we recover a basis of Hermitian operators S_j, enabling us to parametrize any candidate symmetry operator as

$$S = \sum_{j=1}^{k} c_j S_j \quad \text{with} \quad c_j \in \mathbb{R}. \tag{17}$$

B. Minimal perturbation as a semidefinite program

Let $\mathcal{P}(S)$ denote the set of Hermitian perturbations ΔH that restore S according to Eq. (12) or Eq. (9). The minimal perturbation norm is then

$$f(S) \coloneqq \min_{\Delta H \in \mathcal{P}(S)} \|\Delta H\|_\infty. \tag{18}$$

Analytical constructions, as in Ref. [10], provide some $\Delta H \in \mathcal{P}(S)$, but not necessarily the minimal one. Since both restoring conditions are linear in ΔH, Eq. (18) is an instance of spectral-norm minimization under linear constraints. Such problems are convex and can be cast as an SDP [16, pp. 169–170], so $f(S)$ can be computed to global optimality. Introducing an auxiliary variable λ that bounds $\|\Delta H\|_\infty$ from above, the SDP is posed as:

$$\begin{aligned} \min_{\lambda \in \mathbb{R},\ \Delta H \in \mathcal{P}(S)} \quad & \lambda \\ \text{subject to} \quad & \lambda \mathbb{1} \pm \Delta H \succeq 0, \end{aligned} \tag{19}$$

with $f(S)$ as the optimal value $\lambda^\star$. The linear matrix inequalities (LMIs) constrain the eigenvalues of ΔH to the interval $[-\lambda, \lambda]$, while the membership constraint $\Delta H \in \mathcal{P}(S)$ imposes linear constraints, enforcing the physical symmetry. This formulation also avoids the non-differentiability of $\|\cdot\|_\infty$.

To solve Eq. (19) numerically, we implement the optimization using the CVXPY [17] package. Inside this framework, we declare the perturbation ΔH as a Hermitian matrix and λ as a scalar non-negative optimization variable. To evaluate the inner optimization for different candidate symmetries, we parametrize S according to Eq. (17), where the Hermitian basis operators S_j serve as constant matrices and the coefficients c_j act as parameters within the Disciplined Parametrized Programming (DPP) framework [18]. While held fixed during a single SDP evaluation, these parameters are updated dynamically by the outer optimization loop. This formulation naturally complies with the restrictions of the DPP ruleset without requiring workarounds. Crucially, it allows CVXPY to bypass full problem re-canonicalization at each iteration, reducing the repeated formulation overhead to a small fraction of the initial setup time.

We solve the resulting SDP using the MOSEK [19] solver for high numerical precision. However, we switch to the Splitting Conic Solver (SCS) [20], which supports differentiation through a cone program [18], when extracting gradient information of the minimal spectral norm with respect to the coefficients $\nu_j \coloneqq \partial f(S)/\partial c_j$. Further discussion on the choice of solver and numerical precision limits is provided in Sec. V B.

C. Strategies for finding the best symmetry

The objective of the outer optimization is to maximize the speed limit T by optimizing the coefficients c_j in the parametrized symmetry expansion of Eq. (17) over a fixed basis $\{S_1, \dots, S_k\}$ constructed in Sec. III A. In contrast to the inner optimization where the symmetry is fixed, allowing these coefficients to vary renders the problem non-convex. For a fixed symmetry S, the restoring constraint $\Delta H \in \mathcal{P}(S)$ is strictly linear in the perturbation ΔH. However, treating the symmetry as a variable introduces products of the coefficients c_j and ΔH in the commutator of Eq. (12) or Eq. (9). Consequently, the topography of the optimization landscape is dictated by the structure of symmetry space, which itself is determined by the underlying Lie algebras $\mathfrak{h}$ and $\mathfrak{g}$. Because no universal method exists for finding the global maximum in such non-convex landscapes, we employ several complementary strategies to explore the search space.

We first discuss automated methods that operate independently of any *a priori* insight on the underlying physical system. For systems with low-dimensional symmetry spaces, a grid search over the coefficients c_j remains computationally feasible, scaling exponentially as $O(N_s^k)$ for N_s samples per coefficient dimension across the k-dimensional space. This exhaustive mapping of the optimization landscape offers detailed insights into the distribution and shape of extremal areas. For higher-dimensional symmetry spaces where grid search becomes prohibitive, we employ heuristic and black-box optimization techniques. As a parallelizable baseline, we sample the coefficients randomly using the NumPy [21] package. These random samples serve as starting points for standard local optimizers in the SciPy library [22], as well as the specialized heuristic solver ZoMBI [23], which is tailored particularly for needle-in-a-haystack optimization landscapes. Finally, we implement a gradient ascent scheme, leveraging the ability of CVXPY to return the perturbation derivative ν_j as defined in Sec. III B. Differentiating Eq. (13) with the product rule yields the explicit gradient:

$$\frac{\partial T}{\partial c_j} = T\left\{ \frac{\mathrm{Re}\big[\mathrm{Tr}\big([U^{\otimes p}, S]^\dagger [U^{\otimes p}, S_j]\big)\big]}{\|[U^{\otimes p}, S]\|_F^2} - \frac{\mathrm{Tr}(S_j S)}{\|S\|_F^2} - \frac{\nu_j}{\|\Delta H\|_\infty} \right\}. \tag{20}$$

Only the third term in Eq. (20) depends on the perturbation derivative ν_j, which is returned by the solver and is therefore only reliable above its convergence tolerance. To filter out numerical noise, we set $\nu_j = 0$ whenever the magnitude falls below a threshold $|\nu_j| \leq \eta \coloneqq 10^{-4}$, matching the solver's default convergence tolerance. Following a standard gradient ascent scheme, we update the coefficients at each iteration via $c_j \to c_j + \epsilon \partial T/\partial c_j$ with a suitable step size ϵ. The impact of solver precision and tolerance limits is revisited in Sec. V B 2.

In contrast to the automated strategies, our method also accepts physically motivated symmetry operators. This lets us directly use the analytical symmetries constructed in Ref. [10], allowing physical intuition to guide the search rather than relying solely on automated methods. Given such a candidate symmetry S, we first verify that it commutes with the control algebra $\mathfrak{h}$, and optionally project it onto the quotient space V/W. As discussed in Sec. II, this discards the components along W, which only increase $\|S\|_F$ without contributing to the commutator in Eq. (13). The resulting operator is then fed into the inner convex optimization to find the minimal-norm perturbation ΔH that restores this symmetry. The selection of the optimal search strategy depends on the physical system under study, as illustrated in the following section where we apply the method to three distinct physical architectures.

IV. APPLICATION TO EXAMPLE SYSTEMS

We apply the method of Sec. III to three physical control systems. A pair of locally controlled Ising-coupled qubits serves as validation, a spinless particle on a chain shows it can tighten a known analytical bound, and a ring of globally controlled qubits exposes its limitations.

A. CNOT on a pair of locally controlled qubits

We calculate the target-specific speed limit for implementing a CNOT gate on a two-qubit system with Hilbert

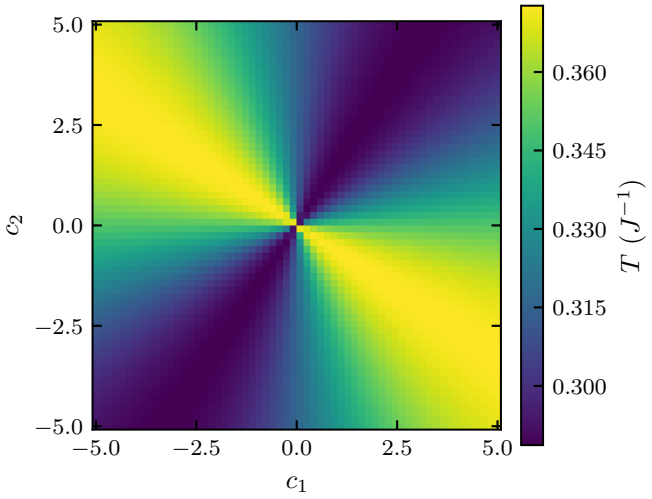

Figure 3. Heatmap of the speed limit T for the CNOT operation between Ising-coupled qubits, obtained from a grid search over the two-dimensional quadratic symmetry $S^{(2)} = c_1 S_1^{(2)} + c_2 S_2^{(2)}$ with $c_1, c_2 \in [-5, 5]$ sampled at 60 points each. The inner loop converges to $\Delta H = -Z_1 Z_2$ with $\|\Delta H\|_\infty = 1$ for every sampled $S^{(2)}$, thus T is maximized purely by the ratio $\|[U^{\otimes 2}, S^{(2)}]\|_F / \|S^{(2)}\|_F$, which peaks along a line through the origin (excluding $S^{(2)} = 0$) at $T_{\max} = 0.373$.

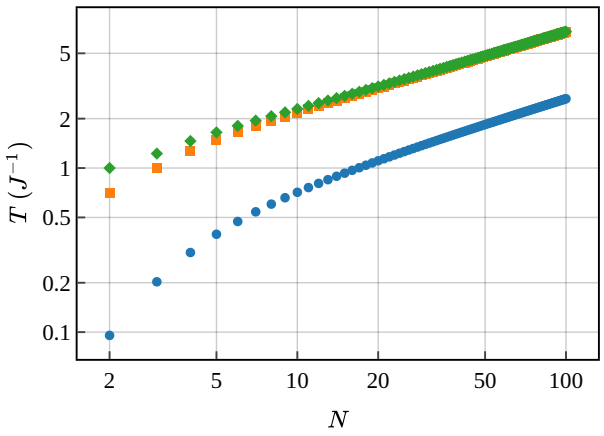

Figure 4. Speed limit T for the SWAP operation between two ends of a spinless-particle chain as a function of the number of sites N (with coupling $J = 1$). The purely analytical bound of Ref. [10] (blue circles, bottom) is compared against two numerically optimized bounds. Injecting the analytical symmetry and optimizing the perturbation via the semidefinite program (SDP) (orange squares, middle) yields an asymptotic tightening by a factor of $T_{\text{num}}/T_{\text{an}} \approx 2.5$. Projecting this symmetry onto the numerically constructed symmetry space prior to the optimization (green diamonds, top) further improves the bound at small N while maintaining the same $O(\sqrt{N})$ scaling.

space $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2$. The Hamiltonian is composed of a longitudinal Ising coupling $H_d = J Z_1 Z_2$ and local controls $H_1 = X_1$, $H_2 = Y_1$, $H_3 = X_2$, $H_4 = Y_2$ on each qubit. For our calculations, we set the coupling strength $J = 1$. Although the controls are purely local and thus incapable of generating entanglement, the system lacks any non-trivial linear symmetry [10]. Explicit computation of the linear symmetry space, as discussed in Sec. III A, confirms this absence: the null space consists solely of the identity operator, rendering the quotient space V/W trivial.

Turning to quadratic symmetries, the computed quotient space yields a two-dimensional basis $\left\{S_1^{(2)}, S_2^{(2)}\right\}$, permitting a thorough, fine-grained grid search over the coefficients c_1, c_2. We feed each candidate symmetry $S^{(2)} = c_1 S_1^{(2)} + c_2 S_2^{(2)}$ directly into our inner optimization SDP to compute the corresponding speed limit, with the results shown in Fig. 3.

For all candidate symmetries, the SDP consistently converges to the perturbation $\Delta H = -Z_1 Z_2$, which fully decouples the qubits and has spectral norm $\|\Delta H\|_\infty = 1$. The speed limit is therefore exclusively dictated by the rate of symmetry breaking on the doubled Hilbert space. Since the SDP always converges to the decoupling perturbation, Eq. (8) simplifies to

$$T \geq \frac{\|[U^{\otimes 2}, S^{(2)}]\|_F}{4\|S^{(2)}\|_F}. \tag{21}$$

Searching this parametrized space numerically yields a maximum speed limit of $T_{\max} \approx 0.373$ as shown in Fig. 3.

This result marginally improves upon the analytical bound $T = \sqrt{2}/4 \approx 0.354$ derived in Ref. [10], narrowing the gap to the exact solution $T_{\text{CNOT}} = \pi/4 \approx 0.785$ [24].

The fact that the SDP converges to the same perturbation for every sampled symmetry indicates that the symmetry-based bound has saturated for this system. Nevertheless, this application highlights the advantage of our approach: the ability to determine speed limits directly using structured search methods on an explicit vector space, without ever needing to solve the system's equations of motion.

B. Spinless particle on a chain

Next, we consider the case of a single spinless particle on a chain with $N \in \mathbb{N}$ sites. Adjacent sites are coupled via the drift Hamiltonian

$$H_d = J \sum_{j=1}^{N-1} \left(|j\rangle \langle j+1| + |j+1\rangle \langle j| \right), \tag{22}$$

where we control the potential at the first site of the chain

$$H_1 = |1\rangle \langle 1|. \tag{23}$$

For our calculations, we set the coupling strength $J = 1$. The Hilbert space $\mathcal{H} = \mathbb{C}^N$ grows only linearly with N.

Our target operation swaps the state between two ends of the chain

$$\mathrm{SWAP} = \mathbb{1} - |1\rangle\langle 1| - |N\rangle\langle N| + |1\rangle\langle N| + |N\rangle\langle 1| \ , \quad (24)$$

which lies within the reachable set of operations, as local control on a boundary node is sufficient for full controllability [25, 26]. Our primary focus is the scaling behavior of the speed limit for this operation as a function of chain length N.

We first analyze the symmetry space of the chain. Unlike in the two-qubit CNOT example, this system admits linear symmetries that can be directly used to calculate speed limits. Since the control Hamiltonian H_1 projects onto the first site, its commutant V is block-diagonal with dimensions 1×1 and $(N-1) \times (N-1)$, yielding $\dim(V) = 1 + (N-1)^2$. Full controllability renders the commutant W of the full DLA trivial, thus $\dim(W) = 1$. Consequently, the quotient space has dimension $\dim(V/W) = (N-1)^2$. Due to this quadratic scaling, a thorough search through this optimization space is not feasible for $N > 3$.

In Ref. [10], an analytical speed limit for the SWAP operation was derived by explicitly constructing a single candidate symmetry and a corresponding perturbation, showing an asymptotic scaling of $O\left(\sqrt{N}\right)$. A natural starting point is to inject the analytical symmetry candidate directly into our method, allowing the SDP solver to optimize the perturbation rather than relying on a hand-constructed one. The middle trace in Fig. 4 demonstrates that this numerical optimization improves upon the original analytical bound by an asymptotically constant factor. We extract this factor by fitting the ratio of the bounds, $R_T = T_{\mathrm{num}}(N)/T_{\mathrm{an}}(N)$, to the function $y(N) = c + a/N + b/N^2$ with $a, b, c \in \mathbb{R}$, which yields $c = 2.50006 \pm 2.3 \times 10^{-5}$. The numerical optimization therefore tightens the analytical speed limit by a factor of approximately 2.5. Additionally projecting this candidate symmetry onto the numerically constructed symmetry space prior to solving the SDP further tightens the bound for small N, as shown in the topmost trace, while preserving the same $O\left(\sqrt{N}\right)$ asymptotic scaling.

In the unprojected case, the symmetry $S^{(1)}_{\mathrm{num}}$ used as input is identical to the analytical candidate $S^{(1)}_{\mathrm{an}}$ for all N. Consequently, the commutator and symmetry-norm factors in Eq. (11) cancel when taking the ratio $T_{\mathrm{num}}/T_{\mathrm{an}}$, leaving only the dependence on the perturbation norms:

$$R_T = \frac{T_{\mathrm{num}}}{T_{\mathrm{an}}} = \frac{\|\Delta H\|_{\infty,\mathrm{an}}}{\|\Delta H\|_{\infty,\mathrm{num}}} = R_{\|\Delta H\|_\infty}. \quad (25)$$

This isolates the source of our improvement, confirming that it stems entirely from the SDP solver finding a perturbation with a significantly smaller spectral norm than the hand-constructed one. The perturbation designed in Ref. [10] was engineered to introduce a degeneracy in H_d, producing an eigenstate that is also invariant under the control H_1, thus restoring the symmetry. In contrast,

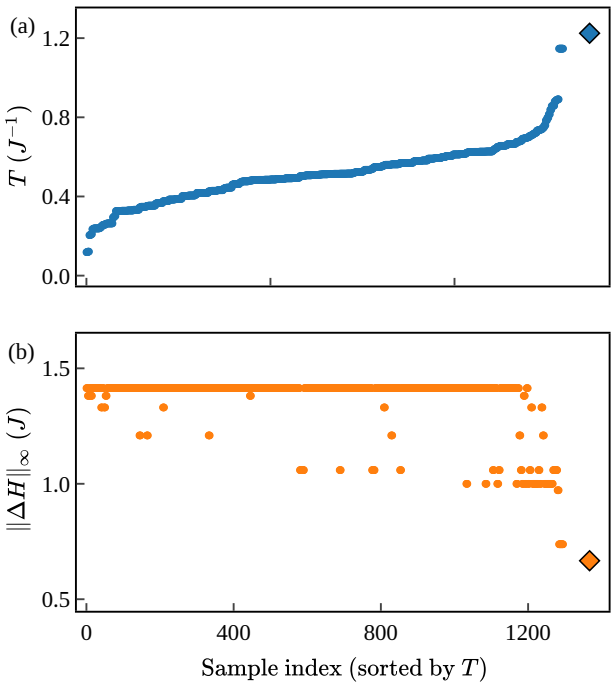

Figure 5. Optimization landscape illustrated by a systematic grid search over the four-dimensional symmetry space for the spinless-particle chain with $N = 3$ sites for $c_j \in [-1, 1]$ with six samples each, totaling 1296 samples. The panels show (a) the calculated speed limit T for the SWAP operation between the ends of the chain and (b) the corresponding minimal spectral norm of the perturbation $\|\Delta H\|_\infty$ obtained by solving the semidefinite program (SDP), with samples along the horizontal axis sorted by ascending T. The best-performing candidate (diamond at the far right) was obtained by directly injecting the physically motivated symmetry rather than from the grid search. The vast majority of grid points yield a constant value near $\|\Delta H\|_\infty \approx 1.4$, with only rare jumps down to ≈ 0.74, whereas the injected symmetry reaches ≈ 0.67.

the numerical perturbation merely rotates the existing eigenvectors of H_d without merging eigenspaces. Because this rotation is a subtler modification, it requires a smaller spectral norm $\|\Delta H\|_\infty$, directly translating into a tighter bound on T.

Motivated by this improvement, we aim to find an even tighter bound by utilizing the automated methods discussed in Sec. III C. Unfortunately, neither heuristic nor black-box methods converge to the known local maximum found with the physically motivated symmetry. Additionally, due to the high dimensionality for large N, the computational cost for the optimization algorithms scales poorly. We therefore apply a directional search using the full gradient of Eq. (20), in which the SDP solver provides the perturbation norm term ν_j. We find ν_j to oscillate consistently around 10^{-9} across all iterations, falling well below the threshold η established in Sec. III C and therefore failing to provide a meaningful search direction. The remaining two terms depend solely on the symmetry $S^{(1)}$

and its commutator with the target unitary, thus fully dictating the ascent direction. As a consequence, the gradients do not steer the optimization toward smaller perturbation norms $\|\Delta H\|_\infty$, yielding speed limits that do not exceed the previous results. To investigate this behavior, we perform a systematic grid search on the four-dimensional symmetry space of the $N = 3$ chain. The resulting speed limit T and perturbation norm $\|\Delta H\|_\infty$, plotted in Figs. 5(a) and 5(b) respectively, reveal an essentially flat $\|\Delta H\|_\infty$-landscape with narrow minima, directly explaining the vanishing gradient components ν_j. Resolving the narrow minimum achieved by the physically motivated symmetry requires deliberately matching its coefficients to 8 digits, illustrating why automated search strategies struggle to locate it, even more so in higher dimensions and why physical intuition remains the more effective approach for this system.

C. Ising ring with global control

As a final example, we apply the method to a ring of $N \in \mathbb{N}$ qubits. The system evolves under a nearest-neighbor Ising drift Hamiltonian with periodic boundary conditions $N + 1 \equiv 1$

$$H_d = J \sum_{n=1}^{N} Z_n Z_{n+1} \,, \tag{26}$$

subject to transverse global control fields

$$H_1 = \sum_{n=1}^{N} X_n \,, \quad H_2 = \sum_{n=1}^{N} Y_n \,. \tag{27}$$

We set the coupling strength $J = 1$. In contrast to the local control setting of Sec. IV A, these fields act uniformly across the ring, leaving the individual qubits indistinguishable to the controls. Consequently, the control algebra $\mathfrak{h}$ is invariant under the full permutation group S_N, whereas the full algebra $\mathfrak{g}$ depends on the adjacency structure of the qubits and is therefore symmetric only under the smaller dihedral group $D_N \subset S_N$ [27].

We evaluate the speed limit for preparing permutation-invariant Dicke states from the all-zero state $|0\rangle \equiv |0\rangle^{\otimes N}$. An N-qubit Dicke state with k excitations is defined as

$$|D_k^N\rangle = \binom{N}{k}^{-1/2} \sum_P P \left\{ |1\rangle^{\otimes k} |0\rangle^{\otimes N-k} \right\} \,, \tag{28}$$

summed over all distinct permutations P of the site indices [11]. We focus on the two-excitation state, abbreviated as $|D\rangle \equiv |D_2^N\rangle$. For this target state, numerical pulse optimization indicates that the minimum preparation time scales superlinearly with N [11], motivating us to examine whether our symmetry-based bounds capture this scaling.

Because H_d, H_1, and H_2 all commute with the action of D_N, the system preserves dihedral symmetry throughout its evolution. Consequently, no sequence of controls can drive the system outside a given symmetry sector. Both the initial state $|0\rangle$ and the target state $|D\rangle$ are permutation-symmetric and reside entirely within the D_N-invariant subspace $\mathcal{H}_{\text{inv}} \subset \mathcal{H}$. Following Ref. [11], we construct an isometry V_{iso} whose columns form an orthonormal basis of $\mathcal{H}_{\text{inv}}$, reducing the effective Hilbert space dimension from 2^N to $d = O(2^N/N)$. We project the Hamiltonians onto this sector via $H \mapsto V_{\text{iso}}^\dagger H V_{\text{iso}}$ for each $H \in \{H_d, H_1, H_2\}$, with state vectors transforming as $|\psi\rangle \mapsto V_{\text{iso}}^\dagger |\psi\rangle$. Within this invariant subspace, the state transfer from the initial state to the Dicke state can be realized by the target unitary

$$U = \mathbb{1} - |0\rangle\langle 0| - |D\rangle\langle D| + |0\rangle\langle D| + |D\rangle\langle 0| \,. \tag{29}$$

To confirm that $U \in \mathcal{R}$, we compute $\mathfrak{g} = \langle iH_d, iH_1, iH_2 \rangle_{\text{Lie}}$ via nested commutators for ring sizes $3 \le N \le 6$ and find $\dim(\mathfrak{g}) = d^2 = \dim(\mathfrak{u}(d))$ in each case. By the Lie algebra rank condition [12], the reachable set is therefore the entire unitary group, $\mathcal{R} = \mathrm{U}(d)$.

We now turn to the symmetry space of the system, beginning with linear symmetries. For $N = 3$, no non-trivial solutions exist, and for $N \ge 4$ all of them commute with the target unitary. By linearity, the entire space yields only the trivial bound $T \ge 0$, forcing us to search for quadratic symmetries. The dimension of the quotient space grows rapidly with N, yielding $\dim(V/W) = 2, 14, 26, 181$ for $N = 3, 4, 5, 6$, respectively. For systems with $N \ge 7$, computing this space ceases to be feasible. We examine the underlying scaling and computational bottlenecks in Sec. V B.

With a numerical basis for the quadratic symmetries at hand, the next step is to determine which symmetry $S^{(2)}$ yields the tightest speed limit. Unlike in the case of the spinless-particle chain of Sec. IV B, no analytical candidate is available, prompting us to employ the automated search strategies of Sec. III C. For $N = 3$, the two-dimensional symmetry space allows for an exhaustive grid search. Across all sampled symmetries $S^{(2)}$, the minimal perturbation norm remains constant at $\|\Delta H\|_\infty = 2$. When plotted over the coefficient space, the speed limit peaks along a diagonal line through the origin (excluding $S^{(2)} = 0$), reaching $T_{\max} \approx 0.194$. This profile closely resembles the CNOT landscape in Fig. 3, despite differing in magnitude. The similarity hints at a shared algebraic structure between the target unitaries, which lies beyond our current scope.

To extend the analysis to $N > 3$, we apply the heuristic and random-sampling optimization strategies. Regardless of the search heuristic or initial distribution, the optimization consistently converges to a constant minimal perturbation norm $\|\Delta H\|_\infty$ for each N, matching the observed behavior for $N = 3$. Furthermore, we observe vanishing gradients $|\nu_j| < \eta$ across the sampled symmetry space. Because the optimal perturbation norm remains constant across all sampled symmetries for any given N, it acts as a fixed factor in the denominator of Eq. (8). Consequently, maximizing the speed limit reduces to maximizing the normalized commutator norm from Eq. (21).

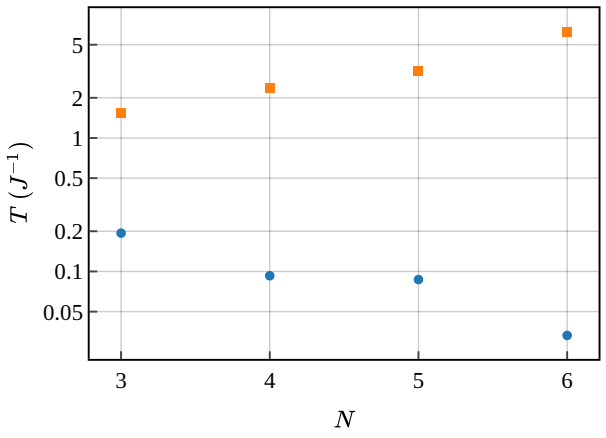

Figure 6. Speed limit and numerically optimized preparation times for the Dicke state $|D_2^N\rangle$ in an Ising-coupled qubit ring with global X and Y controls as a function of the number of qubits N. The symmetry-based lower bound (blue circles) is calculated using the symmetry-based framework. The upper trace (orange squares) represents the minimal preparation times obtained via pulse engineering, extracted from Ref. [11], where the $N = 3$ value is inferred from the $k = 1$ state via an instantaneous global X rotation. The symmetry-based lower bound decreases with N, whereas the pulse optimization results exhibit an increasing scaling behavior.

As shown in Fig. 6, our calculated bounds decrease with N. The bound therefore fails to capture the growth of the preparation time with system size observed in the pulse optimization results of Muratori *et al.* [11].

An inspection of the optimized perturbations ΔH shows that their structure depends on the parity of N. For even N, the solver converges to $\Delta H = -H_d$, canceling the drift entirely. For odd N, geometric frustration breaks the spectral symmetry of the ring, guiding the solver to shift the perturbed drift to the identity $\Delta H = \mathbb{1} - H_d$. Because these perturbed drifts commute trivially with any candidate symmetry, they yield symmetry-agnostic solutions whose spectral properties and consequences for the speed limit scaling are examined in Sec. V A.

V. DISCUSSION

In this section, we investigate the mechanism behind the differing scaling behaviors observed across our example systems. We then examine the computational bottlenecks of our method and discuss our choice of SDP solvers.

A. Origin of the differing scaling trends

Across our numerical calculations, the minimal perturbation norm $f(S)$ defined in Eq. (18) exhibits two distinct convergence behaviors. For both Ising-coupled qubit systems, $f(S)$ is completely independent of the candidate symmetry S, remaining constant across all probed symmetries for any given system size N. In contrast, the spinless particle on a chain yields non-trivial variation in $f(S)$ as S varies. What determines whether $f(S)$ depends on the choice of symmetry, and how does this invariance determine the scaling behavior of the calculated quantum speed limit T?

A natural initial hypothesis is to attribute this discrepancy to the symmetry order. While linear symmetries suffice for the spinless-particle chain, both qubit cases require quadratic symmetries. Moving from linear to quadratic symmetries affects three properties: it lifts the operators to the doubled Hilbert space $\mathcal{H} \otimes \mathcal{H}$, algebraically guarantees the existence of non-trivial symmetries for non-fully controllable systems [13], and introduces an extra factor of $1/2$ in Eq. (8). The existence guarantee and the factor of $1/2$ do not change how T scales with N. The lift to $\mathcal{H} \otimes \mathcal{H}$, however, changes the restoring constraint and may therefore affect the optimal ΔH. Since our examples using linear and quadratic symmetries are also different physical systems, we cannot separate the effect of the symmetry order from that of the system.

In both Ising systems, for each N, the SDP returns the same perturbation for every sampled symmetry. This originates from the restoring constraint Eq. (12) or Eq. (9), which admits two symmetry-agnostic solutions that map directly to the observed perturbations in Secs. IV A and IV C. The first solution,

$$\Delta H_{\text{cancel}} = -H_d\,, \tag{30}$$

cancels the drift entirely. Its spectral norm is given by the spectral radius of the drift Hamiltonian:

$$\|\Delta H_{\text{cancel}}\|_\infty = \max_i |\sigma_i|\,, \tag{31}$$

where σ_i are the eigenvalues of H_d. The second construction,

$$\Delta H_{\text{flatten}} = c\mathbb{1} - H_d\,, \tag{32}$$

with $c \in \mathbb{R}$, shifts the drift to a scalar multiple of the identity, collapsing the spectrum into a single degenerate eigenvalue. This perturbation norm is minimized by centering the shift at the spectral midpoint $c = (\sigma_{\max} + \sigma_{\min})/2$, yielding

$$\min_c \|\Delta H_{\text{flatten}}\|_\infty = \frac{\sigma_{\max} - \sigma_{\min}}{2}. \tag{33}$$

Because both constructions render the drift equal to zero or a scalar multiple of the identity, the perturbed Hamiltonian and its symmetric embedding commute trivially with any candidate symmetry S. We define the minimum of these two norms as the symmetry-independent perturbation ceiling:

$$F(N) \coloneqq \min\left(\max_i |\sigma_i|, \frac{\sigma_{\max} - \sigma_{\min}}{2}\right), \tag{34}$$

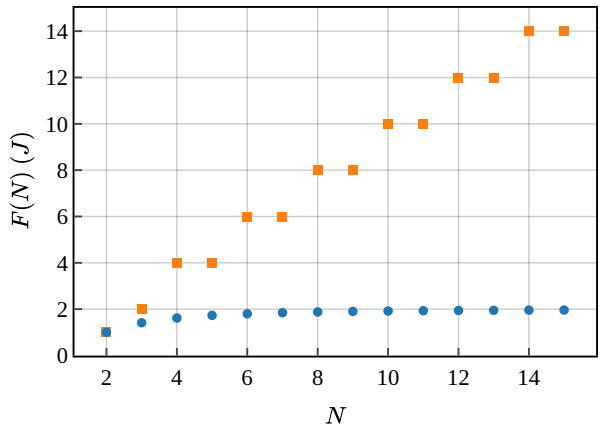

Figure 7. Scaling of the symmetry-independent perturbation ceiling $F(N)$ with system size N. The spinless-particle chain is indicated by the blue circles, while the Ising-coupled qubit ring is indicated by the orange squares. The ceiling for the qubit ring grows monotonically, whereas that of the spinless-particle chain asymptotically converges to $2J$.

which is determined exclusively by the spectrum of H_d. This ceiling establishes an upper bound on the convex optimization of Eq. (18):

$$f(S) \leq F(N)\,. \tag{35}$$

Our numerical optimizations for the Ising-coupled qubits consistently hit this ceiling across the entire sampled space. By contrast, a grid search on the spinless-particle chain successfully identified symmetries with perturbation norms below this bound, as illustrated in Fig. 5. For the qubit examples, even fine grid searches on small systems, like the two-qubit CNOT and the three-qubit ring, found no symmetries yielding $f(S) < F(N)$. However, because our search relies on finite sampling of the symmetry space, we cannot rule out their existence.

To see how $F(N)$ dictates the scaling of the speed limit, consider the combined bound of Eq. (13) with $\|\Delta H\|_\infty = f(S)$. Bounding the commutator ratio by applying the triangle inequality and the unitary invariance of the Frobenius norm yields

$$\frac{\|[U^{\otimes p}, S]\|_F}{\|S\|_F} \leq \frac{\|U^{\otimes p}S\|_F + \|SU^{\otimes p}\|_F}{\|S\|_F} = 2\,. \tag{36}$$

Since this ratio is bounded by a constant, the asymptotic scaling of the speed limit is dictated entirely by $1/f(S)$. By evaluating Eq. (34) directly from the spectrum of H_d, we obtain the exact ceiling $F(N)$ independently of the numerical optimization. As shown in Fig. 7, this ceiling grows linearly with N for the Ising models, directly accounting for the $O(1/N)$ decay in the speed limit T. For the spinless-particle chain, $F(N)$ grows with N but rapidly saturates at $2J$. Consequently, without symmetries that outperform the trivial solutions, the baseline speed limit would decrease before plateauing to an $O(1)$ constant. The achieved $O\left(\sqrt{N}\right)$ scaling on the chain is therefore not a generic consequence of the perturbation ceiling, but requires finding a symmetry family whose minimal perturbation norm shrinks with N.

The scaling behavior of $F(N)$ determines the worst-case scaling trend of the speed limit whenever no symmetry yielding $f(S) < F(N)$ can be found. Whether the investigated qubit systems are fundamentally constrained by their DLAs to $f(S) = F(N)$ across their entire symmetry spaces, or whether favorable symmetries with $f(S) < F(N)$ exist, remains an open problem.

B. Computational bottlenecks and solver choice

1. *Computational bottlenecks*

The procedure outlined in Fig. 2 contains two computationally demanding stages: constructing a basis for the symmetry space via the null-space formulation in Eq. (15), and solving the SDP in Eq. (19) for a given symmetry. While the symmetry basis is computed upfront and can be stored for subsequent runs, the SDP must be solved repeatedly during heuristic and gradient-based searches. In our calculations, the symmetry basis construction accounts for the primary runtime cost, while initial canonicalization of the SDP drives peak memory demand for large systems.

The stacked constraint matrix $\mathcal{M} \in \mathbb{C}^{gm\times m}$ defined in Eq. (15) scales rapidly with system size. Computing its kernel via rank-revealing QR decomposition or SVD scales as $O(gm^3)$, with QR carrying a smaller prefactor [14, p. 278, pp. 284–285]. In terms of the Hilbert space dimension d, this demands $O(gd^6)$ floating-point operations (FLOPs) for linear symmetries and $O(gd^{12})$ for quadratic symmetries. At larger system sizes, dense representations quickly become prohibitive, necessitating sparse matrix formats and sparse linear algebra routines.

In contrast, the SDP formulation does not require vectorization. The commutator equality constraints are enforced directly on $\mathcal{H}$ for linear symmetries and on $\mathcal{H}\otimes\mathcal{H}$ for quadratic symmetries, while the LMIs bounding the spectral norm remain defined on the physical Hilbert space $\mathcal{H}$. Because the outer optimization queries the SDP repeatedly with varying coefficients c_j in Eq. (17), DPP significantly reduces the per-iteration overhead. This benefit is most evident for the largest problem instance considered: the Ising ring example of Sec. IV C with $N = 6$, whose quotient space V/W is 181-dimensional. Initial canonicalization and solving require approximately 13 minutes, whereas subsequent iterations with updated coefficients drop to roughly 16 seconds. During the canonicalization step, memory rather than runtime becomes the limiting factor, with the solver's internal data structures peaking at approximately 94 GB.

Reducing the Hilbert space dimension d directly ad-

dresses both bottlenecks. As demonstrated in Sec. IV C, projecting onto invariant subspaces, such as the D_N-invariant sector of the ring, reduces d from 2^N to $O(2^N/N)$, making larger systems computationally accessible. However, this approach is applicable only when both the Hamiltonian dynamics and the target unitary leave the chosen subspace invariant, which restricts the class of addressable target gates.

2. *Solver choice*

The CVXPY modeling framework [17] supports multiple backend solvers differing in their underlying algorithms and licensing terms. We employ both the open-source solver SCS, CVXPY's default solver for SDPs, which implements Douglas–Rachford splitting [20], and MOSEK, a commercial interior-point solver available under an academic license [19]. Under default configurations, SCS terminates at a moderate-precision residual tolerance of 10^{-4}, whereas MOSEK reaches a high-precision relative duality gap of 10^{-8}. Furthermore, as described in Sec. III C, the SCS solver supports differentiating the optimal solution with respect to problem parameters, enabling gradient-based optimization.

Our selection of solver is pragmatic: because individual solve times remain modest during our calculations, we default to MOSEK with its higher precision for all heuristic and black-box methods that do not require gradients, switching to SCS otherwise. The variation in solver precision does not affect our conclusions, as distinguishing meaningful differences in the optimal perturbations and the resulting speed limit T does not require a resolution beyond 10^{-4}.

VI. CONCLUSION AND OUTLOOK

In this work, we translated the theoretical framework of target-specific, symmetry-based QSLs [10] into a systematic numerical method. Requiring only the drift and control Hamiltonians alongside the target operation as inputs, the procedure decomposes into three primary stages: constructing an explicit basis for the admissible symmetry space, searching this space for an optimal candidate symmetry, and optimizing for the minimal-norm perturbation that restores it. The first and third stages reduce to well-posed computational tasks: a kernel computation and a convex SDP, respectively, both guaranteed to yield the exact or globally optimal solution. In particular, the SDP removes the need for hand-constructed perturbations. The intermediate stage, which strongly influences the tightness of the resulting speed limit, is non-convex and relies on complementary strategies, including gradient ascent, black-box heuristics, and physically motivated candidate injections.

Applying the method to three physical architectures illustrates both its practical utility and current limitations. In the two-qubit CNOT example, an automated search verified that our numerical optimization can reproduce and slightly tighten the hand-derived symmetry bound of Ref. [10]. For the SWAP operation between the ends of a spinless-particle chain, the SDP solver discovered a perturbation mechanism based on rotating existing eigenstates of the drift Hamiltonian rather than forcing spectral degeneracy, tightening the analytical bound by a factor of approximately 2.5 while preserving the $O(\sqrt{N})$ scaling. Conversely, for the Dicke state preparation on a globally controlled qubit ring, the calculated bound decayed as $O(1/N)$ for all sampled symmetries, failing to capture the superlinear scaling observed via pulse optimization [11].

This behavior originates from the SDP solver returning only symmetry-independent perturbations for all sampled candidate symmetries in the qubit architectures. This pattern revealed a symmetry-independent perturbation ceiling, $F(N)$, set by perturbations that cause the perturbed drift Hamiltonian to commute trivially with any operator, specifically by canceling the drift entirely or shifting it to a scalar multiple of the identity. Because this ceiling depends solely on the spectrum of the drift Hamiltonian, it explains the worst-case scaling behavior of the speed limits observed in the qubit systems. It also highlights the value of physically motivated candidate symmetries, which proved decisive for achieving tighter bounds in the spinless-particle case.

With this method in place, an immediate next step is to evaluate target-specific QSLs across other systems and operations, particularly where bounds are currently loose or unknown. Beyond these direct applications, several avenues emerge to extend and advance the methodology itself:

Escaping the perturbation ceiling: Symmetries that satisfy $f(S) < F(N)$ appear to be rare and localized within the symmetry space. Establishing analytical criteria to determine whether a given DLA admits such symmetries would serve a dual purpose: it would prevent fruitless searches when such solutions are algebraically ruled out, while justifying computational investment in challenging searches when tighter bounds are guaranteed to exist. Empirically, using the method to investigate the effect of strategically placed local controls in globally driven systems, such as the Ising ring, could expose additional structure in the symmetry space and help bypass the ceiling.

Numerics-to-theory feedback: The discovery of a non-obvious perturbation mechanism in the spinless-particle chain demonstrates that convex optimization can find solutions missed by intuition. Leveraging such solver-discovered perturbations to extend the repertoire of symmetry-perturbation constructions could complement theoretical efforts, closing the loop between the numerical and analytical approaches.

Algorithmic efficiency: The primary computational bottleneck remains the null-space calculation of the stacked constraint matrix defining the symmetry space. An alternative would be to compute pairwise null-space intersec-

tions [14, pp. 328–329] sequentially across the generators, restricting each new set of constraints to the current candidate subspace rather than decomposing the entire stacked matrix in a single computation. This could reduce the computational cost, thus extending the reach to larger systems.

State-to-state extension: Extending the theoretical framework beyond unitary and simulable-Hamiltonian targets to state-to-state transfers would broaden the class of problems it can address. Such an extension would directly apply to the Dicke state preparation studied in this work, which is fundamentally a state-to-state problem currently handled by choosing one specific unitary among the many that implement this transfer.

ACKNOWLEDGMENTS

This research was supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy with funds from the Hightech Agenda Bayern and is part of the funding measure Munich Quantum Valley. The authors gratefully acknowledge the computational and data resources provided through the joint high-performance data analytics (HPDA) project "terrabyte" of the German Aerospace Center (DLR) and the Leibniz Supercomputing Center (LRZ).

DATA AVAILABILITY

The data that support the findings of this article are available from the corresponding author upon reasonable request.

[1] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[2] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe, EPJ Quantum Technol. **9**, 19 (2022).
[3] L. Mandelstam and I. E. Tamm, The uncertainty relation between energy and time in non-relativistic quantum mechanics, in *Selected Papers* (Springer, Berlin, 1991) p. 115.
[4] N. Margolus and L. B. Levitin, The maximum speed of dynamical evolution, Physica D **120**, 188 (1998).
[5] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, J. Magn. Reson. **172**, 296 (2005).
[6] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, Phys. Rev. A **84**, 022326 (2011).
[7] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, Time-Optimal Quantum Evolution, Phys. Rev. Lett. **96**, 060503 (2006).
[8] T. Koike, Quantum brachistochrone, Philos. Trans. R. Soc., A **380**, 20210273 (2022).
[9] D. Burgarth, J. Borggaard, and Z. Zimborás, Quantum distance to uncontrollability and quantum speed limits, Phys. Rev. A **105**, 042402 (2022).
[10] M. Wiedmann and D. Burgarth, Quantum speed limits from symmetries in quantum control, J. Phys. A: Math. Theor. **59**, 065301 (2026).
[11] A. Muratori, V. M. Stojanovic, E. Cuestas, T. Calarco, and F. Motzoi, *W*- and Dicke-state engineering using optimal global control in nearest-neighbor coupled ring-shaped qubit arrays (2025), arXiv:2512.19545 [quant-ph].
[12] D. D'Alessandro, *Introduction to Quantum Control and Dynamics*, 2nd ed., Advances in Mathematics Series (CRC Press, Boca Raton, 2021).
[13] R. Zeier and T. Schulte-Herbrüggen, Symmetry principles in quantum systems theory, J. Math. Phys. **52**, 113510 (2011).
[14] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins Studies in the Mathematical Sciences (Johns Hopkins Univ. Press, Baltimore, MD, 2013).
[15] E. Provenzi, *From Euclidean to Hilbert Spaces: Introduction to Functional Analysis and Its Applications*, 1st ed. (Wiley, 2021).
[16] S. P. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, 2004).
[17] S. Diamond and S. Boyd, CVXPY: A Python-embedded modeling language for convex optimization, J. Mach. Learn. Res. **17**, 1 (2016).
[18] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter, Differentiable convex optimization layers, in *Advances in Neural Information Processing Systems* (2019) p. 9558.
[19] MOSEK ApS, *MOSEK Optimizer API for Python 11.2.3* (2026).
[20] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, Conic optimization via operator splitting and homogeneous self-dual embedding, J. Optim. Theory Appl. **169**, 1042 (2016).
[21] C. R. Harris *et al.*, Array programming with NumPy, Nature **585**, 357 (2020).
[22] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nat. Methods **17**, 261 (2020).
[23] A. E. Siemenn, Z. Ren, Q. Li, and T. Buonassisi, Fast Bayesian optimization of needle-in-a-haystack problems using zooming memory-based initialization (ZoMBI), npj Comput. Mater. **9**, 79 (2023).

[24] G. Vidal, K. Hammerer, and J. I. Cirac, Interaction cost of nonlocal gates, Phys. Rev. Lett. **88**, 237902 (2002).

[25] X. Wang, P. Pemberton-Ross, and S. G. Schirmer, Symmetry and subspace controllability for spin networks with a single-node control, IEEE Trans. Autom. Control **57**, 1945 (2012).

[26] D. Burgarth, D. D'Alessandro, L. Hogben, S. Severini, and M. Young, Zero forcing, linear and quantum controllability for systems evolving on networks, IEEE Trans. Autom. Control **58**, 2349 (2013).

[27] X. Wang, M. S. Okyay, A. Kumar, and B. M. Wong, Accelerating Quantum Optimal Control of Multi-Qubit Systems with Symmetry-Based Hamiltonian Transformations, AVS Quantum Sci. **5**, 043801 (2023).