Research Square

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

Reductions of QAOA Induced by Classical Symmetries: Theoretical Insights and Practical Implications

Boris Tsvelikhovskiy

tsvelibor@gmail.com

University of California, Riverside

Bao Bach

University of Delaware

Jose Falla

University of Delaware

Ilya Safro

University of Delaware

Article

Keywords:

Posted Date: September 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10706941/v1

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

Additional Declarations: No competing interests reported.

Reductions of QAOA Induced by Classical Symmetries: Theoretical Insights and Practical Implications

Boris Tsvelikhovskiy [1,*] Bao Bach [2] Jose Falla [3] and Ilya Safro [2,3]

[1] *Department of Mathematics, University of California, Riverside, CA 92521, USA*
[2] *Department of Computer and Information Sciences, University of Delaware, Newark, DE 19716, USA*
[3] *Department of Physics and Astronomy, University of Delaware, Newark, DE 19716, USA*

We demonstrate that classical symmetries can be systematically exploited to design the Quantum Approximate Optimization Algorithm (QAOA). Focusing on the MaxCut problem, we analyze reduced QAOA instances obtained by fixing a single variable. We show that the underlying algebraic structure is highly sensitive to this choice. In particular, we construct graph families where this reduction collapses the DLA dimension from exponential to quadratic. These results suggest that symmetry-aware reduction can guide the design of expressive, trainable QAOA circuits. We prove that every graph admits an embedding into a larger graph with $\mathcal{O}(|V|^2)$ vertices and edges, preserving MaxCut solutions while matching the expressivity of multi-angle QAOA after reduction. Within this extension, every state in the reduced Hilbert space, including all optimal solutions, is reachable at finite depth. Conversely, the exponentially growing DLA predicts the onset of barren plateaus.

I. INTRODUCTION

What could be simpler than fixing one binary variable? In a classical optimization problem with a global bit-flip symmetry, this operation merely removes a redundant representation: the core problem and its objective values remain unchanged, and all original optimal solutions are immediately recoverable. One might therefore expect the corresponding quantum algorithm to change only superficially, by using one fewer qubit. We show that this intuition is fundamentally wrong. Specifically, for the Quantum Approximate Optimization Algorithm (QAOA) applied to MaxCut, fixing a single vertex changes the Hamiltonians generating the quantum dynamics.Depending on which vertex is fixed, this reduction can unlock the entire remaining Hilbert space, alter the trainability of the variational landscape, or collapse the algebraic complexity of the dynamics from exponential to quadratic. Thus, a step that is almost trivial classically can produce a fundamentally different quantum algorithm.

This observation has direct implications for one of the central challenges in variational quantum computing: balancing expressivity and trainability. A circuit must explore a sufficiently large state space to reach high-quality solutions, but excessively large and unstructured dynamical spaces can lead to exponentially suppressed gradients and barren plateaus. Our results show that this balance can potentially be controlled before the quantum circuit is executed, through the choice of a classical symmetry representative. Symmetry reduction is therefore not merely a way to save one qubit. It provides a mechanism for redesigning the quantum dynamics while preserving the original optimization problem.

We study this phenomenon through dynamical Lie algebras (DLAs), which characterize the continuous transformations generated by the Hamiltonians of a parameterized quantum circuit. The DLA determines the effective state space accessible to the circuit and provides a rigorous link between circuit architecture, controllability, expressivity, and gradient concentration [7, 15, 16, 22]. A small or highly structured DLA restricts the possible evolution, whereas a sufficiently large DLA can enable arbitrary state preparation within the relevant Hilbert space. At the same time, exponentially large DLAs are closely associated with concentration of the optimization landscape and the emergence of barren plateaus [8, 17, 29]. The DLA therefore offers a natural mathematical framework for understanding the competing requirements of expressivity and trainability.

The *Quantum Approximate Optimization Algorithm* [6] is a prominent variational quantum algorithm for combinatorial optimization. It alternates evolution under a problem Hamiltonian, which encodes the objective function, and a mixer Hamiltonian, which moves amplitude between candidate solutions. QAOA has

* Corresponding author: tsvelibor@gmail.com

been applied to graph problems including MaxCut and Maximum Independent Set [6], community detection [24], and graph partitioning [28], and remains one of the primary candidates for investigating quantum advantage in combinatorial optimization [5, 10, 12, 21]. Despite its seemingly simple structure, however, characterizing or explicitly computing the resulting DLA and its effect on the optimization landscape remains computationally difficult in general.

Previous studies of QAOA with the standard Pauli-X mixer have characterized its DLAs for selected graph families, including paths, cycles, and complete graphs [1, 13]. For Erdős–Rényi random graphs with edge probability $p = 0.5$, the QAOA-MaxCut DLA was shown to be, with probability at least $1 - e^{-\Omega(n)}$, a direct sum of either one or two simple Lie algebras, each of dimension $\Theta(4^n)$ [19]. DLA characterizations have also been obtained for Grover-mixer QAOA [25] and for constrained XY mixers with selected connectivity structures [9, 14]. These studies analyze the dynamics generated by a fixed encoding of the optimization problem.

Symmetry reduction has been used previously to eliminate redundant configurations and decrease the effective size of quantum optimization problems [4, 23, 26, 27]. Its effect on the associated DLA, however, has remained largely unexplored. This raises a fundamental question: does removing a classical symmetry merely reduce the number of qubits, or can it change the expressive and trainability properties of the quantum ansatz itself?

We answer this question for binary optimization problems invariant under a global bit-flip, with MaxCut as the central example. A MaxCut solution is unchanged when all bits are simultaneously inverted, because the two sides of a cut can be exchanged without changing its value. Fixing the bit associated with one vertex therefore selects one representative from each pair of equivalent solutions and reduces the problem from n to $n-1$ variables. Applying this construction for each vertex produces n reduced problems that are equivalent as classical optimization tasks. However, their quantum formulations need not be equivalent. Fixing a vertex removes its associated mixer term and converts every incident two-qubit interaction into a one-qubit term acting on a neighboring vertex. The resulting Hamiltonian generators, and hence their DLAs, depend strongly on which vertex is selected.

Our first main result shows that this reduction can remove a fundamental restriction present in unreduced QAOA. For MaxCut, the global bit-flip symmetry decomposes the full Hilbert space into invariant sectors, and QAOA initialized in the uniform superposition state remains confined to the global bit-flip symmetry-protected subspace [13]. After fixing a vertex v, let Γ_v denote the graph obtained by removing v and adjacent edges. If Γ_v is connected and not bipartite, we prove that the free reduced DLA (generated by the separate terms of the two Hamiltonians) is

$$\mathfrak{g}^{v}_{\Gamma,\mathrm{free}} \simeq \mathfrak{su}(2^{n-1}).$$

The reduced Hilbert space is therefore no longer split into inaccessible symmetry sectors. The corresponding multi-angle QAOA is pure-state controllable: at sufficiently large finite depth, it can prepare any target state in the reduced Hilbert space, up to a global phase. In simple terms, removing one classically redundant bit can unlock the entire remaining quantum state space.

Our second main result shows that this maximal expressivity can also be obtained using the conventional QAOA parametrization. We derive explicit graph-theoretic conditions sufficient for the standard reduced DLA, generated by only the problem and mixer Hamiltonians, to coincide with the free reduced DLA generated by independently controlled local terms. Furthermore, for every connected graph $\Gamma = (V, E)$ and every selected vertex $v \in V$, we construct an extended graph $\widehat{\Gamma}_v$ satisfying

$$|V(\widehat{\Gamma}_v)|,\ |E(\widehat{\Gamma}_v)| \in \mathcal{O}(|V|^2)$$

for which these two reduced DLAs coincide. The extension preserves the optimization problem: an optimal MaxCut solution of the original graph is recovered by discarding the auxiliary bits. This improves on the $\mathcal{O}(|E|^2)$ extension required in the unreduced setting [19], which may scale as $\mathcal{O}(|V|^4)$ for dense graphs.

Our third main result demonstrates that symmetry reduction can produce the opposite, but equally useful, effect: it can drastically simplify the quantum dynamics. For the family of k-armed spider graphs $\mathcal{O}_{m_1,\ldots,m_k}$ shown in Figure 4, the standard unreduced DLA has exponential dimension [19]. Fixing the central vertex separates the graph into independent paths and reduces the corresponding DLA dimension to quadratic growth. This gives an explicit family in which removing one redundant variable changes the algebraic complexity from exponential to polynomial. Because DLA dimension is connected to gradient concentration, such a reduction also predicts a substantially more trainable optimization landscape.

We complement these theoretical results with numerical experiments on asymmetric graphs. Exact calculations for all asymmetric graphs on six and seven vertices support the conjecture that there exists a choice of reduction vertex whose standard reduced DLA is of smaller dimension than the original DLA. For connected asymmetric graphs with 7–15 vertices, exact DLA calculations become prohibitively expensive, and we instead use the variance of the QAOA loss-function gradient as a proxy for DLA dimension [22]. Across the tested instances, appropriately selected reductions consistently produce larger gradient variance than the unreduced circuits, supporting the prediction that symmetry reduction can improve trainability.

These results reveal a previously overlooked distinction between classical and quantum problem reduction. Classically, fixing a symmetry-related variable simply selects an equivalent representation of the same solution. Quantum mechanically, the choice determines the generators of the dynamics and can therefore change which states are reachable, whether the ansatz is fully controllable, and whether its DLA has exponential or polynomial dimension. Symmetry reduction thus provides a principled way to engineer the expressive and trainability properties of QAOA without changing the optimization objective.

Figure 1 summarizes the principal differences between the standard and reduced frameworks.

II. RESULTS

Classical Pre-processing and Reduced QAOA. Let $\mathbb{B}^n := \{0,1\}^n$ denote the set of all binary strings of length n, and let $\mathcal{S}_{2^n}$ denote the symmetric group acting on the 2^n elements of $\mathbb{B}^n$. A significant class of optimization problems entails finding elements in $\mathbb{B}^n$ on which a given function

$$F : \mathbb{B}^n \to \mathbb{R} \tag{1}$$

(i.e., the optimization objective with possible constraints) is either minimized or maximized.

If the function $F(x)$ is invariant with respect to a permutation $g \in \mathcal{S}_{2^n}$, i.e., $F(g^{-1}(x)) = F(x)$ for any $x \in \mathbb{B}^n$, then g is called a *symmetry* of F. Let $G \subseteq \mathcal{S}_{2^n}$ be a subgroup of symmetries of F. Then the set $\mathbb{B}^n$ can be expressed as a disjoint union of G-orbits, $\mathcal{O}_1, \mathcal{O}_2, \ldots, \mathcal{O}_m$, where m is the number of distinct orbits:

$$\mathbb{B}^n = \bigsqcup_{j=1}^{m} \mathcal{O}_j.$$

Rather than working directly with the original combinatorial optimization problem on $\mathbb{B}^n$, we may instead pass to the quotient induced by the group action. Specifically, consider the quotient set

$$\mathbb{B}_G := \mathbb{B}^n / G, \tag{2}$$

where two elements of $\mathbb{B}^n$ are equivalent if and only if they belong to the same G-orbit.

For clarity of exposition, we will henceforth assume that the action of the group G on $\mathbb{B}^n$ is *free*, that is, no non-identity element of G fixes any point of $\mathbb{B}^n$. In this case, all G-orbits have cardinality equal to $|G|$, and therefore the quotient has size

$$|\mathbb{B}_G| = \frac{2^n}{|G|}.$$

Since the objective function F is constant along G-orbits, we obtain a well-defined function

$$\widetilde{F} : \mathbb{B}_G \to \mathbb{R}, \qquad \widetilde{F}(\mathcal{O}_j) := F(x), \tag{3}$$

where $x \in \mathcal{O}_j$ is any representative of the equivalence class. This yields the reduced combinatorial optimization problem of minimizing $\widetilde{F}$ on $\mathbb{B}_G$.

Remark II.1. The solutions of the original problem, i.e., the minima (maxima) of F on $\mathbb{B}^n$, are in $|G|$-to-1 correspondence with the solutions of the reduced problem, namely the minima (maxima) of $\widetilde{F}$ on $\mathbb{B}_G$. Equivalently, the solutions of the original problem are precisely the G-orbits of the solutions of the reduced problem.

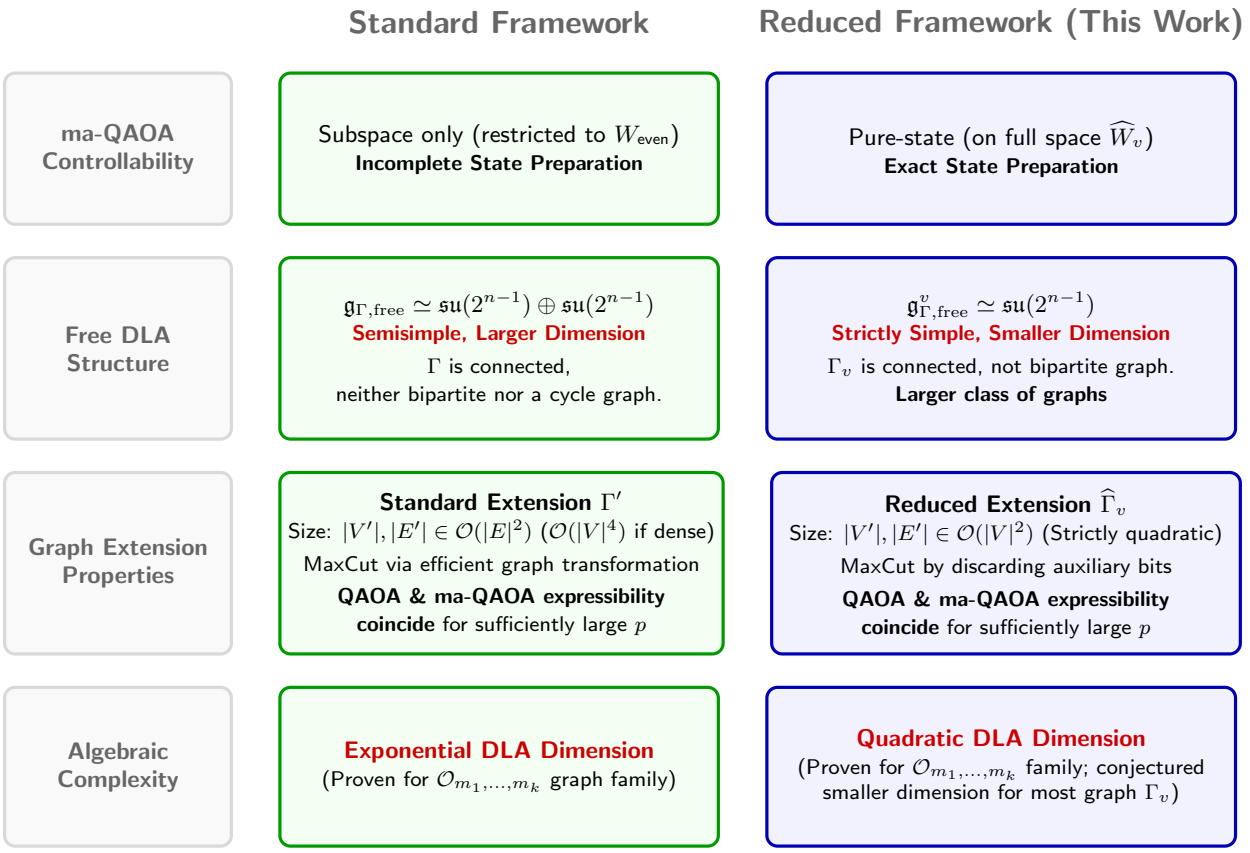


FIG. 1. A comparative overview of the standard QAOA framework versus our proposed reduced QAOA framework. **Controllability (Row 1).** Standard ma-QAOA yields only subspace controllability for the subspace W_{even} (spanned by states with an even number of minus states in the Hadamard basis) for connected graphs that are neither bipartite nor cycles (see Theorem 1 in [13]). Conversely, our reduced ma-QAOA achieves full pure-state controllability over the entire reduced Hilbert space $\widehat{W}_v$ for all connected graphs, except when the reduced graph is disconnected or bipartite (see Theorem II.4). **Free DLA Structure (Row 2).** While the standard framework yields a larger, semisimple free DLA isomorphic to $\mathfrak{su}(2^{n-1}) \oplus \mathfrak{su}(2^{n-1})$ for connected graphs that are neither bipartite nor cyclic (see Theorem 1 in [13]), our reduced framework yields a free DLA that is simple, isomorphic to $\mathfrak{su}(2^{n-1})$, for pairs (Γ, v) whose reduced graph Γ_v is connected and not bipartite (see Theorem II.4). **Graph Extensions (Row 3).** For any connected graph $\Gamma = (V, E)$, the standard framework constructs an extended graph Γ' with size $|V'|, |E'| \in \mathcal{O}(|E|^2)$ such that the standard and free DLAs coincide, enabling MaxCut mapping via an efficient graph transformation (see Theorem 5 in [19]). Our reduced framework requires an extended graph $\widehat{\Gamma}_v$ of strictly quadratic size $|V'|, |E'| \in \mathcal{O}(|V|^2)$ to align the standard and free reduced DLAs, where the original MaxCut solution is reconstructed immediately by discarding auxiliary bits (see Theorem II.5). Notably, on both families of extension graphs, QAOA and ma-QAOA parameter circuits achieve identical expressibility for a sufficiently large depth p. **Algebraic Complexity (Row 4).** The standard DLA dimension exhibits exponential growth for the $\mathcal{O}_{m_1,\ldots,m_k}$ graph family (see [19, Corollary 26]), whereas our reduced DLA dimension grows only quadratically for this family and is conjectured to yield strictly smaller dimensions for any general graph Γ.

QAOA for the Original Problem. Below we provide an overview of the Quantum Approximate Optimization Algorithm in a language tailored to the needs of this paper. For a more comprehensive treatment, the reader is referred to [3, 6, 9].

Let W be a vector space of dimension 2^n with computational basis $\{|x\rangle\}_{x\in\mathbb{B}^n}$, indexed by elements of $\mathbb{B}^n$. A Hamiltonian H_P is said to *represent* a function $F : \mathbb{B}^n \to \mathbb{R}$ if

$$H_P |x\rangle = F(x) |x\rangle \quad \text{for all } x \in \mathbb{B}^n. \tag{4}$$

The quantum formulation of the optimization problem associated with F is obtained via the following classical-to-quantum correspondence:

- Set $\mathbb{B}^n \rightsquigarrow$ Hilbert space $W = (\mathbb{C}^2)^{\otimes n}$;
- Objective function $F \rightsquigarrow$ Linear operator H_P (problem Hamiltonian) acting on W;
- Minima of F on $\mathbb{B}^n \rightsquigarrow$ Lowest energy states of H_P in W;
- Action of G on $\mathbb{B}^n$ that preserves $F \rightsquigarrow$ Unitary representation of G on W that commutes with H_P.

A central component of QAOA is the so-called *mixer Hamiltonian*, denoted H_M. The negative of this operator, $-H_M$, is required to have a unique ground state $|\xi\rangle \in W$. The fundamental idea behind QAOA is to gradually deform H_M into H_P through a series of quantum transformations, in such a way that the ground state at each stage is mapped to a ground state of the next.

The QAOA algorithm initializes the system in the ground state $|\xi\rangle$ of $-H_M$ and applies p alternating evolutions under the problem and mixer Hamiltonians. For classically optimized parameters

$$\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p), \qquad \boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_p),$$

with $\beta_j \in [0, 2\pi)$ and $\gamma_j \in [0, \pi)$, the resulting unitary is

$$U(\boldsymbol{\beta}, \boldsymbol{\gamma}) \coloneqq e^{-i\beta_p H_M} e^{-i\gamma_p H_P} \cdots e^{-i\beta_1 H_M} e^{-i\gamma_1 H_P}. \tag{5}$$

The depth-p QAOA state is therefore

$$|\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle \coloneqq U(\boldsymbol{\beta}, \boldsymbol{\gamma}) \, |\xi\rangle \, .$$

The parameters are optimized classically using $\langle\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})|H_P|\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle$ as the cost function. Figure 2 provides a schematic overview of the procedure.

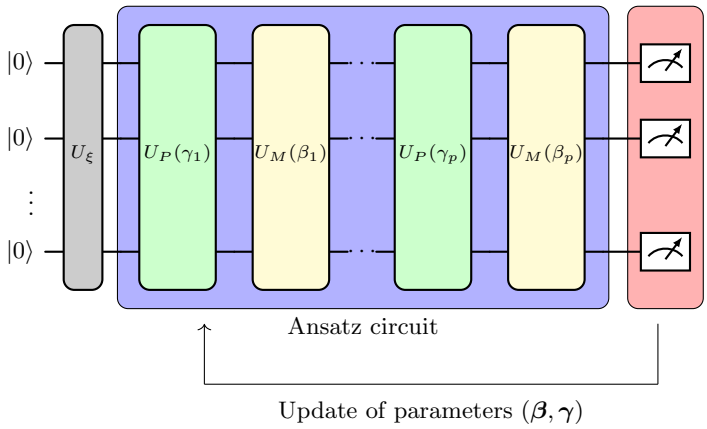


FIG. 2. **Schematic illustration of the QAOA circuit.** The initial state preparation unitary U_ξ with $U_\xi(|0\rangle^{\otimes n}) = |\xi\rangle$ is followed by p alternating applications of unitaries $U_P(\gamma_j) := e^{-i\gamma_j H_P}$ and $U_M(\beta_j) := e^{-i\beta_j H_M}$, generated by the problem Hamiltonian H_P and the mixer Hamiltonian H_M, respectively. At the end of the circuit, the state is measured in the computational basis. Each measurement outcome $x \in \mathbb{B}^n$ is assigned the value $F(x)$ of the objective function, and the empirical mean of these values provides an estimate of $\langle\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})| \, H_P \, |\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle$. This estimate is then used in a classical optimization loop to update the parameters $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ with the goal of minimizing the empirical mean.

Generalities of Reduced QAOA. The action of G on the set $\mathbb{B}^n$ extends to a linear action on the Hilbert space W. Therefore, W is a *representation* of G and can be expressed as a direct sum of isotypic components with respect to this action:

$$W = \bigoplus_i W_i, \tag{6}$$

where the index i enumerates nonisomorphic irreducible representations of G appearing in W. In particular,

$$W^G := \{w \in W \mid g \cdot w = w \ \forall g \in G\} \tag{7}$$

will denote the trivial subrepresentation (the subspace of G-invariants). The actions of G and H_P on W commute:

$$H_P(g(w)) = g(H_P(w)) \quad \forall w \in W, \quad \forall g \in G. \tag{8}$$

Remark II.2. While the dimension of W^G is equal to the number of elements in $\mathbb{B}_G$ (see (2)) there might not be a natural (standard) basis of W^G encoded with the elements for $\mathbb{B}_G$.

However, in certain cases (as described in the next section), we can systematically select a representative $x_j \in \mathbb{B}^n$ for each G-orbit $\mathcal{O}_j \in \mathbb{B}_G$, allowing us to construct a natural reduced Hilbert space with the basis $\{|x_j\rangle\}_{\mathcal{O}_j \in \mathbb{B}_G}$.

Reduced QAOA for Simultaneous Bit-Flip Symmetry.

A convenient representative of each orbit is chosen by fixing the value of a particular bit b_j with $1 \le j \le n$. Recall that the full Hilbert space is $W = (\mathbb{C}^2)^{\otimes n}$, with computational basis vectors indexed by bit strings $x \in \{0,1\}^n$. Fixing the value of the j^{th} bit to either 0 or 1 restricts us to those basis states satisfying $x_j = b_j$. Equivalently, the reduced Hilbert space

$$W_j = \mathbb{C}^2 \otimes \cdots \otimes \underbrace{\mathbb{C}|b_j\rangle}_{j\text{-th position}} \otimes \cdots \otimes \mathbb{C}^2 \simeq (\mathbb{C}^2)^{\otimes(n-1)}. \tag{9}$$

is obtained by contracting the j^{th} tensor factor of W into the one-dimensional coordinate subspace corresponding to the chosen value, 0 or 1. Without loss of generality, we assume the fixed bit is 0.

The reduced problem Hamiltonian H_P^j is obtained directly from the original problem Hamiltonian by replacing the Pauli operator Z_j corresponding to the fixed bit with the identity. This yields a reduced QAOA defined on W_j with problem Hamiltonian H_P^j.

From Standard to Reduced DLAs. Recall that the *standard dynamical Lie algebra* (DLA) associated with a given QAOA instance is defined as the real Lie algebra generated, under the commutator operation, by the skew-Hermitian operators iH_M and iH_P, where H_M and H_P denote the mixer and problem Hamiltonians, respectively. In other words,

$$\mathfrak{g}_{\text{std}} := \langle\, iH_M,\, iH_P \,\rangle_{\text{Lie}}. \tag{10}$$

For comparison, we also consider the corresponding *free dynamical Lie algebra*, defined as the real Lie algebra generated by the individual local terms appearing in the mixer and problem Hamiltonians, iH_M and iH_P. Writing

$$H_M = \sum_j M_j, \qquad H_P = \sum_k P_k,$$

as decompositions into local Hermitian operators, the free DLA is given by

$$\mathfrak{g}_{\text{free}} := \left\langle \{\, iM_j \,\},\ \{\, iP_k \,\} \right\rangle_{\text{Lie}}. \tag{11}$$

We now turn to the central objects of this section: the *reduced dynamical Lie algebras*. Given a combinatorial optimization problem (COP) on n binary variables $b_0, \ldots, b_{n-1}$, we consider the family of reduced problems obtained by fixing one variable at a time. For each $\ell \in \{0, \ldots, n-1\}$, setting $b_\ell = 0$ (or equivalently $b_\ell = 1$, up to relabeling) yields a reduced COP on $n-1$ variables, which in turn defines a reduced QAOA instance with corresponding mixer and problem Hamiltonians, denoted H_M^ℓ and H_P^ℓ.

Definition II.3. The *standard reduced dynamical Lie algebra* associated with the ℓ^{th} reduced problem is defined as

$$\mathfrak{g}_{\text{std}}^\ell := \left\langle\, iH_M^\ell,\, iH_P^\ell \right\rangle_{\text{Lie}}.$$

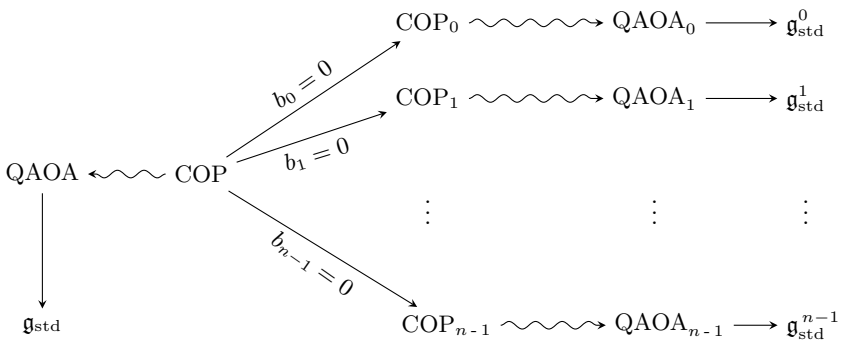


FIG. 3. **Reduction scheme for dynamical Lie algebras.** An n-bit combinatorial optimization problem gives rise to a QAOA instance with standard dynamical Lie algebra $\mathfrak{g}_{\text{std}}$. Fixing a variable b_ℓ produces a reduced optimization problem COP_ℓ, a corresponding reduced QAOA instance, and an associated standard reduced DLA $\mathfrak{g}^\ell_{\text{std}}$.

Figure 3 summarizes these relationships schematically. The original COP gives rise to a single QAOA instance with standard DLA $\mathfrak{g}_{\text{std}}$, while each reduced COP produces a corresponding reduced QAOA instance and standard reduced DLA. The arrows indicate the passage from optimization problems to QAOA Hamiltonians and, subsequently, to their associated dynamical Lie algebras.

Reduced DLAs for the MaxCut Problem. To ground this abstract symmetry reduction in a concrete quantum optimization setting, we focus henceforth on the graph MaxCut problem—a benchmark QAOA instance whose structural invariance under global bit-flips makes it a prime candidate for this framework.

Let Γ be an undirected graph with vertex set $V = \{1, 2, \ldots, n\}$. The *MaxCut* problem asks for a partition of V into two disjoint sets such that the number of edges crossing the partition is maximized.

We encode a cut by a binary string $(x_1, \ldots, x_n) \in \mathbb{B}^n$, where $x_i = 0$ or 1 indicates the side of the cut to which vertex i belongs. An edge $(i, j) \in E$ contributes 1 to the cut if and only if its endpoints lie on opposite sides, and the corresponding objective function is

$$F(x_1, \ldots, x_n) = \sum_{(i,j)\in E} \left[(1 - x_i)x_j + x_i(1 - x_j)\right], \tag{12}$$

which evaluates to 1 precisely when $x_i \neq x_j$, and to 0 otherwise. Thus, $F(x_1, \ldots, x_n)$ counts the total number of edges crossing the cut defined by $(x_1, \ldots, x_n)$. The goal of the MaxCut problem is therefore to find a binary string $(x_1, \ldots, x_n) \in \mathbb{B}^n$ that maximizes the objective function. The function F is invariant under the natural $\mathbb{Z}_2$-action that flips all bits simultaneously,

$$(x_1, \ldots, x_n) \longmapsto (1 - x_1, \ldots, 1 - x_n), \tag{13}$$

To eliminate this global symmetry, we fix one vertex $s \in \{1, \ldots, n\}$ and impose the condition $x_s = 0$. This selects a unique representative from each $\mathbb{Z}_2$-orbit. For any representative

$$(x_1, \ldots, x_{s-1}, 0, x_{s+1}, \ldots, x_n) \in \mathbb{B}_G,$$

the reduced objective function becomes

$$\widetilde{F}(x_1, \ldots, x_{s-1}, 0, x_{s+1}, \ldots, x_n) = \sum_{\substack{(i,j)\in E \\ i,j\neq s}} \Big[(1 - x_i)x_j + x_i(1 - x_j)\Big] + \sum_{(s,j)\in E} x_j. \tag{14}$$

The corresponding (standard) mixer and problem Hamiltonians are

$$H_M = \sum_{w\in V} X_w, \qquad H_P = \sum_{(v,w)\in E} Z_v Z_w. \tag{15}$$

Then the *free* dynamical Lie algebra is

$$\mathfrak{g}_{\Gamma,\text{free}} = \Big\langle \{iX_w \mid w \in V\},\ \{iZ_wZ_{w'} \mid (ww') \in E\} \Big\rangle_{\text{Lie}}, \tag{16}$$

while the *standard* dynamical Lie algebra is

$$\mathfrak{g}_{\Gamma,\text{std}} = \Big\langle X, Z \Big\rangle_{\text{Lie}}, \tag{17}$$

with $X := iH_M$ and $Z := iH_P$. If the bit associated with a vertex $v \in V$ is fixed, the reduced mixer and problem Hamiltonians become

$$H_M^v = \sum_{w \in V \setminus \{v\}} X_w, \qquad H_P^v = \sum_{\substack{(w,w') \in E \\ w,w' \neq v}} Z_wZ_{w'} + \sum_{(v,w) \in E} Z_w. \tag{18}$$

Consequently, the reduced *free* dynamical Lie algebra is

$$\mathfrak{g}^v_{\Gamma,\text{free}} = \Big\langle \{iX_w \mid w \in V \setminus \{v\}\},\ \{iZ_wZ'_w \mid (ww') \in E,\ w,w' \neq v\},\ \{iZ_w \mid (vw) \in E\} \Big\rangle_{\text{Lie}}, \tag{19}$$

while the reduced *standard* dynamical Lie algebra is the algebra

$$\mathfrak{g}^v_{\Gamma,\text{std}} = \Big\langle \mathcal{X}_{\widehat{v}},\ \mathcal{Z}_{\widehat{v}} \Big\rangle_{\text{Lie}}, \tag{20}$$

with generators

$$\mathcal{X}_{\widehat{v}} := iH_M^v \text{ and } \mathcal{Z}_{\widehat{v}} := iH_P^v. \tag{21}$$

Let Γ_v denote the subgraph obtained from Γ by removing the distinguished vertex v and all its incident edges. The free reduced DLAs for the MaxCut problem typically yield the maximal possible algebra; with only a few exceptions determined by the topology of Γ_v, they are isomorphic to the full special unitary algebra on the remaining qubits.

Theorem II.4. Let (Γ, v) be a pair consisting of a graph Γ and a chosen vertex v. Suppose that the reduced graph Γ_v is a connected graph that is not bipartite. Then, the free reduced DLA $\mathfrak{g}^v_{\Gamma,\text{free}}$ is isomorphic to the special unitary Lie algebra on the reduced Hilbert space:

$$\mathfrak{g}^v_{\Gamma,\text{free}} \simeq \mathfrak{su}(W_v) \simeq \mathfrak{su}(2^{n-1}). \tag{22}$$

Unifying Free and Standard Reduced DLAs via Graph Extension. The following result provides a powerful structural insight. It shows that, by extending a given graph in a controlled manner, the free dynamical Lie algebra associated with the reduction of the extended graph admits a natural embedding into the standard reduced dynamical Lie algebra of a suitably enlarged graph. This result is essentially a consequence of Theorem IV.2.

Theorem II.5. Let Γ be a connected graph on n vertices and let $v \in V(\Gamma)$ be any vertex. Denote by

$$j(v) := \max_{w \in V(\Gamma)} \text{dist}_\Gamma(v, w) \tag{23}$$

the eccentricity of v in Γ. Then there exists a graph $\widehat{\Gamma}_v$ satisfying the following properties:

- *Containment.* The graph Γ is a subgraph of $\widehat{\Gamma}_v$.
- *Vertex bound.* The number of vertices satisfies

$$|V(\widehat{\Gamma}_v)| < 1 + 3(j(v) - 1)(n - 1) + \frac{(3n-2)(3n-3)}{2}. \tag{24}$$

- *Edge bound.* The number of edges satisfies

$$|E(\widehat{\Gamma}_v)| < |E(\Gamma)| + 4(n-1) + 3(j(v)-1)(n-1) + \frac{(3n-2)(3n-3)}{2}. \tag{25}$$

- *Coincidence of standard and free DLAs.* The standard and free reduced dynamical Lie algebras of the extended graph coincide; namely,

$$\mathfrak{g}^{v}_{\widehat{\Gamma}_v,\mathrm{std}} = \mathfrak{g}^{v}_{\widehat{\Gamma}_v,\mathrm{free}}, \qquad \text{cf. (45)}. \tag{26}$$

- *Optimization equivalence.* Optimal solutions of the MaxCut problem on Γ are obtained from optimal solutions on $\widehat{\Gamma}_v$ by discarding the bits corresponding to the auxiliary vertices $V(\widehat{\Gamma}_v) \setminus V(\Gamma)$.

Remark II.6. Since $j(v) \leq n-1$ for any connected graph Γ on n vertices, Theorem II.5 implies the bounds

$$|V(\widehat{\Gamma}_v)| = \mathcal{O}(n^2), \qquad |E(\widehat{\Gamma}_v)| = \mathcal{O}(n^2).$$

In particular, the number of edges in the extended graph $\widehat{\Gamma}_v$ is at most quadratic in the number of vertices of the original graph Γ. Thus, the graph extension incurs at most a quadratic overhead in system size.

Expressivity and Pure-State Controllability. The equality of DLAs in Eq. (26) has important implications for pure-state controllability, which we now elaborate on.

If the reduction of the extension graph $\widehat{\Gamma}_v$ at v (meaning the graph obtained after removal of v and its adjacent edges) is a connected graph, and the pair $(\widehat{\Gamma}_v, v)$ satisfies $\mathfrak{g}_{\widehat{\Gamma}_v,\mathrm{free}} \simeq \mathfrak{su}(\widehat{W}_v)$, which holds for all connected graphs except paths, stars, and even cycles (see Theorem II.4), the standard reducedQAOA ansatz attains maximal expressivity within the reduced Hilbert space $\widehat{W}_v$ (see Eq. (9)). Specifically, for any pair of normalized states $|t\rangle, |h\rangle \in \widehat{W}_v$, there exists a depth $L \in \mathbb{Z}_{>0}$ and parameters $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ such that the L-layer QAOA circuit $U(\boldsymbol{\beta}, \boldsymbol{\gamma})$ (defined in Eq. (5)) achieves exact state preparation up to a global phase ϕ:

$$U(\boldsymbol{\beta}, \boldsymbol{\gamma}) |h\rangle = e^{i\phi} |t\rangle . \tag{27}$$

In particular, choosing $|t\rangle$ to be a ground state of the problem Hamiltonian $\widehat{H}_P$ guarantees finding the exact optimal solution with 100% probability.

Remark II.7. In the language of quantum control theory, this maximal expressivity over the state space is known as pure-state controllability.

We summarize these findings in the following theorem.

Theorem II.8. If the reduction of the extension graph $\widehat{\Gamma}_v$ at vertex v is connected and is not a path, a star, or an even cycle, then the reduced Quantum Approximate Optimization Algorithm, QAOA_v, for the MaxCut problem on the extended graph $\widehat{\Gamma}_v$ exhibits pure-state controllability.

Optimization Landscapes and Barren Plateaus. Another important consequence of our results concerns the structure of the optimization landscape associated with variational quantum algorithms. In particular, the dimension and representation-theoretic structure of the underlying dynamical Lie algebra are known to be tightly linked to the possible emergence of barren plateaus, regions of exponentially suppressed gradients in the parameter landscape [7, 15, 16, 18, 20, 22].

The barren plateau phenomenon refers to the exponential suppression of gradients of the loss function with respect to circuit parameters, rendering classical training methods ineffective [18, 20]. For sufficiently deep variational circuits, this trainability bottleneck can be quantified by analyzing the algebraic properties of the reduced system. In particular, we have the following result.

Theorem II.9. Let $\Gamma = (V, E)$ be a graph and $v \in V$ a distinguished vertex such that the standard reducedDLA coincides with the reduced free DLA, i.e., $\mathfrak{g}^{v}_{\Gamma,\mathrm{std}} = \mathfrak{g}^{v}_{\Gamma,\mathrm{free}}$. If the induced subgraph Γ_v satisfies the topological conditions of Theorem II.4, then $\dim(\mathfrak{g}^{v}_{\Gamma,\mathrm{std}}) = 4^{n-1} - 1$, and the variance of the loss function

over the reduced uniform superposition state ρ_v with the normalized reduced problem Hamiltonian $\widehat{H}_P^v$ (see (47) and (48)) is bounded by:

$$\mathrm{Var}\big[\ell_{\boldsymbol{\beta},\boldsymbol{\gamma}}(\rho_v, \widehat{H}_P^v)\big] \;\leq\; \frac{2^{n-1}}{4^{n-1}-1}. \tag{28}$$

In particular, the variance decays exponentially in n, implying the presence of a barren plateau.

Action on Hilbert Spaces. We analyze the structure of the full Hilbert space W, as well as the reduced spaces W_v, under the action of their respective free dynamical Lie algebras.

Theorem II.10. Let Γ be a connected graph on n vertices. The n-qubit Hilbert space $W = (\mathbb{C}^2)^{\otimes n}$ admits the orthogonal decomposition

$$W = W_{\mathrm{even}} \oplus W_{\mathrm{odd}}, \tag{29}$$

where W_{even} (respectively, W_{odd}) is the subspace spanned by the Hadamard basis states $\{|\pm\rangle^{\otimes n}\}$ containing an even (respectively, odd) number of $|-\rangle$ tensor factors. Moreover,

$$\dim(W_{\mathrm{even}}) = \dim(W_{\mathrm{odd}}) = 2^{n-1}. \tag{30}$$

These subspaces are precisely the two irreducible invariant subspaces of the free dynamical Lie algebra $\mathfrak{g}_{\Gamma,\mathrm{free}}$.

Remark II.11. The representations W_{even} and W_{odd} correspond to the trivial and sign representations of the $\mathbb{Z}_2$ symmetry group generated by the global parity operator $X^{\otimes n} \equiv X_1 \cdots X_n$. Because this global bit-flip operator commutes with the DLA generators, the Hilbert space naturally splits into the $+1$ and -1 eigenspaces of $X^{\otimes n}$.

Remark II.12. The decomposition (29) into two irreducible invariant subspaces was previously established in Ref. [13] for *archetypal* graphs, namely connected graphs that are neither bipartite nor cyclic. Theorem II.10 extends this structural result to all connected graphs.

Theorem II.10 admits the following analogue for reduced dynamical Lie algebras. Unlike the full Hilbert space, the reduced Hilbert space is irreducible under the action of the free reduced dynamical Lie algebra.

Theorem II.13. Let Γ be a connected graph and let v be an arbitrary vertex of Γ. The reduced Hilbert space W_v is an irreducible representation of the free reduced dynamical Lie algebra $\mathfrak{g}^v_{\Gamma,\mathrm{free}}$.

Connection to the Conjecture on Natural DLAs. We conclude this part by placing our results in the context of a central conjecture proposed in [13] (Conjecture 1 therein). In that work, the authors introduced the notion of the *natural dynamical Lie algebra*, designed to respect the intrinsic symmetries of the MaxCut problem. Concretely, this algebra is defined as the subalgebra of the free dynamical Lie algebra that commutes with both the action of the graph automorphism group and the global bit-flip symmetry on the Hilbert space W. A precise definition may be found in Section VI of [13].

For our purposes, the key structural fact is that the natural dynamical Lie algebra sits between the standard and free DLAs:

$$\mathfrak{g}_{\Gamma,\mathrm{std}} \subseteq \mathfrak{g}_{\Gamma,\mathrm{nat}} \subseteq \mathfrak{g}_{\Gamma,\mathrm{free}}. \tag{31}$$

Definition II.14. Let V be a representation space of a Lie algebra $\mathfrak{g}$, and let $|\xi\rangle \in V$. The *cyclic representation* $V_{\mathfrak{g},|\xi\rangle}$ generated by $|\xi\rangle$ is the smallest $\mathfrak{g}$-invariant subspace of V containing $|\xi\rangle$. Equivalently,

$$V_{\mathfrak{g},|\xi\rangle} = \mathrm{span}\left\{\rho(x_1)\rho(x_2)\cdots\rho(x_k)\,|\xi\rangle \mid x_1,\ldots,x_k \in \mathfrak{g},\ k \geq 0\right\},$$

where $\rho : \mathfrak{g} \to \mathrm{End}(V)$ denotes the corresponding Lie algebra representation.

Remark II.15. The notion of a cyclic representation is closely related to that of a Krylov subspace in numerical linear algebra. Given a single linear operator $A \in \mathrm{End}(V)$ and a vector $v \in V$, the corresponding Krylov subspace is $\mathcal{K}(A, v) := \mathrm{span}\{v, Av, A^2v, \ldots\}$. In contrast, the cyclic representation $V_{\mathfrak{g},|\xi\rangle}$ is generated by the action of an entire Lie algebra rather than a single operator. Thus, it may be viewed as the natural Lie-algebraic generalization of a Krylov subspace, obtained by allowing arbitrary finite products of the representation operators $\rho(x)$, where $x \in \mathfrak{g}$.

Let

$$|\xi_n\rangle := \underbrace{|+ + \cdots +\rangle}_{n} \tag{32}$$

denotes the uniform superposition over the standard basis states indexed by $\mathbb{B}^n$. Denote by m and M the respective dimensions of the cyclic subrepresentations of W generated by $|\xi_n\rangle$ under the actions of $\mathfrak{g}_{\Gamma,\mathrm{std}}$ and $\mathfrak{g}_{\Gamma,\mathrm{nat}}$.

Remark II.16. One motivation for focusing on these subspaces is rooted in the structure of the Quantum Approximate Optimization Algorithm. The state $|\xi_n\rangle$ is the standard choice of initial state for QAOA, and the unitaries appearing in the QAOA circuit are generated by elements of $\mathfrak{g}_{\Gamma,\mathrm{std}}$. Consequently, the entire QAOA evolution starting from $|\xi\rangle$ remains confined to the cyclic subrepresentation up to the final measurement.

In other words, this subspace captures the effective portion of the Hilbert space that is dynamically accessible to QAOA when initialized in the uniform superposition state. Knowing its dimension therefore provides insight into the true degrees of freedom governing the algorithm's evolution, as opposed to the full ambient Hilbert space.

The conjecture proposed in [13] can be reformulated as follows.

Conjecture II.17. For most graphs Γ on n vertices, the quantity

$$\Delta := M - m \tag{33}$$

is bounded by a polynomial in n. Equivalently, the difference between the dimensions of the cyclic representations generated by $|\xi_n\rangle$ for $\mathfrak{g}_{\Gamma,\mathrm{std}}$ and $\mathfrak{g}_{\Gamma,\mathrm{nat}}$ grow at most polynomially with the system size.

While a reduced analog of the natural dynamical Lie algebra has not yet been formally defined, a natural candidate would be the (largest with respect to inclusion) subalgebra of $\mathfrak{g}^v_{\Gamma,\mathrm{free}}$ commuting with the action of the automorphism group of the reduced graph Γ_v. This yields an algebra squeezed between the reduced DLAs $\mathfrak{g}^v_{\Gamma,\mathrm{std}}$ and $\mathfrak{g}^v_{\Gamma,\mathrm{free}}$.

The results of the present work imply that an analogue of Conjecture II.17 holds for the reduced framework associated with the extension graphs constructed in Theorem II.5. Indeed, Theorem II.5 shows that, for a suitable extension $\widehat{\Gamma}_v$, one has the equality of dynamical Lie algebras $\mathfrak{g}^v_{\widehat{\Gamma}_v,\mathrm{std}} = \mathfrak{g}^v_{\widehat{\Gamma}_v,\mathrm{free}}$, while Theorem II.13 establishes that the reduced Hilbert space W_v is an irreducible representation of $\mathfrak{g}^v_{\Gamma,\mathrm{free}}$. Consequently, the cyclic representation generated by the reduced initial state $|\xi_v\rangle$ (see (46)) coincides with the entire reduced Hilbert space W_v for both the standard and free reduced dynamical Lie algebras.

Therefore, the reduced analogue of Δ vanishes for the family of extension graphs $\widehat{\Gamma}_v$ constructed in Theorem II.5. In particular, the analogue of Conjecture II.17 holds for the reduced framework considered in this work, associated with the natural extension of any graph Γ and distinguished vertex v.

Exponential vs. Quadratic Growth: The k-Armed Spider Graphs. We present a family of graphs for which the dimension of the standard dynamical Lie algebra grows exponentially in the number of vertices, while there exists a distinguished vertex for which the dimension of the corresponding standard *reduced* DLA is only quadratic in the number of vertices. This provides a concrete and explicit separation between the full and reduced settings.

The family under consideration consists of the so-called k-armed spider graphs, denoted $\mathcal{O}_{m_1,\ldots,m_k}$, which were studied in [19]. Such a graph is obtained by taking a central vertex v and attaching to it k path graphs $P_{m_1},\ldots,P_{m_k}$, each joined to v at one of its endpoints. In other words, $\mathcal{O}_{m_1,\ldots,m_k}$ consists of k disjoint paths whose initial vertices are identified with the central vertex v (see Figure 4).

The reduction of $\mathcal{O}_{m_1,\ldots,m_k}$ at the central vertex v removes v and disconnects the graph into a disjoint union of the path graphs $P_{m_1} \sqcup \cdots \sqcup P_{m_k}$:

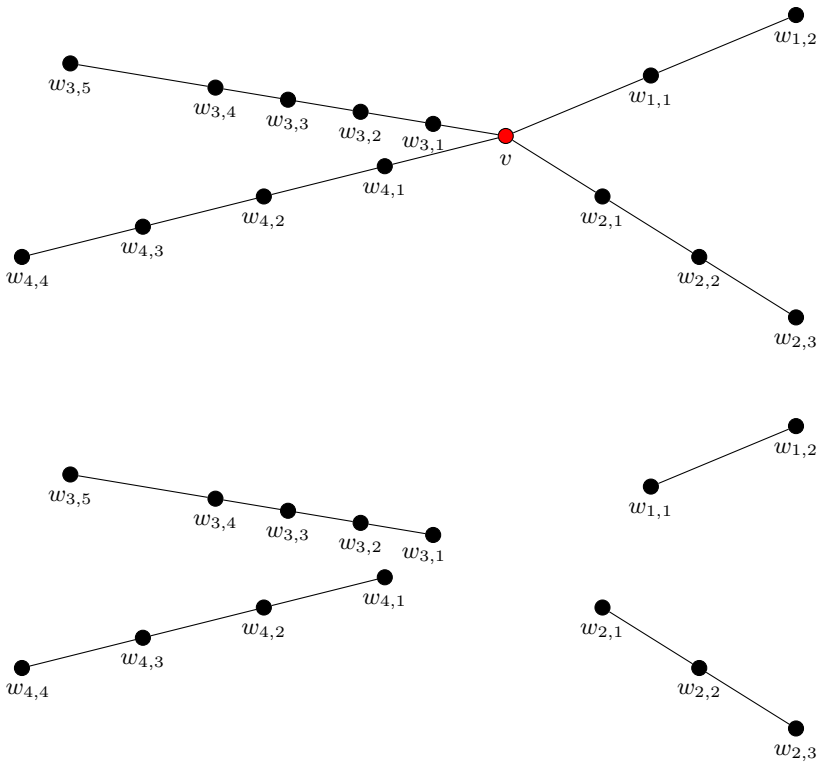


FIG. 4. Top: the 4-armed spider graph $\mathcal{O}_{2,3,4,5}$ with central vertex v and 4 path graphs P_2, P_3, P_4 and P_5 attached. Bottom: the reduced graph $\mathcal{O}^v_{2,3,4,5}$ obtained by removing the central vertex v, resulting in a disjoint union of the 4 path graphs.

Assume now that the lengths $m_1, \ldots, m_k$ are pairwise distinct. It was established in [19, Corollary 26] that in this case the standard and free DLAs of the full spider graph coincide, namely $\mathfrak{g}_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{std}} = \mathfrak{g}_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{free}}$. Consequently, Theorem 1 of [13] allows to conclude that $\dim(\mathfrak{g}_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{std}})$ is exponential in n.

We show that the analogous statement holds for the reduced DLAs at the central vertex v:

$$\mathfrak{g}^v_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{std}} = \mathfrak{g}^v_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{free}} = \mathfrak{g}_{P_{m_1},\mathrm{free}} \oplus \cdots \oplus \mathfrak{g}_{P_{m_k},\mathrm{free}}. \tag{34}$$

Thus, the reduced algebra decomposes as a direct sum of the free DLAs associated with the individual arms and has dimension

$$\dim\left(\mathfrak{g}^v_{\mathcal{O}_{m_1,\ldots,m_k},\mathrm{std}}\right) = 2\sum_{j=1}^{k} m_j^2 + n \leq 2\left(\sum_{j=1}^{k} m_j\right)^2 + n = 2n^2 + n. \tag{35}$$

This establishes a quadratic upper bound in the system size.

In contrast, for the full (non-reduced) spider graph, the standard DLA grows exponentially with the number of vertices, as shown in [19]. The k-armed spider graphs therefore provide an explicit family in which symmetry reduction at a carefully chosen vertex dramatically alters the algebraic complexity: from exponential growth in the full system to quadratic growth in the reduced one.

Empirical Tests of DLA Dimension Reduction Across Asymmetric Graphs. We perform a numerical study of the dimensions of the standard and reduced dynamical Lie algebras for asymmetric graphs. Across all instances examined, we observe that symmetry reduction can lead to a substantial decrease in the dimension of the associated dynamical Lie algebra. These findings motivate the conjecture that for any asymmetric graph, there exists a choice of reduction vertex for which the dimension of the corresponding standard reduced dynamical Lie algebra is strictly smaller than that of the original DLA.

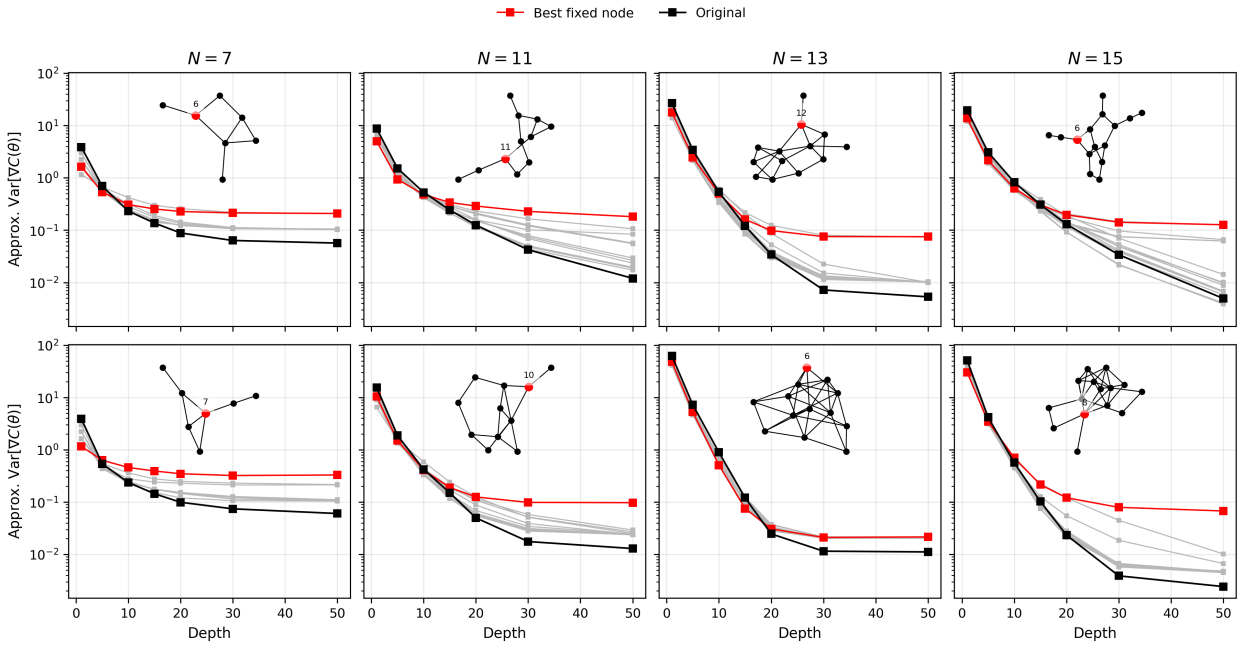


FIG. 5. Approximate variance of the loss-function gradient for selected asymmetric graphs. Each column shows two representative examples for a given number of graph vertices. Black data points correspond to the approximate variance in the original graph instance as a function of QAOA depth. Gray data points correspond to the approximate variance for a given reduced QAOA instance as a function of QAOA depth, with red data points corresponding to the approximate variance for the reduced QAOA instance that results in the largest approximate variance at depth 50. For each depth, the approximate variance is computed over 400 parameter initializations. Additionally, each subplot contains the rendered original graph instance, where the vertex that results in the largest approximate variance at depth 50 is labeled and highlighted in red.

Conjecture II.18. Let Γ be an asymmetric graph. Then there exists a vertex $v \in V(\Gamma)$ such that the dimension of the corresponding standard reduced dynamical Lie algebra is strictly smaller than that of the full standard DLA:

$$\dim(\mathfrak{g}^{v}_{\Gamma,\mathrm{std}}) \;<\; \dim(\mathfrak{g}_{\Gamma,\mathrm{std}}).$$

This conjecture has been verified for all asymmetric graphs on 6 and 7 vertices.

Determining the exact dimensions of standard and reduced dynamical Lie algebras for QAOA applied to MaxCut on large asymmetric graphs quickly becomes computationally intractable. To probe larger systems, we therefore exploit the established relationship between the dimension of the dynamical Lie algebra and the variance of the QAOA loss function. As shown in [22], the variance is inversely related to the DLA dimension; in particular, an increased variance is indicative of a smaller dynamical Lie algebra.

Using this proxy, we compute the loss-function variance for both standard and reduced QAOA circuits on a collection of connected asymmetric graphs with 7–15 vertices. In all cases studied, the reduced QAOA exhibits a consistently higher variance than the corresponding unreduced instance, providing further numerical support for Conjecture II.18, as evidenced in Figure 5.

Reduced DLAs for Grover-mixer QAOA. In this subsection, we compare the reduced and standard dynamical Lie algebras arising in the variant of QAOA that employs the *Grover mixer*, defined as the projector onto the initial state:

$$G_M \;:=\; |\xi\rangle\langle\xi| \;=\; H^{\otimes n}\,|0\cdots0\rangle\langle 0\cdots0|\,H^{\otimes n}, \tag{36}$$

where $|\xi\rangle = |+\cdots+\rangle$ is taken to be the uniform superposition of all computational basis states.

We briefly recall the setup and refer the reader to [2] and [25] for a more detailed discussion. We consider a general combinatorial optimization problem of minimizing an objective function F of the form given in (1).

The variant of QAOA, employing the mixer Hamiltonian G_M from (36) was introduced in [2] and is known as the *Grover-mixer QAOA* (GM-QAOA).

Notice that the range of the reduced objective function $\widetilde{F}(x)$ coincides with the range of $F(x)$ (each value appears with multiplicity reduced by a factor of two). Consequently, by Theorem III.1 of [25], the dynamical Lie algebra $\mathfrak{g}_\xi$ generated by the GM-QAOA Hamiltonians, together with all of its bit-reduced versions $\mathfrak{g}_\xi^b$ (obtained by fixing any bit b to the value 0 or 1), are all isomorphic to

$$\mathfrak{su}(r) \oplus \mathfrak{u}(1) \oplus \mathfrak{u}(1), \tag{37}$$

where r denotes the number of distinct objective values attained by $F(x)$ and $\widetilde{F}(x)$.

We now look at the situation for QAOA implementing Grover search as in [11]. Consider the case of searching for a single marked element $\omega \in \mathbb{B}^n$. The problem Hamiltonian is the rank-one projector

$$H_P = |\omega\rangle \langle\omega| ,$$

while the mixer H_M is chosen to be the standard X-mixer $B = \sum_{j=1}^{n} X_j$. These generate the *standard dynamical Lie algebra*

$$\mathfrak{g}_{\text{std}} := \langle iB,\ iH_P \rangle_{\text{Lie}}.$$

By Corollary V.3 of [25], this algebra is isomorphic to

$$\mathfrak{su}(n+1) \oplus \mathfrak{u}(1) \oplus \mathfrak{u}(1), \tag{38}$$

while all corresponding standard reduced algebras $\mathfrak{g}_{\text{std}}^b$ (obtained by restricting the bit b to a fixed bit value) satisfy

$$\mathfrak{g}_{\text{std}}^b \ \cong\ \mathfrak{su}(n) \oplus \mathfrak{u}(1) \oplus \mathfrak{u}(1).$$

Remark II.19. The above results highlight a distinctive feature of the Grover-mixer version of QAOA: its standard reduced dynamical Lie algebras are isomorphic to the original ones. In contrast, for QAOA with the standard X-mixer, bit reductions can alter the structure of the DLA in many different ways.

Supplementary Results. We study the dimensions of standard and reduced dynamical Lie algebras for several specific families of graphs. Notably, for star graphs $K_{1,n}$, we show that the standard reduced dynamical Lie algebra corresponding to the central vertex is isomorphic to the 3-dimensional Lie algebra $\mathfrak{su}(2)$ for all n. In contrast, numerical computations for $n \leq 11$ indicate that the dimension of the corresponding standard (unreduced) dynamical Lie algebra grows with n.

For cycle graphs C_n, as well as for the boundary vertices of path graphs P_n with $n \geq 5$, we find that the dimensions of the corresponding standard reduced dynamical Lie algebras exceed those of the standard dynamical Lie algebras.

III. DISCUSSION

We have shown that classical symmetry reduction can fundamentally alter the dynamical structure of QAOA. Under explicit and efficiently verifiable graph-theoretic conditions, the standard reduced dynamical Lie algebra coincides with the corresponding free reduced DLA, implying that the standard reduced ansatz attains the same expressive power as the unconstrained multi-angle variant on the reduced Hilbert space. Moreover, this regime is universal in the sense that, for any connected graph and any choice of reduction vertex, it can be enforced through an explicit graph extension requiring only $\mathcal{O}(|V|^2)$ additional vertices and edges, while preserving the MaxCut solutions of the original instance under a simple projection.

Our results show that symmetry reduction reshapes the dynamical properties of QAOA in multiple complementary ways. Compared to the standard framework, the reduced architecture achieves stronger controllability guarantees, admits graph extensions with quadratic vertex complexity rather than quadratic edge

complexity, and, for suitable graph families, exhibits a dramatic reduction in dynamical Lie algebra dimension from exponential to quadratic growth. More broadly, our numerical experiments on asymmetric graphs indicate that reductions producing smaller DLAs consistently exhibit larger gradient variance, suggesting that symmetry reduction can improve trainability without degrading optimization performance.

Taken together, these results identify reduced dynamical Lie algebras as a natural structural framework for understanding the interplay between expressivity and trainability in variational quantum algorithms. Since DLA dimension is closely connected to the variance of the QAOA loss landscape, our findings provide a principled explanation for when symmetry reduction is expected to mitigate barren plateaus. They also suggest a practical workflow in which inexpensive graph-theoretic diagnostics or DLA dimension estimates guide the selection of advantageous reduction vertices before quantum optimization is performed.

Several directions remain open. An important next step is the development of efficient classical algorithms for computing or estimating the dimensions and structural decompositions of reduced dynamical Lie algebras beyond the graph classes considered here. Another promising direction is to extend the present framework to other combinatorial optimization problems and symmetry classes, where reduction may similarly reshape the associated dynamical Lie algebras. More broadly, the constructive techniques developed in this work suggest that symmetry-aware modification of optimization instances can serve as a systematic design principle for tailoring the expressive and trainability properties of near-term variational quantum algorithms.

IV. METHODS

A key distinguishing feature of the standard dynamical Lie algebras arising from reduced QAOA, compared to their classical counterparts, is that the generator corresponding to the problem Hamiltonian is non-homogeneous. Specifically, it comprises a linear part,

$$\mathcal{Z}_{\widehat{v},1} := i \sum_{(vw)\in E(\Gamma)} Z_w,$$

supported on the immediate neighbors of the distinguished vertex v, alongside a degree-2 interaction part coming from the remaining edges (see (18)). As a first step, we show that these parts can be isolated separately within the standard reduced DLA $\mathfrak{g}^{v}_{\Gamma,\mathrm{std}}$.

Crucially, this structural decomposition allows us to prove that a specific, important class of distinguished elements is contained within these standard reduced dynamical Lie algebras. We begin by establishing the necessary notation.

Definition IV.1. Let v be a vertex in a graph Γ, and let $k \geq 1$ be an integer. We define the set of vertices at graph distance k from v as

$$\mathcal{N}_{v,k} := \{\, w \in V(\Gamma) : \mathrm{dist}(v,w) = k \,\}. \tag{39}$$

We further partition this neighborhood into subsets of vertices with even and odd degrees, respectively:

$$\mathcal{N}_{v,k,\mathrm{even}} := \{w \in \mathcal{N}_{v,k} \mid \deg(w) \equiv 0 \pmod 2\}, \tag{40}$$
$$\mathcal{N}_{v,k,\mathrm{odd}} := \{w \in \mathcal{N}_{v,k} \mid \deg(w) \equiv 1 \pmod 2\}. \tag{41}$$

For each $k \in \mathbb{Z}_{>0}$ and $v \in V(\Gamma)$, we define the corresponding distance-restricted Pauli-X operators:

$$\begin{aligned} \mathcal{X}_{v,k} &:= i \sum_{w\in\mathcal{N}_{v,k}} X_w, \\ \mathcal{X}_{v,k,\mathrm{even}} &:= i \sum_{w\in\mathcal{N}_{v,k,\mathrm{even}}} X_w, \\ \mathcal{X}_{v,k,\mathrm{odd}} &:= i \sum_{w\in\mathcal{N}_{v,k,\mathrm{odd}}} X_w. \end{aligned} \tag{42}$$

These operators represent the sums of Pauli-X operators supported on vertices at a fixed graph distance k from v, optionally restricted by vertex parity.

Theorem IV.2. Let $\Gamma = (V, E)$ be a connected graph and let $v \in V(\Gamma)$. Assume that for every layer $j \geq 1$ and every vertex $w \in \mathcal{N}_{v,j}$, the number of neighbors of w in the preceding layer $\mathcal{N}_{v,j-1}$ is odd, i.e.,

$$|\{\, u \in \mathcal{N}_{v,j-1} \mid (u, w) \in E \,\}| \equiv 1 \pmod 2. \tag{43}$$

Then, the standard reduced dynamical Lie algebra $\mathfrak{g}^v_{\Gamma,\mathrm{std}}$ contains the elements $\mathcal{X}_{\widehat{v},k}$, $\mathcal{X}_{v,k,\mathrm{even}}$, and $\mathcal{X}_{v,k,\mathrm{odd}}$ as defined in (42) for all $k \in \mathbb{Z}_{>0}$.

The key intuition behind the verification of this result is to initialize with the linear operator $\mathcal{Z}_{\widehat{v},1}$ and iteratively evaluate alternating commutators with the X and degree-2 Z generators. This procedure recursively "slices" the Lie algebra elements according to their distance from the distinguished vertex.

Remark IV.3. The conclusion of Theorem IV.2 extends to any vertex w satisfying $iX_w \in \mathfrak{g}^v_{\Gamma,\mathrm{std}}$ and for which the parity condition (43), centered at w, holds.

Combining Theorem IV.2 with Remark IV.3 provides a mechanism for establishing the presence of additional local Pauli-X generators in the reduced dynamical Lie algebra. In particular, for certain pairs (Γ, v), this allows one to prove that

$$iX_w \in \mathfrak{g}^v_{\Gamma,\mathrm{std}} \qquad \text{for every } w \in V(\Gamma) \setminus \{v\}.$$

By commuting these singleton X-generators with the degree-one Z-generator, and by taking nested commutators of pairs of singleton X-generators supported on edges of Γ with the degree-two Z-generators, one can straightforwardly obtain the remaining Z-generators of both the free reduced dynamical Lie algebra on Γ_v and the free reduced dynamical Lie algebra of the original graph Γ with distinguished vertex v.

Thus, we establish the containment of the free dynamical Lie algebra on the reduced graph Γ_v within the standard reduced dynamical Lie algebra on the original graph Γ:

$$\mathfrak{g}^v_{\Gamma,\mathrm{std}} \supseteq \mathfrak{g}_{\Gamma_v,\mathrm{free}}, \tag{44}$$

and equality between the standard and free reduced dynamical Lie algebras:

$$\mathfrak{g}^v_{\Gamma,\mathrm{std}} = \mathfrak{g}^v_{\Gamma,\mathrm{free}}. \tag{45}$$

The construction of an *extension graph* provides a concrete realization of this mechanism and stands as one of our core architectural results. Specifically, given a connected graph Γ and a vertex v, we construct its extension to satisfy the hypotheses of Theorem II.5 by appending path graphs to the original vertices. We precisely calibrate the lengths of these paths so that their terminal vertices lie at pairwise distinct distances from the distinguished vertex v within the extended graph.

By Theorem IV.2, this structural property ensures the existence of the corresponding distance-restricted X-elements for the newly added leaves in the dynamical Lie algebra. Furthermore, by Remark IV.3, this argument can be applied recursively at these leaves. Proceeding inductively from the leaves back toward the original graph Γ, we systematically recover the remaining X-type generators. We provide the full, detailed construction of the extension graph and verification of the properties listed in Theorem II.5.

Next we sketch the derivation of the variance bound stated in Theorem II.9.

For sufficiently large circuit depth p, the variance of the loss function asymptotically concentrates and can be evaluated analytically via the structural properties of the underlying dynamical Lie algebra (see [22]).

Let the initial state be the reduced uniform superposition state

$$|\xi_v\rangle := \frac{1}{\sqrt{2^{n-1}}} \sum_{s \in \mathbb{B}^n \cap W_v} |s\rangle \tag{46}$$

on $n-1$ qubits, denoted by

$$\rho_v := |\xi_v\rangle \langle \xi_v| , \tag{47}$$

and let the normalized reduced problem Hamiltonian be given by

$$\widehat{H}^v_P := \frac{1}{|\lambda_{\max}|} H^v_P, \tag{48}$$

where $\lambda_{\max}$ is the maximum eigenvalue of the reduced problem Hamiltonian (corresponding to the maximum value of the reduced objective function $\widetilde{F}$ in (14)).

Consider any pair (Γ, v) satisfying the algebraic identity $\mathfrak{g}^v_{\Gamma,\mathrm{std}} = \mathfrak{g}^v_{\Gamma,\mathrm{free}}$, such that the induced subgraph Γ_v avoids the exceptional topologies enlisted in Theorem II.4. Under these conditions, the standard reducedDLA attains its maximal scaling:

$$\mathfrak{g}^v_{\Gamma,\mathrm{std}} \simeq \mathfrak{su}(2^{n-1}), \tag{49}$$

which is a simple Lie algebra with dimension $\dim(\mathfrak{g}^v_{\Gamma,\mathrm{std}}) = 4^{n-1} - 1$.

To estimate the variance using the general framework of [22], we evaluate the respective DLA components. Because ρ_v is a pure state, its projection satisfies

$$\mathcal{P}_{\mathfrak{g}^v_{\Gamma,\mathrm{std}}}(\rho_v) = \mathrm{Tr}(\rho_v^2) = 1. \tag{50}$$

Furthermore, the projection of the normalized operator $\widehat{H}^v_P$ onto the special unitary algebra satisfies the general upper bound

$$\mathcal{P}_{\mathfrak{g}^v_{\Gamma,\mathrm{std}}}(\widehat{H}^v_P) \;\leq\; 2^{n-1}. \tag{51}$$

Substituting the dimension of $\mathfrak{su}(2^{n-1})$ and the individual projection evaluations (50) and (51) into the variance approximation formula, we establish the upper bound:

$$\mathrm{Var}\big[\ell_{\boldsymbol{\beta},\boldsymbol{\gamma}}(\rho_v, \widehat{H}^v_P)\big] \;\approx\; \frac{\mathcal{P}_{\mathfrak{g}^v_{\Gamma,\mathrm{std}}}(\rho_v)\,\mathcal{P}_{\mathfrak{g}^v_{\Gamma,\mathrm{std}}}(\widehat{H}^v_P)}{\dim(\mathfrak{g}^v_{\Gamma,\mathrm{std}})} \;\leq\; \frac{2^{n-1}}{4^{n-1}-1}. \tag{52}$$

Acknowledgments. We would like to thank Rui Mao for a careful reading of an earlier version of the manuscript and for identifying a counterexample to the corresponding version of Theorem IV.2. This research was supported in part by NSF Award No. 2427042.

[1] Jonathan Allcock, Miklos Santha, Pei Yuan, and Shengyu Zhang. On the dynamical Lie algebras of quantum approximate optimization algorithms. *Quantum*, 10:2119, May 2026. 2

[2] A. Bärtschi and S. Eidenbenz. Grover mixers for QAOA: Shifting complexity from mixer design to state preparation. *IEEE International Conference on Quantum Computing and Engineering*, pages 72–82, 2020. 13, 14

[3] K. Blekos, D. Brand, A. Ceschini, C. Chou, R. Li, K. Pandya, and A. Summer. A review on quantum approximate optimization algorithm and its variants. *Physics Reports*, 1068:1–66, 2024. 4

[4] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Physical review letters*, 125(26):260505, 2020. 2

[5] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and et al. Variational quantum algorithms. *Nat. Rev. Phys.*, 3(9):625–644, 2021. 2

[6] E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm. *arXiv:1411.4028*, 2014. 1, 2, 4

[7] E. Fontana, D. Herman, S. Chakrabarti, N. Kumar, R. Yalovetzky, J. Heredge, S. H. Sureshbabu, and M. Pistoia. Characterizing barren plateaus in quantum ansätze with the adjoint representation. *Nat. Commun.*, (15), 2024. 1, 9

[8] M. L. Goh, M. Larocca, L. Cincio, M. Cerezo, and F. Sauvage. Lie-algebraic classical simulations for quantum computing. *Physical Review Research*, 7(3):033266, 2025. 1

[9] S. Hadfield, Z. Wang, B. O'Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansätz. *Algorithms*, 12(34), 2019. 2, 4

[10] Dylan Herman, Cody Googin, Xiaoyuan Liu, Yue Sun, Alexey Galda, Ilya Safro, Marco Pistoia, and Yuri Alexeev. Quantum computing for finance. *Nature Reviews Physics*, 5(8):450–465, 2023. 2

[11] Z. Jiang, E. G. Rieffel, and Z. Wang. Near-optimal quantum circuit for Grover's unstructured search using a transverse field. *Physical Review A*, 95(6):062317, 2017. 14

[12] I. Joseph, Y. Shi, M. Porter, A. Castelli, V. Geyko, F. Graziani, S. Libby, and J. DuBois. Quantum computing for fusion energy science applications. *Physics of Plasmas*, 30(1), 2023. 2

[13] S. Kazi, M. Larocca, M. Farinati, P. Coles, M. Cerezo, and R. Zeier. Analyzing the quantum approximate optimization algorithm: ansätze, symmetries, and Lie algebras. *PRX Quantum*, 6, 2025. 2, 4, 10, 11, 12

[14] S. Kordonowy and H. Leipold. The lie algebra of XY-mixer topologies and warm starting qaoa for constrained optimization. *arXiv preprint arXiv:2505.18396*, 2025. 2

[15] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. Coles, and M. Cerezo. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6, 2022. 1, 9

[16] M. Larocca, N. Ju, D. García-Martín, P. Coles, and M. Cerezo. Theory of overparametrization in quantum neural networks. *arXiv:2109.11676*, 2021. 1, 9

[17] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, and Z. Holmes. Barren plateaus in variational quantum computing. *Nat. Rev. Phys.*, pages 1–16, 2025. 1

[18] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, and Z. Holmes. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, pages 1–16, 2025. 9

[19] R. Mao, P. Yuan, J. Allcock, and S. Zhang. Qaoa-maxcut has barren plateaus for almost all graphs. *arXiv:2512.24577*, 2026. 2, 4, 11, 12

[20] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):1–6, 2018. 9

[21] C. Outeiral, M. Strahm, J. Shi, G. Morris, S. Benjamin, and C. Deane. The prospects of quantum computing in computational molecular biology. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 11(1):e1481, 2021. 2

[22] M. Ragone, B. N. Bakalov, F. Sauvage, A. F. Kemper, C. Marrero, M. Larocca, and M. Cerezo. A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits. *Nat. Commun.*, 15, 2024. 1, 3, 9, 13, 16, 17

[23] R. Shaydulin, Stuart Hadfield, Tad Hogg, and Ilya Safro. Classical symmetries and the quantum approximate optimization algorithm. *Quantum Information Processing*, 20:1–28, 2021. 2

[24] R. Shaydulin, H. Ushijima-Mwesigwa, I. Safro, S. Mniszewski, and Y. Alexeev. Network community detection on small quantum computers. *Advanced Quantum Technologies*, 2(9):1900029, 2019. 2

[25] B. Tsvelikhovskiy, M. Nuyten, and B. N. Bakalov. Provable avoidance of barren plateaus for the quantum approximate optimization algorithm with grover mixers. *arXiv:2509.10424*, 2025. 2, 13, 14

[26] B. Tsvelikhovskiy, I. Safro, and Y. Alexeev. Symmetries and dimension reduction in quantum approximate optimization algorithm. *arXiv preprint arXiv:2309.13787*, 2023. 2

[27] B. Tsvelikhovskiy, I. Safro, and Y. Alexeev. Equivariant quantum approximate optimization algorithm. *IEEE Transactions on Quantum Engineering*, pages 1–14, 2026. 2

[28] H. Ushijima-Mwesigwa, R. Shaydulin, C. Negre, S. Mniszewski, Y. Alexeev, and I. Safro. Multilevel combinatorial optimization across quantum architectures. *ACM Transactions on Quantum Computing*, 2(1):1–29, 2021. 2

[29] M. T. West, J. Heredge, M. Sevior, and M. Usman. Provably trainable rotationally equivariant quantum machine learning. *PRX Quantum*, 5, 2024. 1

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementaryinformationproofs.pdf