

Optimization-free recursive Quantum Approximate Optimization Algorithm for the binary paint shop problem

Gary J Mooney ^{1,*} Jedwin Villanueva ¹ Bhaskar Roy Bardhan,² Joydip Ghosh ² Charles D Hill ^{1,3}
and Lloyd C L Hollenberg ^{1,†}

¹*School of Physics, University of Melbourne, Parkville 3010, VIC, Australia*

²*Research and Advanced Engineering, Ford Motor Company, Dearborn, Michigan 48124, USA*

³*School of Mathematics and Statistics, University of Melbourne, Parkville 3010, VIC, Australia*



(Received 2 September 2025; accepted 13 February 2026; published 24 March 2026)

The Quantum Approximate Optimization Algorithm (QAOA) is a leading candidate for near-term quantum advantage, yet its practical impact is hindered by limited performance on symmetric local Hamiltonians and the costly optimization of variational parameters. The recursive-QAOA (RQAOA) introduced by Bravyi *et al.* [*Phys. Rev. Lett.* **125**, 260505 (2020)] addresses the first limitation while also reducing circuit size, and parameter transfer techniques can be used to bypass the optimization loop effectively. In this work, we combine these two ideas to develop an optimization-free RQAOA and evaluate its performance on the binary paint shop problem (BPSP)—an optimization problem found in manufacturing where a sequence of cars must be painted under constraints while minimizing the number of color changes. The BPSP can be formulated as an Ising ground state problem with a symmetric local Hamiltonian in the form of MAX-CUT and properties well suited for the application of QAOA parameter transfer. We benchmark QAOA and RQAOA with parameter transfer against classical solvers and heuristics, and investigate their resilience to suboptimal parameters. For circuit optimization, we use reverse-causal cones and introduce a method of trimming outer two-qubit gates. To estimate the classical resources needed to simulate these quantum algorithms, we compute entanglement entropy and bond dimensions using matrix product state methods. We also compare circuit sizes and measurement counts across implementations. Our results show that RQAOA is inherently robust to parameter deviations, maintaining near-optimal solutions without noticeable degradation under parameter transfer while substantially reducing quantum resource requirements compared to QAOA. This highlights a viable route toward scalable quantum optimization without the overhead of the classical optimization loop and its challenges with barren plateaus.

DOI: [10.1103/xv79-w3rs](https://doi.org/10.1103/xv79-w3rs)

I. INTRODUCTION

With quantum computing technology rapidly advancing within the noisy intermediate-scale quantum (NISQ) era [1], an important research question receiving a lot of attention is what real-world problems can quantum computing be advantageous for in the near future. A promising candidate for this endeavour is optimization. Optimization problems are generally broadly applicable and often difficult to solve exactly, in many cases they can be efficiently mapped to quantum algorithms [2–5]. A common quantum algorithmic approach to optimization is the Quantum Approximate Optimization Algorithm (QAOA) [6–8]. It is a popular quantum-classical hybrid algorithm that can find good approximations to the ground states of Ising Hamiltonians that are encoding the

solutions to a quadratic unconstrained binary optimization (QUBO) problem instance. The QAOA circuits are relatively short compared to mainstay circuit model algorithms like Shor’s [9] and Grover’s [10] algorithms, making them well suited for analysis on NISQ computing. It remains to be seen whether general quantum optimization approaches, such as QAOA, can compete with classical heuristics that have been tailored over many decades to specific classes of optimization problems like the traveling salesman problem (TSP) [11–16]. The recursive Quantum Approximate Optimization Algorithm (RQAOA) introduced by Bravyi *et al.* [17] has shown potential. Motivated by alleviating inherent computational limits present in QAOA with symmetric local Hamiltonians, it produces outstanding improvements in solution qualities over baseline QAOA. In addition, it provides a trade-off between quantum circuit resources and number of circuits by reducing the circuit to reverse-causal cones (RCCs), which calculate $\langle ZZ \rangle$ correlations without needing to generate the full QAOA ansatz state. However, a noticeable drawback in RQAOA is that it uses QAOA as a subroutine, executing it for the initial problem graph and subsequent graphs generated by incrementally reducing its size until it is small enough to classically solve. Each application of QAOA involves optimizing its parameters, which often requires hundreds of circuits for even

*Contact author: mooneyg@unimelb.edu.au

†Contact author: lloydch@unimelb.edu.au

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

modest problem sizes and QAOA depths while also suffering from barren plateaus [18–20].

In this work, we demonstrate that combining RQAOA with precomputed QAOA parameters overcomes these challenges, achieving near-optimal performance while remaining robust to parameter noise. This robustness suggests that parameter tuning for RQAOA is scalable, providing strong solutions without extensive parameter optimization.

The binary paint shop problem (BPSP) is a combinatorial optimization challenge important both economically and environmentally to the automotive industry, with the goal to paint sequences of cars under certain constraints while minimizing the number of times the color needs to be swapped [21]. It has a simpler formulation than the general paint shop problem (PSP) used in actual car factories; however, it is still NP-complete and APX-hard, making it an appealing problem for foundational studies [22,23]. The BPSP instances can conveniently be mapped efficiently to symmetric Ising graphs that approach a regular structure as problem sizes increase. Specifically, the problem Ising graphs approach four-regular with unit-couplings (± 1 weights), which is equivalent to MAX-CUT [24] on a four-regular graph with ± 1 weighted edges, a ubiquitous and well-studied problem in theoretical computer science. Under a fixed-angle conjecture, this regularity is well suited for the use of precomputation techniques to find good estimates to the optimized QAOA parameters, typically well within 1% approximation ratio on equivalent MAX-CUT problems [25,26]. Previous work on BPSPs used precomputed parameters to calculate QAOA depth 1 and 2 expectations on problems of up to 100 qubits [4,27]. For graphs with irregular structure, the median optimal parameters calculated from many smaller instances can be transferred to unseen larger instances using the QAOAKit Python library [28,29].

Building on this motivation, we introduce and study RQAOA in combination with precomputed parameters as an approach to solve the BPSP. We benchmark the solution qualities obtained by precomputing and optimizing QAOA and RQAOA against classical heuristics and exact solvers, comparing both fixed angle and QAOAKit methods of obtaining precomputed parameters for QAOA and for each recursive step of RQAOA. We demonstrate the substantial improvement that RQAOA provides over QAOA, producing near-optimal solutions on generated BPSP instances up to 20 car bodies (mapped to 20 qubits). We find that RQAOA with precomputed parameters obtained from QAOAKit perform effectively as good as optimal. Precomputed fixed-angle parameters appears to still perform well even though the problem graph regularity is compromised during recursive reductions in RQAOA. Additionally, we observe that RQAOA exhibits a level of noise robustness in QAOA parameters, and that the typical displacements between QAOAKit and optimal parameters fall within this regime. This highlights RQAOA combined with precomputed parameters as a powerful yet scalable approach for BPSP and likely non-BPSP Ising graphs as well. Motivated by the question of algorithm scalability, we also investigate the classical and quantum resource requirements for RQAOA. In particular, we obtain quantum circuit metrics like CNOT counts, CNOT depths and qubit counts, as well as circuit counts and matrix product state (MPS) entanglement entropies and bond dimensions for varying problem




FIG. 1. A diagram of the QAOA and its quantum-classical hybrid nature. The ansatz state $|\beta, \gamma\rangle$ is prepared by running the ansatz circuit with an initial selection of parameters β and γ . The measurement results of the ansatz circuit are used to classically calculate the Hamiltonian H_{prob} energy corresponding to the problem encoding, which acts as the objective function for optimization. The energy calculation guides the adjustment of parameters β and γ based on the optimization strategy, leading to the preparation of an ansatz circuit. This process is repeated until the parameters β and γ have been optimized, at which point the corresponding ansatz state $|\beta, \gamma\rangle$ can be measured to provide approximate solutions to the target QUBO problem.

sizes and QAOA depths. We consider the circuits and their RCCs, and after compilation onto a 27-qubit IBM quantum device layout.

Shortly after this paper was uploaded to the arXiv, another study by Vijendran *et al.* [30] also applied RQAOA to the BPSP. They observed that as the problem size increases, the performance of RQAOA at $p = 1$ with optimized parameters gradually declines. Taken together with our finding that RQAOA appears robust to parameter noise, this suggests that while RQAOA with $p = 1$ performs strongly on BPSP instances at modest scales (around 1000 car bodies), understanding how to maintain this performance for increasingly larger systems remains an open and interesting question.

II. METHODS

A. The quantum approximate optimization algorithm

Here, we give a brief overview of QAOA. For deeper insight into the algorithm and how it is implemented, we include a lengthy derivation and further details in Appendix A.

The QAOA [6], shown in Fig. 1, is a form of variational quantum eigensolver (VQE) [31] based on the variational principle of quantum mechanics. It is a quantum-classical hybrid algorithm for approximating the ground state of a problem Ising Hamiltonian H_{prob} (where we will be using qubits in place of spins) that encodes solutions to a QUBO problem instance. A parametrized QAOA ansatz state is constructed from information about the problem Hamiltonian

$$|\beta, \gamma\rangle := U(\beta, \gamma)|+\rangle^{\otimes N} \quad (1)$$

$$= \prod_{l=1}^p U_{\text{mix}}(\beta_l) U_{\text{prob}}(\gamma_l) |+\rangle^{\otimes N}, \quad (2)$$




FIG. 2. An example of the BPSP consisting of a sequence of eight car instances and four body types. There are two cars for each body type and they must be painted different colors from within the selection of two available colors (red and blue in this case). The aim is to assign colors in a way that minimizes the number of color changes in the sequence. The suboptimal coloring shown is obtained using the classical greedy heuristic (described in Sec. II G 1) and the optimal is obtained using the IBM CPLEX optimizer [32], resulting in four and two color changes, respectively. This diagram is based on Fig. 1 in Streif *et al.* [4].

where p is the number of QAOA layers, $\beta := \{\beta_1, \beta_2, \dots, \beta_p\}$ and $\gamma := \{\gamma_1, \gamma_2, \dots, \gamma_p\}$ are the QAOA parameters, $U_{\text{mix}}(\beta_i)$ is the mixer operator, and $U_{\text{prob}}(\gamma_i)$ is the phase operator encoding the problem Hamiltonian. By tuning the QAOA parameters to minimize the energy

$$E = \langle \beta, \gamma | H_{\text{prob}} | \beta, \gamma \rangle, \quad (3)$$

we obtain a state with an energy close to that of the ground state. The lowest energy is desired because the ansatz state energy is always an upper bound to the ground state energy due to the variational principle of quantum mechanics.

B. The binary paint shop problem

The PSP consists of a set of optimization problems in the automotive industry relating to painting instances of cars in sequences various colors satisfying certain constraints (a special case of the minimum bijection problem where the graph is a path [33]). The problem was introduced by Epping *et al.* [34], where it was shown to be NP-complete in both the number of colors and the number of car instances in a sequence. There are three main variants of the problem that are typically used. These are, listed in order of increasing granularity, the multicar multicolor PSP, the multicar PSP [35], and the BPSP. Each problem assumes that the order of cars in the sequence is predetermined by some external factor. This is because the sequence order is more importantly optimized with respect to the factory stages before and after the paint shop, which are typically the body shop stage and the trim/chassis/assembly shop stage, respectively [34,36]. In the multicar multicolor PSP, there can be any number of car instances of the same body type in the sequence, where there are color count requirements for any number of colors for each body type. For example, there could be five body type A cars and three body type B cars in a sequence, where body type A cars require two to be red, two to be blue, and one to be green, while body type B cars require two to be red and one to be blue. This version of the problem is general in nature and can be directly used to model a variety of paint shop problems in modern car factories. The multicar PSP is more specific with the added constraint that there can only be up to two colors per car body type. This version of the problem is directly relevant to factories for optimizing the choice of filler coats. The BPSP, shown in Fig. 2, adds a further constraint where there can only

be up to two cars of each body type in the sequence, and they each are required to be painted different colors (e.g., red and blue). The BPSP can be formally defined as the following:

Definition 1 (binary paint shop problem). Let $S = \{s_1, \dots, s_n\}$ be a nonempty sequence of car instances, let $B = \{b_1, \dots, b_m\}$ be a nonempty set of car body types, and let $\Omega : S \rightarrow B$ be a map that labels the car body type for each car instance in S . A tuple (S, B, Ω) is an input to the BPSP when S contains exactly two car instances of each car body type in B under the map Ω . That is, for all $b_k \in B$, there exists exactly two distinct car instances $s_i, s_j \in S$ such that $\Omega(s_i) = \Omega(s_j) = b_k$. Let $\mathcal{C}$ be a set of colors. A *coloring* is a map $f : S \rightarrow \mathcal{C}$ that assigns a color from $\mathcal{C}$ to each car instance in S . Each coloring has a property called the number of color changes Δ_C , which is the number of neighboring car instances $s_i, s_{i+1} \in S$ where $f(s_i) \neq f(s_{i+1})$. For the BPSP with input (S, B, Ω) , a coloring f_2 maps to a set of two colors $\mathcal{C}_2 = \{1, 2\}$ and has the constraint that each pair of car instances in S that are assigned the same car body type from B under Ω are assigned different colors in $\mathcal{C}_2$ under f_2 . That is, for all distinct pairs $s_i, s_j \in S$ such that $\Omega(s_i) = \Omega(s_j)$, it must follow that $f_2(s_i) \neq f_2(s_j)$. The number of color changes for f_2 , can be calculated as $\Delta_C = \sum_{i=1}^{n-1} |f_2(s_i) - f_2(s_{i+1})|$ since $|f_2(s_i) - f_2(s_{i+1})| = 1 - [f_2(s_i) = f_2(s_{i+1})]$ when $f_2(\cdot) \in \{1, 2\}$, where $[\alpha]$ is the Iverson bracket of the statement α (1 when true, 0 otherwise). The BPSP can now be stated as the following. Given an input (S, B, Ω) , find a coloring f_2 that minimizes the number of color changes Δ_C , that is,

$$f_2^* := \arg \min_{f_2} \sum_{i=1}^{n-1} |f_2(s_i) - f_2(s_{i+1})|, \quad (4)$$

$$\Delta_C^* := \sum_{i=1}^{n-1} |f_2^*(s_i) - f_2^*(s_{i+1})|, \quad (5)$$

where f_2^* and Δ_C^* denote an optimal coloring and corresponding minimized color change count, respectively.

Even though the BPSP is simpler than the multicar variants of the problem, it is still NP-complete and APX-hard [22,23], making it an interesting problem for analysis that could lead to ideas for the more general problems.

Example 1 (BPSP). Let (S, B, Ω) be the input for an eight-car BPSP instance, that is, there is a sequence of eight cars S that are each assigned a car body from a set B such that two

copies of each body are present in the sequence. The sequence of cars S and set of car bodies B can be written as

$$S = \{s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8\},$$

$$B = \{\bullet, \blacktriangle, \blacksquare, \blacklozenge\}.$$

The car bodies of the cars in the sequence are defined through the relation Ω as

$$\begin{aligned}\Omega(s_2) &= \Omega(s_7) = \bullet, \\ \Omega(s_1) &= \Omega(s_3) = \blacktriangle, \\ \Omega(s_5) &= \Omega(s_8) = \blacksquare, \\ \Omega(s_4) &= \Omega(s_6) = \blacklozenge.\end{aligned}$$

The sequence of car bodies is used to summarize the problem

$$\Omega(S) = \{\blacktriangle, \bullet, \blacktriangle, \blacklozenge, \blacksquare, \blacklozenge, \bullet, \blacksquare\}.$$

The goal is to find a coloring $f_2 : S \rightarrow \mathcal{C}_2$, where $\mathcal{C}_2 = \{1, 2\} \equiv \{\mathbf{R}, \mathbf{B}\}$, that minimizes the number of color changes Δ_C such that the color of cars that have the same body are different, that is,

$$\begin{aligned}f_2(s_2) &\neq f_2(s_7), & f_2(s_1) &\neq f_2(s_3), \\ f_2(s_5) &\neq f_2(s_8), & f_2(s_4) &\neq f_2(s_6).\end{aligned}$$

There are many approaches to this problem. An example solution from a classical greedy heuristic (described in Sec. II G 1) is

$$f'_2(S) = \{\mathbf{R}, \mathbf{R}, \mathbf{B}, \mathbf{B}, \mathbf{B}, \mathbf{R}, \mathbf{B}, \mathbf{R}\}, \quad \Delta'_C = 4.$$

The optimal solution can be obtained using an exact solver such as CPLEX [32], resulting in

$$f^*_2(S) = \{\mathbf{R}, \mathbf{B}, \mathbf{B}, \mathbf{B}, \mathbf{B}, \mathbf{R}, \mathbf{R}, \mathbf{R}\}, \quad \Delta^*_C = 2.$$

A succinct presentation can be used to display both the problem input and the solution by labeling the car instances by their bodies and drawing them with their solution colors,

$$\{\blacktriangle, \bullet, \blacktriangle, \blacklozenge, \blacksquare, \blacklozenge, \bullet, \blacksquare\}. \quad (6)$$

C. Mapping the BPSP to QAOA

To solve combinatorial optimization problems using QAOA, they are first transformed into Ising Hamiltonian energy minimization problems. These Hamiltonians have already been obtained in previous studies for the BPSP [4] and the multicar PSP [35]. We describe the transformation for the BPSP here with emphasis on careful justification for the construction of the Ising Hamiltonian. Remarkably, it is possible to map the BPSP to QAOA using a single qubit for each car body type, rather than for each car instance. With such a mapping, the state of each qubit corresponds to the color assigned to the first occurrence of the two car instances with the same corresponding body type. Couplings between qubits in the Ising Hamiltonian are determined by iterating through the sequence of car instances and observing the relational positions of the car body types. For each adjacent pair of car instances in the sequence, if the pair are both first or both second occurrences of their assigned body

type, then they can be influenced to have the same color by adding a coupling between their corresponding qubits with a -1 coupling strength. If one car instance of the pair is the first occurrence of their assigned body type, while the other is the second occurrence, then the cars can be influenced to have the same color by adding a $+1$ coupling between their corresponding qubits. In the case that both car instances of the adjacent pair have the same body type, then the term becomes a constant since $Z_i Z_i = 1$. Finally, a value of $1/2$ is added to each term in the Ising Hamiltonian's sum over couplings so that the energy corresponds to the number of color changes. After following these steps and refactoring the expression to sum over pairs of adjacent body types, we obtain a BPSP Ising Hamiltonian that directly represents an Ising graph

$$H_{\text{BPSP}} = \frac{1}{2} \sum_{(i,j) \in E} J_{ij} Z_i Z_j + \frac{C}{2}, \quad (7)$$

where J_{ij} is the total coupling strength between the pair of qubits corresponding to body types i and j , E are the edges in the Ising Hamiltonian's corresponding Ising graph with nonzero coupling, and C specifies an integer constant. A detailed derivation of this expression is included in Appendix B.

The Ising graph approaches four-regular and $J_{ij} = \pm 1$ couplings with the distribution $P(J_{ij} = -1) = 2/3$ and $P(J_{ij} = 1) = 1/3$ as the problem size increases [4]. Note that the BPSP Ising Hamiltonian is in the same form as the Hamiltonian for integer weighted MAX-CUT,

$$H_{\text{MaxCut}} = -\frac{1}{2} \sum_{(i,j) \in E} J_{ij} (1 - Z_i Z_j) \quad (8)$$

$$= \frac{1}{2} \sum_{(i,j) \in E} J_{ij} Z_i Z_j + \frac{C'}{2}, \quad (9)$$

where C' is an integer constant different to C from Eq. (7).

D. The QAOA depth and its limitations

When choosing the QAOA depth p , there is a trade-off to consider. Higher values of p reduce the error caused by the time discretization of the Trotterized approximation; however, higher values of p also require a higher number of quantum circuit gates to implement, which increases the amount of noise in the computation. There is also an algorithmic performance limitation that occurs when p is too low with respect to the number of qubits N . The QAOA Hamiltonian is local, meaning that it is a sum of operators that are bounded in the number of qubits that they act on, which for QAOA is two qubits. Due to this limitation, measured qubits of a QAOA ansatz at QAOA depth p can only be correlated with qubits that are within the p -neighborhood of it, which includes qubits within p distance from the measured qubit in the graph representing the problem Ising Hamiltonian. If the p -neighborhood does not span all of the qubits, then there can be bounds on the algorithmic performance. In Bravyi *et al.* [17], it was shown that there exist problem graphs G such that when $p < (\log_2(N)/3 - 4)/D$, where $D \geq 3$ is the maximum degree of vertices in G , the QAOA variational energy for Z_2 -symmetric problem Hamiltonians (such as for MAX-CUT)




FIG. 3. A graphic highlighting the RQAOA reduction process. Eliminating variables by choosing edge correlations to round at each RQAOA reduction step can be thought of as conducting a classical greedy depth-first search on a decision tree. The traversed direction at each step is chosen by rounding the ZZ correlation of the pair of qubits with the largest correlation magnitude, a depth-first search heuristic where the ZZ correlations are calculated using a quantum device. By rounding a ZZ correlation to ± 1 , a logical relation is established between the qubits, enabling one of them to be removed from the problem Ising Hamiltonian and assigned a state classically once the reduction process completes. The available choice of edges to round, $e_k^{(i)}$ in the diagram, are the edges in the reduced Ising graph during step i , where $k \in \{1, 2, \dots, |E^{(i)}|\}$, and the number of nodes and edges in the Ising graph during step i are denoted by $|V^{(i)}|$ and $|E^{(i)}|$, respectively. There are many paths in the decision tree that lead to optimal solutions. For instance, let Γ be a distribution of problem instances in which all states have an equal likelihood of being solution states when selecting a problem instance at random from Γ . Given a random problem instance from Γ , during each RQAOA step, if an edge is assigned a ZZ correlation rounding of ± 1 with a uniform probability distribution, then there is a 50% probability that the choice is correct, leading the path to progress toward a solution state (ignoring solution degeneracy).

on G is upper bounded by $(5/6 + \sqrt{D-1}/3D)E_{\max}$, where $E_{\max}$ is the exact largest energy of the problem Hamiltonian. For more typical instances of the max-independent set problem, it was shown in Ref. [37] that when qubits have an average degree of d , the QAOA solution approximation is bounded away from optimal on average by a constant for fixed d when p is less than a small degree-dependent multiple of $\log_2(N)$. In a follow-up paper [38], a similar result was shown for d -regular bipartite graphs for the MAX-CUT and max-independent set problems. In separate studies, it was shown that the performance of constant depth QAOA for bounded-degree problem Ising graphs cannot achieve the same scaling as the best classical algorithms for increasing problem sizes [39,40]. However, these QAOA depth limitation studies conclude that when the p -neighborhood for each qubit

covers the whole problem graph, then there is no indication that the algorithm's power is limited.

E. The recursive quantum approximate optimization algorithm

The RQAOA was introduced by Bravyi *et al.* [17] in 2019 to overcome the above limitations relating to the QAOA Hamiltonian being local. It was shown to outperform standard QAOA for frustrated Ising models on random three-regular graphs. It was later shown analytically for $p = 1$ to produce an approximation ratio of 1 for MAX-CUT on $2n$ -vertex complete graphs [41], while baseline QAOA's is strictly less than $1 - 1/(8n^2)$. An overview diagram of RQAOA is provided in Fig. 3. It works by initially performing QAOA and optimizing the ansatz state with respect to the parameters as usual. Then the correlations between each of the adjacent qubits in the




FIG. 4. (a) An example of the RCC for a line Ising graph QAOA circuit. The ZZ correlation $M_{0,1}$ only requires qubits 0 and 1 to be measured. Only the gates that contribute to the qubit 0 and 1 measurements are needed. The grayed out gates can be removed since they can be iteratively commuted to the end of the circuit. (b) The shorthand definition used for $R_{ZZ}(\theta_{(i,j)}^{(l)})$ gates applied to qubits i and j within QAOA layer l . (c) Circuits showing the equivalence of single-qubit noise gates replacing trimmed two-qubit gates introduced in Sec. II E 2. The final noise gate has $R_Z(\theta_{(i,j)}^{(l)})$ applied in half of the shots and $R_Z(-\theta_{(i,j)}^{(l)})$ applied in the other half.

problem graph are calculated as $M_{i,j} = \langle \beta, \gamma | Z_i Z_j | \beta, \gamma \rangle$. The largest magnitude correlation is chosen and a constraint is imposed on it to effectively round it to ± 1 ,

$$Z_j = \text{sgn}(M_{i,j}) Z_i. \quad (10)$$

Substituting this into the problem Hamiltonian allows one of the variables to be eliminated, resulting in a smaller problem graph. The incident edges of the nodes corresponding to the two variables are merged as a result, with a minus sign multiplied to the edge weights depending on the relation. Then, QAOA is applied again to this smaller graph to find the largest magnitude correlation and the process is repeated until the problem graph is reduced to be small enough to classically solve exactly. This process typically requires $N - 1$ steps to fully reduce the problem graph to a single node, where N is the number of car bodies. In some cases, a single step will decrease the reduced graph node count by more than 1. This is because when merging incident edges of logically related variables that are incident to the same node, the two edges combine into one with a resulting weight that can possibly cancel to zero. This will effectively remove the edge from the graph, which can possibly cause a single qubit to become isolated from the rest of the graph, enabling it to be removed.

The RQAOA can be thought of as a greedy tree search, where directions are chosen based on measured correlations. Using the largest correlation as the direction is only a simple method for making a decision. Significant improvements can be achieved by implementing more advanced methods. This

has been demonstrated by using reinforcement learning [42], problem-specific classical subroutines [43], quantum random access codes [44], and successive interference cancellation combined with majority vote [45].

1. Reverse-causal cones

The RQAOA allows us to avoid needing to prepare the full ansatz state. This is because RQAOA only needs to know the largest correlation magnitude between each adjacent pair of qubits to round a correlation to ± 1 when reducing the graph and the parametrized ansatz state can be optimized at each reduction step by calculating the energy at each optimization iteration, which is just a weighted sum of the correlations. When measuring a correlation, only the operations that connect to the pair of qubits within their $(p - l + 1)$ -neighborhood for each QAOA layer l can influence their measured state. An example is shown in Fig. 4. The set of these operations for a target pair of qubits (i, j) is called the RCC and is the set of gates not commuting with the operator product $Z_i Z_j$, which can be determined by working backwards through the ansatz state preparation circuit [27]. The process for constructing it is the following. In the last QAOA layer $l = p$, include all two-qubit operators and involved qubits that include the target pair of qubits. In the second last layer $l = p - 1$, include all two-qubit operators that involve any qubit that has been included so far. Repeat this procedure until all layers have been processed. The number of RCC circuits used to calculate the energy of an Ising graph is equal to

its edge count, which for an N -node regular graph, scales as $\mathcal{O}(N)$. The total number of RCC circuits for RQAOA is then

$$\text{RCC circuit count} = \sum_{i=1}^{N-1} |E^{(i)}|, \quad (11)$$

where $|E^{(i)}|$ is the number of edges in the Ising graph during reduction step i . The number of edges monotonically decreases with respect to graph reduction, thus the total number of RCC circuits for RQAOA on a regular problem graph scales as $\mathcal{O}(N^2)$.

2. Trimming RCC outer two-qubit gates

In short-depth QAOA circuits, there are often cases when the RCC's first layer uses extra qubits that are not included in the second layer. These extra qubits and the $R_{ZZ}(\theta)$ gates that connect them can be trimmed and replaced by equivalent single-qubit noise gates applied to the corresponding connected qubits [as shown in Fig. 4(c)]. The noise gates can be implemented by splitting the circuit shots into two halves, which each apply either the $R_Z(\theta)$ or $R_Z(-\theta)$ gate based on the equal probability of measuring a 0 or 1 on the trimmed neighbor qubit before the $R_{ZZ}(\theta)$ gate is applied. This means that by distributing the shots equally across each possibility of 0 and 1, the two-qubit gates and corresponding qubits can be replaced with their corresponding single-qubit $R_Z(\pm\theta)$ gates. These trimmed RCC circuits are generally easier to implement and have lower levels of gate error, along with a lower quantum resource cost overhead when compiling to hardware.

RCC circuits with multiple trimmed qubits can be implemented in practice by constructing a circuit for each bitstring that represents the possible measured states of the removed qubits. Each of these trimmed circuits have corresponding $R_Z(\pm\theta)$ gates applied based on their associated removed qubit bit value in the bitstring. Using this approach, the number of unique trimmed RCC circuits, Γ , used to implement an RCC circuit is upper bounded by 2^k , where k is the number of removed qubits. As for the number of removed qubits, given a d -regular problem Ising graph with QAOA depth p , the number of qubits in the first layer of an RCC is upper bounded by $\sum_{l=0}^{p-1} 2(d-1)^l$, assuming there is the maximum of $\sum_{l=0}^{p-1} 2(d-1)^l$ qubits included in the second layer. The number of removed qubits k is thus upper bounded as

$$k \leq 2(d-1)^p. \quad (12)$$

Therefore, the number of unique trimmed RCC circuits Γ is upper bounded as

$$\Gamma \leq 2^{2(d-1)^p}. \quad (13)$$

It is also limited when increasing depth p begins to make the RCC qubits saturate the problem Ising graph, which reduces the possible additional qubits that can be in the first layer and not the second. For an untrimmed RCC circuit with N number of shots, we evenly divide the shots between each of the trimmed RCC circuits, resulting in $N/2^k$ shots each. The correlation $M_{i,j}$ is then calculated as the average over correlations measured from each of the trimmed RCC circuits.

An added benefit of dividing the shots in this way is that the variance of the measured correlation is reduced because it no longer depends on the probabilistic nature of measurement outcomes of the outer qubits.

It is tempting to further remove the noise gates entirely, because reducing circuit noise is often the goal in quantum computing. It turns out that the noise-gate choices of $R_Z(\theta)$ or $R_Z(-\theta)$ from a single trimmed qubit have the same effect on the ZZ-correlation due to symmetry, and is different to removing the noise gate altogether. Thus, when there is only a single trimmed qubit, the noise gate can be replaced by only one of its possible choices of gate. For any additional trimmed qubits, the full noise gates are required. This is because different combinations of noise-gate angle choices can induce different correlation effects across qubits that affect the ZZ-correlation observable.

F. Bypassing the QAOA optimization loop using precomputed parameters

The QAOA typically involves a classical optimization loop over the QAOA parameters to search for extrema of the cost function where the cost is evaluated as the QAOA ansatz state energy at each iteration using a quantum device. This optimization process typically requires many evaluations of the ansatz energy since the volume of parameter space grows exponentially with the QAOA depth p [46] and the energy landscape is prone to vanishing gradients due to barren plateaus, which under certain conditions can exponentially decrease the gradient with respect to qubit count [18,20]. This is particularly a concern in NISQ era since the noisy stochastic nature of quantum devices increases the difficulty in navigating the energy landscape, and barren plateaus can be further induced by noise [19].

There are a variety of methods to alleviate the problem of parameter optimization [47–49]. It turns out that on families of problem graphs that share certain characteristics, the optimized QAOA parameters concentrate about fixed values [25,26,29,46,50–52]. The mean optimal parameters for random problem instances appear to be independent of problem size and the variance tends to decrease as problem size increases [27,52]. For unweighted MAX-CUT problem graphs, the median QAOA parameters computed over all graphs of the same size, appear to give good approximation ratios for typical instances [51]. This allows for the possibility of calculating QAOA parameters that can be transferred to unseen instances, enabling a good approximation to be achieved without requiring a classical optimization loop. The convergence of QAOA parameters and the transferability between different instances can be explained and predicted based on the types of local subgraphs within the problem graph [50]. By observing the energy landscape of all local subgraphs of certain families of problem graphs (such as all even-regular or odd-regular graphs) separately there are clear patterns that have each subgraph's global extrema overlap in similar regions of the QAOA parameter space. This indicates that certain QAOA parameters will cause the QAOA ansatz energy for all of the local subgraphs to be close to a global minima (or maxima) at the same time, minimizing (or maximizing) the total QAOA ansatz energy. For MAX-CUT on

regular graphs with fixed weights, fixed parameters can be calculated by analyzing local subgraphs that are worst case under a no small cycles conjecture, producing solutions with approximation ratios better than some large worst-case guarantee [53]. Comparing with the Goemans-Williamson (GW) algorithm on generic graphs [54]—the classical MAX-CUT polynomial algorithm based on semidefinite programming with the best approximation ratio guarantee $\alpha = 0.87856$ —these precomputed QAOA parameters have been shown to give a performance guarantee on three-regular graphs better than the GW algorithm for $p \geq 11$ with the guarantee approaching unity with a gap proportional to $1/\sqrt{p}$ [25]. This performance guarantee also applies to general d -regular graphs with approximation ratio proportional to $1/\sqrt{d}$ for fixed p . There are also results showing that averaging optimal QAOA parameters over many random three-regular MAX-CUT instances give parameters that provide good approximations to other larger three-regular MAX-CUT instances [55]. Additionally, a neural network approach has been shown to calculate problem-specific parameters that provide near-optimal approximations for individual MAX-CUT problem graphs [49]. Fixed parameters for MAX-CUT on regular graphs can be efficiently calculated in the limit of infinite problem size via a classical variation of QAOA that is based on tensor network techniques called tree-QAOA [27]. Furthermore, analytical formulas have been derived to calculate $p = 1$ QAOA expectation values for general MAX-CUT graphs [26,56,57] and to further analytically optimize with respect to $p = 1$ QAOA parameters for classes of Ising graphs [26]. The QAOA expected energy for typical instances of the Sherrington-Kirkpatrick model has also been formulated in the infinite size limit for arbitrary values of p with a time complexity of $\mathcal{O}(16^p)$ [58]. The parameter transfer methods and results mentioned so far mostly target unweighted MAX-CUT problem graphs. However, it has been shown that median parameters over typical unweighted MAX-CUT instances can be transferred to weighted MAX-CUT problem graphs by rescaling γ based on average weight magnitude and average node degree of the target graphs [26,29]. For a large dataset of random instances with up to 20 nodes and various weight distributions, it was shown that the mean approximation ratio decreased by only 0.02 when using parameter transfer from a single vector (calculated as the median optimized values over all nine-node nonisomorphic unweighted MAX-CUT graphs) of QAOA parameters compared to directly optimizing the instances in the dataset. This reduction further improved to only 0.012 when evaluating QAOA using parameter transfer from ten additional random samples from a pretrained metadistribution of optimized QAOA parameters based on kernel density estimation [59]. There are also techniques to transfer an arbitrary set of optimal parameters into the adequate domain using symmetries of the problem graph [60].

There are two main approaches to estimating fixed parameters for the BPSP that we will be using in this work. The first is to use the results for four-regular graphs with ± 1 couplings obtained in Ref. [4] calculated using tensor networks on tree-QAOA [27]. The second approach is to estimate the fixed parameters as the median over many smaller random instances using the QAOAKit Python library [28].

Approximation measure

To compare methods fairly, we aim for a measure of approximation quality that is independent of problem instance and preferably independent of the objective function's properties, such as whether higher values are better or whether the values span positive and negative values. To keep comparisons simple, we construct a measure scaled between 0 and 1 where 1 means that the approximation returns optimal solutions, while 0 means that the approximation is not useful. Using the following definition, also used in Refs. [29,61], the approximation measure for the approximation value is the fraction between the worst case value and the best possible value of the objective function:

$$\text{approx. measure} = \frac{(\text{worst}) - (\text{approx. val.})}{(\text{worst}) - (\text{best})}. \quad (14)$$

This can also be thought of as one minus the residual energy density [48]. The “worst value” can be specified as the value that is furthest away from the optimal solution, or it can be specified as the average value over inputs that are randomly chosen from a uniform distribution, in which case the approximation measure represents how useful the approximation is compared to random guessing.

G. Classical approaches

Here, we describe three classical algorithms for solving the BPSP that are used as points of comparison in our analyses: greedy heuristic, recursive greedy heuristic, and exhaustive search. We use the following problem instance as an example to help describe them:

$$\Omega(S) = \{\blacktriangle, \bullet, \blacktriangle, \blacklozenge, \blacksquare, \blacklozenge, \bullet, \blacksquare\}.$$

1. Greedy heuristic

The first car in the sequence is assigned red (which consequently assigns blue to its paired car). Then, iterating through the sequence, each unassigned car is assigned the same color as the previous car in the sequence. For example,

$$\begin{aligned} \Omega(S) &\rightarrow \{\color{red}\blacktriangle, \bullet, \color{blue}\blacktriangle, \blacklozenge, \blacksquare, \blacklozenge, \bullet, \blacksquare\} \\ &\rightarrow \{\color{red}\blacktriangle, \color{red}\bullet, \color{blue}\blacktriangle, \blacklozenge, \blacksquare, \blacklozenge, \color{blue}\bullet, \blacksquare\} \\ &\rightarrow \{\color{red}\blacktriangle, \color{red}\bullet, \color{blue}\blacktriangle, \color{blue}\blacklozenge, \blacksquare, \color{red}\blacklozenge, \color{blue}\bullet, \blacksquare\} \\ &\rightarrow \{\color{red}\blacktriangle, \color{red}\bullet, \color{blue}\blacktriangle, \color{blue}\blacklozenge, \color{blue}\blacksquare, \color{red}\blacklozenge, \color{blue}\bullet, \color{red}\blacksquare\}, \end{aligned}$$

resulting in a color change count of $\Delta_c = 4$.

2. Recursive greedy heuristic

Each car body pair is added to the sequence one by one, in the same order they appear in the problem car sequence. Each pair's colors are assigned to minimize the number of color

TABLE I. The precomputed QAOA parameters for four-regular graphs with unit couplings, referred to as “BPSP Precomp” in Fig. 5. The data was obtained from Streif *et al.* [4], which uses the approach presented in Ref. [27]. We include the values here for ease of reference.

QAOA depth	β_1	γ_1	β_2	γ_2	β_3	γ_3	β_4	γ_4
1	-0.392 69	0.523 58						
2	-0.534 11	0.407 84	-0.282 96	0.739 74				
3	-0.587 94	0.354 50	-0.423 18	0.651 38	-0.223 01	0.754 26		
4	-0.604 98	0.315 00	-0.477 80	0.587 54	-0.361 27	0.673 22	-0.187 53	0.771 20

changes directly after they are added. For example,

$$\begin{aligned}
 \Omega(S) &\rightarrow \{\triangle, \triangle\} \\
 &\rightarrow \{\triangle, \bullet, \triangle, \bullet\} \\
 &\rightarrow \{\triangle, \bullet, \triangle, \triangle, \triangle, \bullet\} \\
 &\rightarrow \{\triangle, \bullet, \triangle, \triangle, \triangle, \bullet, \triangle, \triangle\},
 \end{aligned}$$

resulting in a color change count of $\Delta_c = 3$.

3. Exhaustive search

This algorithm finds exact solutions to the BPSP by mapping the BPSP instance to an Ising Hamiltonian using the method shown in Sec. II C then uses an eigensolver to find the eigenvector with the smallest eigenvalue, which corresponds to the minimum number of color changes. In this case, the NumPy minimum eigensolver simply sorts the diagonal elements and returns the smallest, essentially performing an exhaustive search. The returned solution is

$$\Omega(S) \rightarrow \{\triangle, \bullet, \triangle, \triangle, \triangle, \bullet, \triangle, \triangle\},$$

resulting in a color change count of $\Delta_c = 2$.

III. RESULTS

A. RQAOA performance with precomputed QAOA parameters

We coded RQAOA in Python and used a local IBM quantum state vector simulator to obtain results for random BPSP instances, comparing the quality of solutions with QAOA and classical approaches. To maximally observe the effectiveness of the quantum algorithm itself, we do not use a classical solver once the RQAOA is reduced to a certain size (like is suggested in the original paper). Instead, the graph reduction continues until one qubit remains, at which point the solution is obtained through the logical relations recorded at each reduction step. We generate a sample of 20 random BPSP instances for each car body count ranging from 4 to 20. The performance of each approach is measured using two sample averages on the solutions: the color change count and the approximation measure [described in Eq. (14)]. The sample error bars correspond to standard errors. We compare a total of nine approaches to the BPSP, three for RQAOA, three for QAOA, and three classical algorithms. For each reduction step of RQAOA and for QAOA, the three methods of choosing the parameters are: optimization, QAOAKit precomputations, and tailored precomputations for four-regular unit-coupling graphs (labeled “BPSP Precomp” and listed in Table I). The three classical algorithms are the greedy heuristic, the recur-

sive greedy heuristic, and the exhaustive search for optimal solutions. Energy evaluations in RQAOA and QAOA are calculated using a state vector simulator with 4096 shots and all optimized data are generated using the Nelder-Mead optimizer with a parameter absolute-error tolerance of 0.0001 and initial QAOA parameters obtained using QAOAKit.

The performance results are shown in Fig. 5. We can see that all approaches to choosing parameters for RQAOA perform near-optimally, even for QAOA depth $p = 1$. The RQAOA QAOAKit Precomp performs just as well as QAOAKit optimized, and occasionally performs marginally better (as shown for 14 car $p = 1$ and 17 car $p = 2$ approximation measures). There are occasional dips in solution quality when using precomputed BPSP parameters; however, overall, it still performs close to optimally. This is surprising, because although the problem graph is initially close to four-regular with unit couplings, its structure can drastically change as the graph is reduced during the steps of RQAOA.

To better understand the effects of unoptimal QAOA parameters, we added variation to optimized parameters to observe the effects on performance, with results shown in Fig. 6. Using the same simulation settings used for Fig. 5, the data are calculated for various standard deviations by initially obtaining QAOA parameters through optimization, then adding a variation that is sampled from a normal distribution with the specified standard deviation. For RQAOA, parameters are varied at each reduction step by resampling. Each data point is averaged over 100 random BPSP instances for each car body count. The vertical lines are the standard deviations of displacements between QAOAKit parameters and optimized, calculated over all instances, and all reduction steps for RQAOA. The data shows that RQAOA is less sensitive to parameter variation than QAOA, indicating a form of noise robustness in the parameters. We suspect that RQAOA also has a level of robustness to quantum device noise as well; however, that was not directly tested for in these simulations. The data also shows that the variation caused by QAOAKit is comfortably within the region where RQAOA gives optimal results. As the problem size increases up to 14 car bodies, we observe no notable changes to the performance in relation to the amount of parameter variation. However, the QAOAKit variations for RQAOA, shown by blue vertical lines, appear to decrease for increasing problem sizes. This is likely because the QAOAKit deviation from optimal is largest for small graphs that are highly irregular, while larger BPSP instances have lower proportions of RQAOA reduction steps with Ising graphs that fall in this regime. We suspect that this will eventually begin to increase again for large problem sizes due to the extrapolation precision of QAOAKit being limited by data obtained on relatively small problem graphs. Additionally,




FIG. 5. The quality of BPSP solutions from a variety of algorithms plotted against the number of car bodies. We compare between three approaches to RQAOA, three approaches to QAOA, and three classical algorithms. For RQAOA and QAOA, we show results for optimized parameters as well as parameters obtained using QAOAKit [28], and parameters obtained by precomputation of four-regular graphs with unit couplings specifically, from Ref. [27]. The classical algorithms being compared with are the greedy heuristic, recursive greedy heuristic, and the exhaustive search. For each car body count, a set of 20 BPSP instances are chosen at random from a uniform distribution. Each data point is calculated as the mean value obtained from the corresponding set of instances and the error bars are calculated as the standard error. (Top) The mean color change counts for QAOA depths 1–3. The RQAOA values shown for the three choices of parameters are almost always equal to the optimal values from the exhaustive search. (Bottom) The mean approximation measures for QAOA depths 1–3. The approximation measure for a solution of an instance is calculated as the color change count scaled linearly between the optimal count (corresponding to 1) and the least optimal (corresponding to 0), where the extremas are obtained using exhaustive search.

for all problem sizes used in these simulations, we observe that RQAOA produces optimal solutions. For large enough problem sizes where RQAOA no longer consistently gives optimal solutions, it is possible that this noise robust region will no longer appear. Although for such large problems, we believe that the performance of RQAOA with QAOAKit precomputed parameters will remain superior to QAOA, even with optimized parameters.

B. Classical and quantum resource requirements

To better understand how RQAOA scales in terms of quantum resource requirements, we have recorded a variety of resource metrics and compiled the results.

In Fig. 7, we compare the CNOT counts, CNOT depths, and qubit counts within QAOA RCC and full circuits, for QAOA depths 1–4 and problem sizes ranging from 4 to 18. In particular, we compare the trimmed RCC circuits against the full circuits. Additionally, we compare metrics for circuits before and after compilation to the 27-qubit *ibmq_montreal* device using Qiskit’s transpiler (v0.26.2) with the StochasticSwap hardware mapping algorithm with any initial SWAP gates removed (since qubits can be sim-

ply relabeled in code). All data points are calculated as the average over 8 random BPSP instances. The results show that for QAOA depths 1 and 2, the RQAOA RCC circuit resource requirements are substantially lower than the full circuits and are approximately constant with respect to increasing numbers of car bodies. This is likely because the included numbers of qubits in the trimmed RCCs, upper bounded by $\sum_{l=0}^{p-1} 2 \times 3^l$ (equalling 2 and 8 for $p = 1$ and $p = 2$, respectively), are much fewer than the problem graph qubit counts. Thus, increasing the size of the problem graph does not have much effect on the span of the RCC. For QAOA depth 3, the number of qubits in the trimmed RCCs is upper bounded by 32 and begins to reliably span the full problem graphs. This leads to the RCC circuit resources growing for increasing car body count throughout our simulations of up to 18 qubits. As the QAOA depth increases, the relative resource differences between RCCs and full circuits appear to diminish, although the reductions in CNOT counts and depths for transpiled circuits remains significant. For higher QAOA depths, we observe that the CNOT count and CNOT depth tends to scale linearly with respect to car body count. Similarly, for 18 car bodies, we notice the




FIG. 6. Plots showing the effect on performance caused by deviations of QAOA parameters away from optimal. The data are calculated by first obtaining QAOA parameters through optimization, then adding variation that is sampled from a normal distribution with the specified standard deviation. For RQAOA, parameters are varied at each reduction step by resampling. The vertical lines are the standard deviations of displacements between QAOAKit parameters and optimized, calculated over all instances, and all reduction steps for RQAOA. There are 100 random BPSP instances for each car body count. Each circuit is executed on a state vector simulator with 4096 shots and all optimized data are generated using the Nelder-Mead optimizer with a parameter absolute-error tolerance of 0.0001. All initial parameters for optimization are obtained using QAOAKit. Error bars are calculated as the standard error.

CNOT count and CNOT depth scaling linearly with respect to QAOA depths.

Further investigating quantum resource requirements, Fig. 8 compares the number of individual circuits required to obtain solutions between the various QAOA-based algorithms. The quantum algorithms presented are the precomputed-parameter and optimized-parameter implementations of QAOA, RQAOA using RCCs, and RQAOA without using RCCs (full). Choosing to optimize the $2p$ parameters instead of precompute them introduces a multiplicative factor to the number of circuits that is similar in quantity to the circuit counts shown for QAOA full optimized. By assuming precomputed parameters to ignore this overhead, the typical circuit counts for $2N$ -car BPSP instances scales as $N - 1$ using RQAOA full, $\mathcal{O}(N^2)$ using RQAOA with RCCs, and is equal to 1 using QAOA. The circuit counts for full RQAOA with precomputed parameters are expected to exceed the circuit counts for fixed 20-qubit optimized QAOA at minimum qubit counts of 100, 200, and 400 for QAOA depths 1, 2, and 3, respectively. Considering that circuit counts for optimized QAOA also appear to be increasing with respect to the qubit count, we expect that circuit counts for full RQAOA with precomputed parameters will always be favorable over optimized QAOA in practice. We find that RQAOA using RCCs and precomputed parameters requires higher numbers of circuits than QAOA with optimization. This provides a trade-off between the number of quantum circuits and their individual resource requirements. Note that the RCC circuits for this data omit the outer two-qubit gate trimming discussed in Sec. II E 2.

With the use of RCCs, low-depth RQAOA circuits exceed the performance of full QAOA while being considerably

smaller in size. A natural question to ask here is how does the classical simulatability of quantum RCC circuits compare to full QAOA circuits? We compare the resource usages within classically simulated QAOA circuits in the form of entanglement entropies and bond dimensions of MPS representing quantum states [62–65]. By efficiently encoding a state’s entanglement structure between qubits, MPS can substantially reduce both storage requirements and computational runtime of quantum algorithm simulations. An MPS expresses a quantum state as a compressed one-dimensional product of tensors, where each tensor corresponds to a single qubit. Neighboring qubits are connected by an internal bond index representing shared quantum correlations. A pure state can be transformed into MPS form by performing a series of Schmidt decompositions on contiguous bipartitions along the qubit line, each implemented via a singular value decomposition (SVD). Dimensions of the bond indices associated with singular values of zero (or below some cutoff threshold) are removed since they do not significantly contribute to the state. The *bond dimension*, defined as the maximum number of remaining dimensions (Schmidt rank) among all bonds, quantifies the maximal information or entanglement capacity across any cut of the line of qubits and directly determines the computational resources required for MPS simulation. The *entanglement entropy* for a bipartition $A|B$ is given by the von Neumann entropy,

$$S(A|B) = - \sum_{\alpha} \lambda_{\alpha}^2 \log(\lambda_{\alpha}^2), \quad (15)$$

where λ_{α} are the nonzero singular values obtained in the SVD between partitions.




FIG. 7. Quantum resource requirements for QAOA RCC and full circuits. The plots compare the circuit CNOT counts, CNOT depths, and qubit counts for QAOA depths 1–4 and car body counts 4–18. Each data point is the average over eight random problem instances. The circuits are mapped to IBM quantum’s heavy-hex hardware using Qiskit’s transpiler (v0.26.2), which uses the StochasticSwap mapping algorithm, and any initial SWAP gates are removed.

Figure 9 shows these measures computed on full and RCC QAOA circuits of depths 1–4 on up to 18 car bodies, averaged over eight random BPSP instances. The entanglement entropy indicates the level of quantum entanglement required to represent the state, reported as the highest von Neumann entropy across contiguous cuts of the MPS (and across each RCC in RCC circuits). The bond dimension reflects the corresponding classical memory cost and is given by the largest Schmidt rank among the cuts. To investigate the trade-off between MPS precision and classical resource usage, we introduced cutoff thresholds of 0, 0.005, and 0.01 when truncating singular values, and recorded the total probability weight excluded, calculated as the sum of discarded squared singular values summed across all bonds. For reference, the theoretical maximum values of entanglement entropy and bond dimension along a one-dimensional MPS are $\lfloor N/2 \rfloor$ and $2^{\lfloor N/2 \rfloor}$, respectively (9 and 512 for 18 car bodies), shown as thin black lines in the plots. The results show that for QAOA depth 1, while the bond dimension appears to increase exponentially for the full circuits with respect to car body count, it remains constant when using RCC circuits. The entanglement entropy appears to increase linearly for the full circuits and is roughly constant when using RCC circuits. As the QAOA depth increases, the curves for full and RCC circuits appear to converge. This is expected since the number of qubits spanned by RCCs ap-

proaches the number of car bodies, N , as the depth increases. As the cutoff gets larger, the bond dimensions measured in our data reduce considerably, especially for RCC circuits, while the total probabilities excluded from quantum states remain relatively small. The truncation of singular values induces an approximation error that is upper bounded by two times the total discarded probability $\mathcal{E}$ [64], implying a pure state fidelity bound of $F \geq 1 - \mathcal{E}$. For these small values of excluded probability, the quantum state maintains a high state fidelity above 0.99 for almost all instances, which would have at most marginal effects on the performance of RQAOA. The variation of entanglement entropy across cutoff thresholds is considered to be negligible since the average difference in values peak at approximately 1%. Therefore, it appears that the combination of RCC circuits and bond dimension cutoffs together help to significantly reduce classical memory costs while maintaining high levels of entanglement entropy and quantum state fidelity. Although the changes in entanglement entropies and state probabilities are small in our data, they will continue to increase as the problem sizes grow. We expect that to maintain small deviations, the cutoff threshold will need to increase to compensate. It would be interesting to see how these results continue to scale with increasing cutoff thresholds and problem sizes. In particular, as the excluded probabilities become larger, how are the RQAOA solution




FIG. 8. The number of circuits used for each quantum algorithm. The algorithms compared are the precomputed parameter and optimized parameter implementations of QAOA, RQAOA using RCCs, and RQAOA without using RCCs (full). Each circuit is executed on a state vector simulator with 4096 shots, and all optimized data are generated using the Nelder-Mead optimizer with a parameter absolute-error tolerance of 0.0001. All precomputed and initial parameters are obtained using QAOAKit. For each car body count, values are averaged over 20 random BPSP instances with error bars calculated as the standard error. A log-scale y axis is chosen to visually separate the data, which scales polynomially in the number of car bodies.

qualities affected? This could help to understand the benefits of combining RCCs and MPS cutoffs for classical simulations of RQAOA and potentially lead to an increased capacity for simulating larger problem sizes.

IV. CONCLUSION

We implemented baseline RQAOA in conjunction with precomputed parameters to investigate the impact on performance for BPSP instances compared to using optimized parameters at every RQAOA reduction step. We demonstrated the substantial improvement that RQAOA provides over QAOA, producing near-optimal solutions over all BPSP instances generated up to sequences of 40 cars consisting of 20 car bodies, which are mapped to 20-qubit problem Ising graphs. We found that RQAOA with precomputed parameters from the QAOAKit Python package performed effectively the same whether they were subsequently optimized or not. To further investigate this, we added variations to optimized parameters and recorded their effect on the quality of solutions for BPSPs with up to 14 car bodies. We found that RQAOA exhibits a level of noise robustness in QAOA parameters. We suspect this extends to robustness in the noise of quantum devices as well; however, we leave this as an open question for future research. From this data, we confirmed that the standard deviations of displacements between QAOAKit precomputed parameters and optimized parameters are comfortably within the region of tolerance where RQAOA produces its best solutions.

Additionally, we performed a variety of comparisons between classical and quantum resource requirements, including

circuit counts, quantum matrix product state bond dimensions and entanglement entropies, and circuit resources such as qubit counts, CNOT counts, and CNOT depths. In particular, trimmed RCC circuits, used for measuring ZZ correlations in RQAOA, were compared with full circuits, with consideration for transpilation onto the 27-qubit *ibmq_montreal* device. We observe that RCCs substantially reduce the circuit resources for QAOA depths 1 and 2 on up to 18 car body BPSP instances, while the benefits begin to diminish as the QAOA depth increases, although the reductions in CNOT counts and depths for transpiled circuits remain significant. For the quantum circuit counts, we find that RQAOA with precomputed parameters typically requires a number of circuits equal to the number of car bodies, which is significantly lower than what QAOA requires to optimize parameters using the Nelder-Mead optimizer. We expect this relation will continue to hold for increasing problem sizes because the circuit counts for QAOA also appears to grow. We find that RQAOA using RCCs and precomputed parameters requires higher numbers of circuits than QAOA with optimization, hence providing a trade-off between the number of quantum circuits and their individual resource requirements. For bond dimensions and entanglement entropies, we find that with QAOA depth 1, they increase exponentially and linearly respectively for full QAOA circuits; however, both remain roughly constant for individual RCCs. As the QAOA depth increases, the curves for full and RCC circuits appear to converge.

Our simulations demonstrate that bypassing the optimization loop using RQAOA with precomputed parameters is an effective approach to substantially improve solution qualities over optimized QAOA while also significantly reducing




FIG. 9. Resource usages for 1D MPS simulations of QAOA circuits with and without RCCs. The names “Full” and “RCC” in the legend correspond to calculations done on full circuits and circuits split into RCCs, respectively. For circuits with RCCs, displayed values are calculated from the RCC with the highest entanglement entropy. The entanglement entropy of the state is recorded as the highest across all contiguous cuts along the MPS line of qubits. The bond dimension is the highest number of nonzero Schmidt coefficients (Schmidt rank) calculated across the cuts. The thin black lines indicate the theoretical maximum values. The cutoff thresholds 0, 0.005, and 0.01 were used to truncate Schmidt coefficients in the Schmidt decomposition of the quantum state. The probability excluded is the total amount of probability removed from the quantum state after truncation. The entanglement entropy variation across cutoff values is considered negligible since the changes peak at approximately 1%. All values are the average over eight random BPSP instances.

circuit counts. Since RQAOA appears to perform well despite its reduction steps significantly modifying the form of the original problem graphs, it is promising that this approach will continue to outperform QAOA on non-BPSP Ising graphs as well.

ACKNOWLEDGMENTS

This work was supported by the University of Melbourne through the establishment of an IBM Quantum Network Hub at the University. C.D.H. was supported by a research grant from the Laby Foundation. The authors gratefully acknowledge the support of the University of Melbourne’s Zero Emission Energy Laboratory (ZEE Lab) and the Victorian Higher Education State Investment Fund (VHESIF).

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not techni-

cally feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: DERIVATION OF QAOA AND ITS IMPLEMENTATION

Here, we give a detailed overview of QAOA and a derivation indicated by the original paper [6] to help intuitively understand how it works.

The problem Ising Hamiltonian encoding a QUBO problem can be written as

$$H_{\text{prob}} = \sum_{(j,k)} J_{jk} Z_j Z_k + \sum_j h_j Z_j, \quad (\text{A1})$$

where Z_j is the Pauli-Z operator on qubit j , h_j is the local field strength on qubit j , J_{jk} is the coupling strength between qubits j and k (where negative values correspond to ferromagnetic couplings), and the first sum is over all pairs of qubits (j, k) .

The ± 1 eigenvalues of Z_j operators correspond to the binary values in the QUBO problem instance via $b_j = (1 - Z_j)/2$.

The QAOA ansatz state on N qubits is prepared as

$$|\beta, \gamma\rangle := U(\beta, \gamma)|+\rangle^{\otimes N}, \quad (\text{A2})$$

where $U(\beta, \gamma)$ is the QAOA time-evolution operator, which is in the form of two alternating unitary operators corresponding to the problem Hamiltonian and a mixer Hamiltonian,

$$U(\beta, \gamma) := U(\beta_1, \gamma_1, \beta_2, \gamma_2, \dots, \beta_p, \gamma_p) \quad (\text{A3})$$

$$= \prod_{l=1}^p U_{\text{mix}}(\beta_l) U_{\text{prob}}(\gamma_l), \quad (\text{A4})$$

where p is the number of QAOA layers, $U_{\text{mix}}(\beta)$ is a mixer Hamiltonian time-evolution operator (mixer operator), and $U_{\text{prob}}(\gamma)$ is the problem Hamiltonian time-evolution operator (phase operator).

The QAOA ansatz state can be thought of as being constructed from a time discretization and Trotterized approximation of the time-evolution operator for the time-dependent Hamiltonian used in quantum annealing (QA) [66],

$$H(t) = b(t)H_{\text{mix}} + c(t)H_{\text{prob}}, \quad (\text{A5})$$

where $b, c : [0, T] \rightarrow [0, 1]$ are monotonic functions, where b is decreasing from 1 to 0 and c is increasing from 0 to 1. By the adiabatic theorem of quantum mechanics, if the system evolves slowly enough from an initial ground state of the system (which is H_{mix} at $t = 0$), the system will remain in the ground state [67,68]. The mixer Hamiltonian H_{mix} is chosen to be

$$H_{\text{mix}} = - \sum_j X_j, \quad (\text{A6})$$

such that the initial state $|+\rangle^{\otimes N}$ is its ground state. The system time-ordered evolution operator can be written as

$$U(t) = \mathcal{T} \exp \left(-i \int_0^t H(t') dt' \right) \quad (\text{A7})$$

$$= \mathcal{T} \exp \left(-i \int_0^t [b(t')H_{\text{mix}} + c(t')H_{\text{prob}}] dt' \right), \quad (\text{A8})$$

where $\mathcal{T}$ is the time-ordering operator. The time evolution can be divided into $p + 1$ small discrete time steps, $\Delta t := T/(p + 1)$, with the assumption that the Hamiltonian is approximately constant during each step. This gives the relation

$$U(t + \Delta t) \approx e^{-iH(t)\Delta t} U(t), \quad (\text{A9})$$

which can be iterated backwards through time to give

$$U(T) \approx e^{-iH(T-\Delta t)\Delta t} e^{-iH(T-2\Delta t)\Delta t} \dots e^{-iH(0)\Delta t} \quad (\text{A10})$$

$$= \prod_{l=1}^{p+1} e^{-i[b_l H_{\text{mix}} + c_l H_{\text{prob}}]\Delta t}, \quad (\text{A11})$$

where $b_l := b((p + 1 - l)\Delta t)$ and $c_l := c((p + 1 - l)\Delta t)$. Applying the Trotter formula then results in

$$U(T) \approx \prod_{l=1}^{p+1} [e^{-i\Delta t b_l H_{\text{mix}}} e^{-i\Delta t c_l H_{\text{prob}}}], \quad (\text{A12})$$

By making the parameter substitutions $\beta_l := -b_l \Delta t$ and $\gamma_l := c_l \Delta t$, and noting that the $l = p + 1$ layer can be omitted since the initial state that it acts on is its ground state, we arrive at an expression that matches the form of the QAOA time-evolution operator shown in Eq. (A4),

$$U(\beta, \gamma) = \prod_{l=1}^p [e^{i\beta_l H_{\text{mix}}} e^{-i\gamma_l H_{\text{prob}}}], \quad (\text{A13})$$

Here, the mixer operator is

$$U_{\text{mix}}(\beta) = e^{i\beta H_{\text{mix}}} \quad (\text{A14})$$

$$= \prod_j e^{-i\beta X_j}, \quad (\text{A15})$$

where $\beta \in [0, \pi]$ due to symmetry, and the phase operator is

$$U_{\text{prob}}(\gamma) = e^{-i\gamma H_{\text{prob}}} \quad (\text{A16})$$

$$= \left(\prod_j e^{-i\gamma h_j Z_j} \right) \left(\prod_{(j,k)} e^{-i\gamma J_{jk} Z_j Z_k} \right), \quad (\text{A17})$$

where $\gamma \in [0, 2\pi]$.

The QAOA ansatz state $|\beta, \gamma\rangle$ defined in Eq. (A2) can be prepared using an ansatz circuit (or QAOA circuit) by implementing the operators using the following quantum gate representations. Operators in the form of $\exp(-i\alpha X_j/2)$ and $\exp(-i\alpha Z_j/2)$ can be performed using the single-qubit rotation gates $R_x^j(\alpha)$ and $R_z^j(\alpha)$, respectively where x and z are the Pauli axes of rotation, j is the target qubit, and α is the angle of rotation. Operators in the form of $\exp(-i\alpha Z_j Z_k/2)$ can be performed by applying the sequence of gates $\text{CNOT}_k^j \cdot R_z^k(\alpha) \cdot \text{CNOT}_k^j$, where CNOT_k^j is a controlled-not gate with control qubit j and target qubit k .

APPENDIX B: REFACTORIZING THE BPSP ISING HAMILTONIAN

By following the process in Sec. II C, we obtain the BPSP Ising Hamiltonian as

$$H_{\text{BPSP}} = \frac{1}{2} \sum_{i=1}^{|S|-1} (J'_{s_i, s_{i+1}} Z_{\Omega(s_i)} Z_{\Omega(s_{i+1})} + 1), \quad (\text{B1})$$

where S is the sequence of car instances, $J'_{s_i, s_{i+1}}$ is the coupling strength assigned by the relation between the adjacent cars $s_i, s_{i+1} \in S$, $\Omega(s_i)$ is the body type of the car instance s_i , and $Z_{\Omega(s_i)}$ is the Pauli-Z operator on the qubit that corresponds to body type $\Omega(s_i)$. In order to obtain the Ising Hamiltonian shown in Eq. (7) that directly represents an Ising graph, the above expression can be simplified by factoring the operator pairs $Z_{\Omega(s_i)} Z_{\Omega(s_{i+1})}$ in the sum such that the sum is over the pairs of adjacent body types $b_i \in B$ in the sequence, instead of the adjacent car instances. The steps for this simplification

are as follows, noting that $[c]$ is the Iverson bracket of the statement c (1 when true, 0 otherwise),

$$H_{\text{BPSP}} = \frac{1}{2} \sum_{k=1}^{|B|} \sum_{j=1}^k \left(\sum_{i=1}^{|S|-1} J'_{s_i, s_{i+1}} [(\Omega(s_i), \Omega(s_{i+1})) = (b_j, b_k)] \right) Z_{b_j} Z_{b_k} + \frac{|S| - 1}{2} \quad (\text{B2})$$

$$= \frac{1}{2} \left(\sum_{(i,j) \in E} J_{ij} Z_i Z_j + \sum_{k=1}^{|B|} [\exists l \in \mathbb{N}, (\Omega(s_l), \Omega(s_{l+1})) = (b_k, b_k)] \right) + \frac{|S| - 1}{2} \quad (\text{B3})$$

$$= \frac{1}{2} \sum_{(i,j) \in E} J_{ij} Z_i Z_j + \frac{A + |S| - 1}{2} \quad (\text{B4})$$

$$= \frac{1}{2} \sum_{(i,j) \in E} J_{ij} Z_i Z_j + \frac{C}{2}, \quad (\text{B5})$$

where B is the set of body types, b_i is the body type corresponding to qubit i , s_i is the i th car instance in the sequence S , $J'_{s_i, s_{i+1}}$ is the coupling strength assigned by the relation between the adjacent cars $s_i, s_{i+1} \in S$, J_{ij} is the total coupling strength between the pair of qubits corresponding to body types i and j , E are the edges in the Ising Hamiltonian's corresponding Ising graph with nonzero coupling, A is the

number of adjacent pairs of cars in the sequence S with the same body type, $\mathbb{N}$ denotes the set of natural numbers, and $C = A + |S| - 1$ specifies an integer constant. Equation (B3) is obtained by summing over the set of edges in the Ising graph while separating the $j = k$ terms (where adjacent cars have the same body type) and relabelling the indices such that i and j refer to body types that correspond to nodes in an Ising graph.

-
- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [2] S. Brandhofer, D. Braun, V. Dehn, G. Hellstern, M. Hüls, Y. Ji, I. Polian, A. Singh Bhatia, and T. Wellens, Benchmarking the performance of portfolio optimization with QAOA, *Quantum Inf. Proc.* **22**, 25 (2022).
 - [3] M. T. Khumalo, H. A. Chieza, K. Prag, and M. Woolway, An investigation of IBM quantum computing device performance on combinatorial optimisation problems, *Neural Computing Appl.* **37**, 611 (2025).
 - [4] M. Streif, S. Yarkoni, A. Skolik, F. Neukart, and M. Leib, Beating classical heuristics for the binary paint shop problem with the quantum approximate optimization algorithm, *Phys. Rev. A* **104**, 012403 (2021).
 - [5] P. Kl. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, Improving variational quantum optimization using CVaR, *Quantum* **4**, 256 (2020).
 - [6] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028*.
 - [7] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, A review on quantum approximate optimization algorithm and its variants, *Phys. Rep.* **1068**, 1 (2024).
 - [8] A. Abbas, A. Ambainis, B. Augustino, A. Bärtschi, H. Buhrman, C. Coffrin, G. Cortiana, V. Dunjko, D. J. Egger, B. G. Elmegreen, *et al.*, Challenges and opportunities in quantum optimization, *Nat. Rev. Phys.* **6**, 718 (2024).
 - [9] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings of 35th Annual Symposium on Foundations of Computer Science* (IEEE, Santa Fe, NM, 1994), pp. 124–134.
 - [10] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing* (ACM, Pennsylvania, Philadelphia, 1996), pp. 212–219.
 - [11] D. S. Johnson, G. Gutin, L. A. McGeoch, A. Yeo, W. Zhang, and A. Zverovitch, Experimental analysis of heuristics for the ATSP, *The Traveling Salesman Problem and its Variations*, edited by G. Gutin and A. P. Punnen, Combinatorial Optimization, Vol. 12 (Springer, Boston, MA, 2007), pp. 445–487.
 - [12] D. Applegate, W. Cook, and A. Rohe, Chained Lin-Kernighan for large traveling salesman problems, *INFORMS J. Comput.* **15**, 82 (2003).
 - [13] K. Helsgaun, An effective implementation of the Lin-Kernighan traveling salesman heuristic, *Eur. J. Oper. Res.* **126**, 106 (2000).
 - [14] S. Arora, Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems, *J. ACM* **45**, 753 (1998).
 - [15] D. S. Johnson and L. A. McGeoch, The traveling salesman problem: A case study, in *Local Search in Combinatorial Optimization*, edited by E. H. L. Aarts and J. K. Lenstra (John Wiley and Sons, London, 1997), pp. 215–310.
 - [16] M. L. Fredman, D. S. Johnson, L. A. McGeoch, and G. Ostheimer, Data structures for traveling salesmen, *J. Algorithms* **18**, 432 (1995).
 - [17] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to variational quantum optimization from symmetry protection, *Phys. Rev. Lett.* **125**, 260505 (2020).
 - [18] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).

- [19] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nat. Commun.* **12**, 6961 (2021).
- [20] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [21] R. Zhang, Environment-aware production scheduling for paint shops in automobile manufacturing: A multi-objective optimization approach, *Int. J. Environ. Res. Public Health* **15**, 32 (2018).
- [22] P. Bonsma, T. Epping, and W. Hochstättler, Complexity results on restricted instances of a paint shop problem for words, *Discrete Appl. Math.* **154**, 1335 (2006).
- [23] F. Meunier and A. Sebő, Paintshop, odd cycles and necklace splitting, *Discrete Appl. Math.* **157**, 780 (2009).
- [24] M. R. Garey, D. S. Johnson, and L. Stockmeyer, Some simplified NP-complete problems, in *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, STOC '74 (Association for Computing Machinery, Seattle, Washington, 1974), pp. 47–63.
- [25] J. Wurtz and D. Lykov, Fixed-angle conjectures for the quantum approximate optimization algorithm on regular MaxCut graphs, *Phys. Rev. A* **104**, 052419 (2021).
- [26] A. Ozaeta, W. van Dam, and P. L. McMahon, Expectation values from the single-layer quantum approximate optimization algorithm on Ising problems, *Quantum Sci. Technol.* **7**, 045036 (2022).
- [27] M. Streif and M. Leib, Training the quantum approximate optimization algorithm without access to a quantum processing unit, *Quantum Sci. Technol.* **5**, 034008 (2020).
- [28] R. Shaydulin, K. Marwaha, J. Wurtz, and P. C. Lotshaw, QAOAKit: A toolkit for reproducible study, application, and verification of the QAOA, in *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)* (IEEE, St. Louis, MO, 2021), pp. 64–71.
- [29] R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, Parameter transfer for quantum approximate optimization of weighted maxcut, *ACM Trans. Quantum Comput.* **4**, 1 (2023).
- [30] V. Vijendran, D. E. Koh, P. K. Lam, and S. M. Assad, Classical and quantum heuristics for the binary paint shop problem, [arXiv:2509.15294](https://arxiv.org/abs/2509.15294).
- [31] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [32] IBM ILOG CPLEX, V12.1: User's manual for CPLEX, Int. Bus. Mach. Corp. **46**, 157 (2009).
- [33] M. Karpinski, Approximability of the minimum bisection problem: An algorithmic challenge, in *Mathematical Foundations of Computer Science* (Springer, Berlin, Heidelberg, 2002), pp. 59–67.
- [34] T. Epping, W. Hochstättler, and P. Oertel, Complexity results on a paint shop problem, *Discrete Appl. Math.* **136**, 217 (2004).
- [35] S. Yarkoni, A. Alekseyenko, M. Streif, D. Von Dollen, F. Neukart, and T. Bäck, Multi-car paint shop optimization with quantum annealing, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Broomfield, CO, 2021), pp. 35–411.
- [36] Y.-H. Han, C. Zhou, B. Bras, L. McGinnis, C. Carmichael, and P. J. Newcomb, Paint line color change reduction in automobile assembly through simulation, in *Winter Simulation Conference* (2003), Vol. 2, p. 1204.
- [37] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case, [arXiv:2004.09002](https://arxiv.org/abs/2004.09002).
- [38] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples, [arXiv:2005.08747](https://arxiv.org/abs/2005.08747).
- [39] J. Basso, D. Gamarnik, S. Mei, and L. Zhou, Performance and limitations of the QAOA at constant levels on large sparse hypergraphs and spin glass models, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Denver, CO, 2022), pp. 335–343.
- [40] M. B. Hastings, Classical and quantum bounded depth approximation algorithms, [arXiv:1905.07047](https://arxiv.org/abs/1905.07047).
- [41] E. Bae and S. Lee, Recursive QAOA outperforms the original QAOA for the MAX-CUT problem on complete graphs, *Quantum Inf. Process.*, **23**, 78 (2024).
- [42] Y. J. Patel, S. Jerbi, T. Bäck, and V. Dunjko, Reinforcement learning assisted recursive QAOA, *EPJ Quantum Technol.* **11**, 6 (2024).
- [43] J. R. Finžgar, A. Kerschbaumer, M. J. A. Schuetz, C. B. Mendl, and H. G. Katzgraber, Quantum-informed recursive optimization algorithms, *PRX Quantum* **5**, 020327 (2024).
- [44] R. Kondo, Y. Sato, R. Raymond, and N. Yamamoto, Recursive quantum relaxation for combinatorial optimization problems, *Quantum* **9**, 1594 (2025).
- [45] B. Gulbahar, Majority voting with recursive QAOA and cost-restricted uniform sampling for maximum-likelihood detection in massive mimo, *IEEE Trans. Wireless Commun.* **24**, 2620 (2025).
- [46] J. Cook, S. Eidenbenz, and A. Bärttschi, The quantum alternating operator ansatz on maximum k-vertex cover, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Denver, CO, 2020), pp. 83–92.
- [47] M. Alam, A. Ash-Saki, and S. Ghosh, Accelerating quantum approximate optimization algorithm using machine learning, in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (IEEE, Grenoble, France, 2020), pp. 686–689.
- [48] M. M. Wauters, E. Panizon, G. B. Mbeng, and G. E. Santoro, Reinforcement-learning-assisted quantum optimization, *Phys. Rev. Res.* **2**, 033446 (2020).
- [49] O. Amosy, T. Danzig, O. Lev, E. Porat, G. Chechik, and A. Makmal, Iteration-free quantum approximate optimization algorithm using neural networks, *Quantum Mach. Intell.* **6**, 38 (2024).
- [50] A. Galda, X. Liu, D. Lykov, Y. Alexeev, and I. Safro, Transferability of optimal QAOA parameters between random graphs, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Broomfield, CO, 2021), pp. 171–180.
- [51] P. C. Lotshaw, T. S. Humble, R. Herrman, J. Ostrowski, and G. Siopsis, Empirical performance bounds for quantum approximate optimization, *Quantum Inf. Proc.* **20**, 403 (2021).
- [52] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance,

- mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [53] J. Wurtz and P. Love, MaxCut quantum approximate optimization algorithm performance guarantees for $p > 1$, *Phys. Rev. A* **103**, 042612 (2021).
- [54] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [55] F. G. S. L. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances, [arXiv:1812.04170](https://arxiv.org/abs/1812.04170).
- [56] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for MaxCut: A fermionic view, *Phys. Rev. A* **97**, 022304 (2018).
- [57] S. A. Hadfield, *Quantum Algorithms for Scientific Computing and Approximate Optimization* (American Physical Society, 2018).
- [58] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, *Quantum* **6**, 759 (2022).
- [59] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, Learning to optimize variational quantum circuits to solve combinatorial problems, *AAAI* **34**, 2367 (2020).
- [60] I. Lyngfelt and L. García-Álvarez, Symmetry-informed transferability of optimal parameters in the quantum approximate optimization algorithm, *Phys. Rev. A* **111**, 022418 (2025).
- [61] J. Villanueva, G. J. Mooney, B. R. Bardhan, J. Ghosh, C. D. Hill, and L. C. L. Hollenberg, Hybrid quantum optimization in the context of minimizing traffic congestion, [arXiv:2504.08275](https://arxiv.org/abs/2504.08275).
- [62] A. Dang, C. D. Hill, and L. C. L. Hollenberg, Optimising matrix product state simulations of Shor's algorithm, *Quantum* **3**, 116 (2019).
- [63] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117 (2014).
- [64] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [65] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [66] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [67] M. Born and V. Fock, Beweis des adiabatensatzes, *Z. Phys.* **51**, 165 (1928).
- [68] T. Kato, On the adiabatic theorem of quantum mechanics, *J. Phys. Soc. Jpn.* **5**, 435 (1950).