

Quantum-Informed Recursive Optimization Algorithms

Jernej Rudi Finžgar^{1,2,*}, Aron Kerschbaumer^{1,2,†}, Martin J.A. Schuetz^{3,4}, Christian B. Mendl^{2,5}
and Helmut G. Katzgraber³

¹BMW AG, Munich, Germany

²Department of Computer Science, Technical University Munich, School of CIT, Garching, Germany

³Amazon Quantum Solutions Lab, Seattle, Washington, USA

⁴AWS Center for Quantum Computing, Pasadena, California, USA

⁵Technical University of Munich, Institute for Advanced Study, Garching, Germany



(Received 21 September 2023; accepted 21 February 2024; published 3 May 2024)

We propose and implement a family of quantum-informed recursive optimization (QIRO) algorithms for combinatorial optimization problems. Our approach leverages quantum resources to obtain information that is used in problem-specific classical reduction steps that recursively simplify the problem. These reduction steps address the limitations of the quantum component (e.g., locality) and ensure solution feasibility in constrained optimization problems. Additionally, we use backtracking techniques to further improve the performance of the algorithm without increasing the requirements on the quantum hardware. We showcase the capabilities of our approach by informing QIRO with correlations from classical simulations of shallow circuits of the quantum approximate optimization algorithm, solving instances of maximum independent set and maximum satisfiability problems with hundreds of variables. We also demonstrate how QIRO can be deployed on a neutral atom quantum processor to find large independent sets of graphs. In summary, our scheme achieves results comparable to classical heuristics even with relatively weak quantum resources. Furthermore, enhancing the quality of these quantum resources improves the performance of the algorithms. Notably, the modular nature of QIRO offers various avenues for modifications, positioning our work as a template for a broader class of hybrid quantum-classical algorithms for combinatorial optimization.

DOI: [10.1103/PRXQuantum.5.020327](https://doi.org/10.1103/PRXQuantum.5.020327)

I. INTRODUCTION

Quantum optimization has been identified as a promising area of research towards practical quantum advantage. On noisy intermediate-scale quantum (NISQ) devices, much effort has been dedicated to studying hybrid quantum-classical algorithms such as the quantum approximate optimization algorithm (QAOA) [1]. Importantly, QAOA is a *local* algorithm. This means that at any constant circuit depth, only qubits that are separated by less than a certain distance in the interaction graph of the optimization problem are able to communicate. Together with

the solution space properties of certain optimization problems such as the \mathbb{Z}_2 symmetry, or the overlap gap property [2], locality was shown to severely limit the performance of QAOA [3–6].

Because introducing *nonlocal* updates in quantum algorithms comes with additional hardware requirements, applying nonlocal updates classically has been proposed as an alternative in recursive QAOA (RQAOA) [3,7]. Here, values of variables are iteratively frozen by rounding the correlations between variables as measured in the quantum state prepared by QAOA. Hence, as variables are removed from the optimization problem, the distances between nodes in the interaction graph are reduced iteratively. As such, the nonlocal effects introduced via the new connections between previously unconnected nodes counterbalance the locality inherent to QAOA.

Building upon RQAOA, here we propose a family of hybrid quantum-classical algorithms, dubbed *quantum-informed recursive optimization* (QIRO). In QIRO, information generated by quantum resources is used to recursively reduce the size of the optimization problem by means of problem-specific classical optimization routines;

*jernerj-rudi.finzgar@tum.de

†Present address: Institute of Science and Technology Austria (ISTA), Am Campus 1, 3400 Klosterneuburg, Austria.

‡These authors contributed equally.

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.

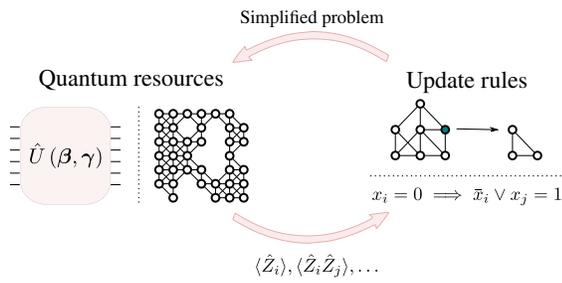


FIG. 1. Schematic visualization of the core principles of the *quantum-informed recursive optimization* (QIRO) algorithm presented in this work. Quantum resources (e.g., QAOA) are used to obtain information, e.g., in the form of one-point and two-point correlations. This information is then used to simplify the problem through problem-specific classical update rules. Here we exemplify the update rules with simple examples for the maximum independent set and satisfiability problems—details on the update rules can be found in Sec. III B. The (simplified) problem obtained by means of the classical update rules is then used to restart the cycle. The algorithm terminates when the problem has been fully simplified.

see Fig. 1 for a schematic illustration. This scheme allows us to leverage decades of research in (classical) combinatorial optimization [8], and tailor the classical subroutines to the particular optimization problem of interest, thereby enhancing the algorithm’s performance.

Moreover, the inclusion of problem-specific update rules comes with additional benefits. It allows us to broaden the scope of our algorithm beyond gate-based architectures to analog devices, as shown in our experiments with the QuEra Aquila neutral atom analog quantum processor. This approach is reminiscent of previous approaches leveraging spin-freezing schemes on quantum annealers and in various Monte Carlo methods [9]. Furthermore, for problems with hard constraints, the update rules can offer options to enforce feasibility by design, even in the presence of noise.

Finally, we propose *backtracking* as a strategy to further improve the performance of QIRO, by attempting to identify and correct nonideal decisions made at earlier stages of the algorithm. Backtracking provides a way to enhance algorithmic performance without necessitating an increase in circuit depth, as is commonplace in quantum optimization algorithms. Taken together, the robustness of QIRO, its applicability to analog devices, and the possibility to improve its efficacy without demanding better quantum resources make QIRO a promising NISQ algorithm.

In this work we provide concrete implementations for two paradigmatic NP-hard combinatorial optimization problems, namely, the maximum independent set and maximum satisfiability. However, our work should be seen mainly as a proof-of-concept demonstration of a general template for designing hybrid quantum-classical algorithms for combinatorial optimization. We study the

performance of QIRO by means of large-scale simulations of low-depth QAOA circuits with up to two hundred variables. Furthermore, we use QIRO to solve the maximum independent set problem on the QuEra Aquila neutral atom quantum device accessed via Amazon Braket. Finally, we compare the performance of QIRO against commonly used classical and quantum optimization techniques. This allows us to assess the role of the quantum resources used in QIRO, and showcase the value of classical subroutines.

The remainder of the paper is organized as follows. In Sec. II we introduce the two problem classes studied in this work, in Sec. III we provide details about QIRO, and in Sec. IV we study the performance of QIRO on the maximum independent set and maximum satisfiability problems. We discuss the implications of our findings in Sec. V and we conclude by suggesting potential extensions to our framework in Sec. VI.

II. PROBLEM CLASSES

Because *problem-specific* update rules are at the core of QIRO, we first introduce the specific optimization problems considered in this work. We formalize these problems, and provide mappings to quantum Hamiltonians \hat{H}_c , which can in turn be implemented on different quantum devices [10,11].

A. Maximum independent set

The maximum independent set (MIS) problem is a paradigmatic NP-hard combinatorial optimization problem. It is of commercial relevance in areas such as network design or traffic optimization [12]. For a graph $G = (V, E)$ with vertex set V and an edge set E , the MIS problem amounts to finding the largest subset $S \subseteq V$ such that no two vertices in S are adjacent [see Fig. 5(a) below for an example]. For a binary vector $\mathbf{x} \in \{0, 1\}^{|V|}$, finding the MIS is equivalent to finding the ground state of the (classical) Hamiltonian

$$H(\mathbf{x}) = - \sum_{i \in V} x_i + \lambda \sum_{(i,j) \in E} x_i x_j. \quad (1)$$

Here, $\lambda > 1$ is a penalty term to enforce the independence constraint [10]. However, one should be mindful of the pitfalls of enforcing the independence constraint as a soft penalty term, because low-energy states (other than the ground states) of $H(\mathbf{x})$ might not necessarily obey the independence constraint. Performing the mapping $z_i = 2x_i - 1$ to Ising variables $z_i = \pm 1$, and promoting the spin variables to quantum operators $z_i \rightarrow \hat{Z}_i$ translates Eq. (1) to a quantum Hamiltonian suitable for quantum devices [10].

Moreover, the MIS problem on a family of graphs called unit disk graphs (UDGs) has been shown to have a natural mapping to neutral atom quantum devices [13,14]. There,

the Rydberg blockade mechanism—which prevents simultaneous excitations of nearby atoms—enforces the independence constraint. Finding the MIS of a general UDG, despite being a specialized family of graphs, remains an NP-hard problem [13] and has commercial applications (e.g., in network design [15]). For additional details, we refer the reader to Appendix D.

B. Maximum satisfiability

The satisfiability (SAT) problem is arguably one of the best-studied combinatorial optimization problems and was the first problem proven to be NP-complete [16]. Given a propositional logic formula ϕ of n Boolean variables $x_i \in \{0, 1\}$, the SAT problem is a decision problem that asks for the existence of an assignment of the variables x_i that satisfies ϕ . Each clause can then be written as

$$\tilde{C}_j = \bigvee_{k=1}^{K_j} \ell_{jk},$$

where literals ℓ_{jk} are either a variable x_{j_k} or its negation \bar{x}_{j_k} for $j_k \in \{1, \dots, n\}$, and K_j is the length of the j th clause. For a fixed $K_j = K$, the corresponding problem of whether or not the Boolean formula with m clauses

$$\phi = \bigwedge_{j=1}^m \tilde{C}_j$$

has a satisfying assignment is referred to as the K -SAT problem.

Here, we focus on the optimization version of satisfiability, referred to as maximum satisfiability (MAX SAT). In MAX SAT, the goal is to find an assignment of the variables x_i that minimizes the number of violated clauses. The MAX- K -SAT problem deals with propositional logic formulae with *at most* K literals per clause. An important parameter of a (MAX-) K -SAT problem is the clause-to-variable ratio $\alpha := m/n$, which strongly influences the algorithmic hardness of random problem instances [17]. Research indicates that there is a phase transition in computational complexity [18]. While for the K -SAT problem an easy-hard-easy transition is observed [19,20], MAX- K -SAT exhibits a transition from easy underconstrained to hard overconstrained problems at a certain α_c [21]. In what follows we limit ourselves to $K = 2$ (clauses contain either one or two literals), as MAX-2-SAT is already NP-hard [22], despite its decision version (2-SAT) being solvable in linear time [23]. Here, the phase transition occurs at $\alpha_c = 1$ [24].

We use the fact that, for Boolean (binary) variables $\bar{x}_i \equiv (1 - x_i)$, solving the MAX-2-SAT problem for a given formula ϕ is equivalent to finding a ground state of the

(classical) Hamiltonian

$$H(\phi) = \sum_{j=1}^m C_j = \sum_{j=1}^m \prod_{k=1}^{K_j} (1 - \ell_{jk}). \quad (2)$$

It is important to keep in mind that the ℓ_{jk} are nothing but the Boolean variables x_{j_k} (or their negations \bar{x}_{j_k}). Hence, by using the transformation into Ising variables outlined in the case of the MIS problem in Sec. II A, we can likewise encode the MAX-2-SAT problem onto quantum devices.

III. ALGORITHM

We now provide details about QIRO. Essentially, in QIRO quantum resources are utilized to inform update rules that recursively simplify the problem at hand (see Fig. 1). The idea is a generalization of previous work, in particular RQAOA [3,7] (also described in Appendix B). However, in contrast to RQAOA, the update steps used in QIRO are problem specific. This comes with several advantages and facilitates cross-pollination with ideas from classical optimization [8].

At each iteration of the QIRO algorithm, for a given problem Hamiltonian, we first prepare a low-energy quantum state, generically in the form of a superposition of low-energy candidate solutions to the optimization problem. Next, we use this quantum state to extract information (e.g., correlations between variables), which is then passed to a classical update step with the goal of simplifying the problem. We design this update step to complement the quantum part of the algorithm by performing operations better suited to classical hardware. This may include nonlocal steps that address the identified limitations of (local) quantum algorithms, or enforcing hard constraints. As such, the classical update step (and the specific design thereof) is crucial to the performance of QIRO, and is typically problem specific. As such, it might need to be modified by the end user, depending on their needs. However, templates for the update steps may be borrowed from the literature on classical combinatorial optimization [25]. The procedure of quantum state preparation and the subsequent classical update step is repeated until the size of the problem is reduced sufficiently, such that, for example, we are able to solve it exactly by a brute-force search or other means if the reduction does not yield problem sizes amenable to exact solvers. Finally, we propose the use of backtracking to identify and rectify problem reductions that led to an unfavorable outcome, thus providing an additional strategy to obtain enhanced solutions without requiring deeper quantum circuits.

In what follows, we provide details on the individual components of QIRO. We begin by describing the quantum state preparation methods utilized in this work, followed by the proposed update rules for the MIS and MAX-2-SAT

problems. Next, we introduce the general strategy of backtracking and describe the particular implementation used in this work. Finally, we combine the building blocks to present the complete QIRO algorithm.

A. Quantum state preparation

It is desirable that the quantum state that is prepared at the beginning of each iteration has a low expectation value with respect to the cost Hamiltonian \hat{H}_c . Therefore, we hope that the information guiding the classical reduction steps is extracted from (a superposition of) configurations corresponding to good candidate solutions to the optimization problem. A plethora of protocols for preparing low-energy states may be found in the literature; in this manuscript we consider QAOA [1] and adiabatic state preparation [26,27], as outlined next.

1. QAOA

Let us first introduce QAOA. For both the MAX-2-SAT and MIS problems considered in this manuscript, the cost Hamiltonian can be written in the general (Ising) form $\hat{H}_c = \sum_{i,j} J_{ij} \hat{Z}_i \hat{Z}_j + \sum_i h_i \hat{Z}_i$, where we have omitted any constant terms. Using the cost Hamiltonian \hat{H}_c from Sec. II, and defining the mixer Hamiltonian as $\hat{H}_{\text{mix}} = -\sum_{i=1}^n \hat{X}_i$, the QAOA circuit is given by an alternating application of the time evolution generated by these two operators. Specifically, taking the initial state as $|\psi_0\rangle := |+\rangle^{\otimes n}$, the state produced with QAOA at depth p is

$$|\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle = e^{-i\beta_p \hat{H}_{\text{mix}}} e^{-i\gamma_p \hat{H}_c} \dots e^{-i\beta_1 \hat{H}_{\text{mix}}} e^{-i\gamma_1 \hat{H}_c} |\psi_0\rangle.$$

The $2p$ parameters $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ are then determined by a classical optimization routine such that $\langle \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) | \hat{H}_c | \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) \rangle$ is minimized.

In this work, we predominantly focus on the lowest QAOA depth ($p = 1$), with the goal to establish a lower bound for the performance of QIRO. It is natural to expect that the improved quality of the quantum information extracted from higher-depth ($p > 1$) QAOA circuits (in the absence of noise) would enhance the performance of the algorithm. Our $p > 1$ experiments on small problem instances shown in Appendix E corroborate this intuition.

For the state produced by QAOA at $p = 1$, the expectation values of the single- and two-body terms comprising \hat{H}_c can be calculated classically in polynomial time, enabling us to perform numerical simulations of systems with several hundred variables [3,7,28]. We note that there is no method known (to us) how this could be done efficiently for $p > 1$. For details about our implementation, we refer the reader to Appendix A.

2. Adiabatic state preparation

We next turn our attention to adiabatic state preparation. We present a general description of the protocol

and defer the details pertaining to our implementation of the MIS problem on neutral atom quantum processors to Appendix D.

At the beginning of an adiabatic protocol, we start in an easy-to-prepare ground state of a simple Hamiltonian, e.g., \hat{H}_{mix} . We then slowly transform the Hamiltonian from \hat{H}_{mix} to \hat{H}_c . If the rate of change is slow enough compared to the smallest instantaneous spectral gap (between the instantaneous ground and first excited states), the adiabatic theorem guarantees that at the end of the protocol, the ground state of \hat{H}_c is obtained [26,27]. Inspired by adiabatic protocols, the quantum annealing computational paradigm has been the subject of widespread interest in recent decades [29]. In quantum annealing one forgoes the strict requirement of adiabaticity and, in general, uses analog protocols to prepare near-ground states. These can, in turn, be used to inform the classical update steps.

B. Update rules

We now discuss example update rules for the considered optimization problems. We begin by preparing a low-energy quantum state $|\psi\rangle$ from which information is distilled. In QIRO this information comes in the form of one-point $\langle \psi | \hat{Z}_i | \psi \rangle$ and two-point correlations $\langle \psi | \hat{Z}_i \hat{Z}_j | \psi \rangle$, which we store in matrix M (see Algorithm 1 below). Potential generalizations, such as using expectation values of other observables, are left for future research.

If the lowest depth ($p = 1$) of QAOA is used to generate these correlations, the only nonzero values of M in an ideal numerical simulation as described in Appendix A (i.e., free of shot and environmental noise) correspond to pairs of variables that are connected in the original interaction graph, i.e., to pairs of variables with nonzero coupling coefficients J_{ij} . Therefore, we design our update rules to only consider correlations between such connected variables. Inspired by RQAOA [3,7], we next use these correlations to successively simplify the problem with each update step. However, in contrast to RQAOA, the QIRO update steps are problem specific in that they exploit the structure of the given optimization problem. As such, at any stage of the algorithm, the problem that remains to be solved is a valid instance of the same class of optimization problems. This will be crucial when deploying QIRO to solve the MIS problem on neutral atom quantum hardware.

In the two specific cases detailed in the remainder of this section, only the largest correlation in terms of its absolute value is used to inform the update step, with ties broken at random. Intuitively, the largest correlation corresponds to a variable (pair) that appears in a certain configuration with high probability in the low-energy state $|\psi\rangle$. Thus, it is reasonable to use this information to inform the update step. For example, if a large positive correlation $\langle \hat{Z}_i \hat{Z}_j \rangle \approx 1$ is measured, it is likely that most of the low-energy configurations have x_i and x_j take the same

ALGORITHM 1. Reduce.

Input: Problem P , partial solution S .

Output Simplified prob. P' , updated partial solution S' .

- 1: Prepare low-energy quantum state $|\psi\rangle$.
 - 2: Store correlations in $M \in \mathbb{R}^{n \times n}$:
 $\forall i \in [n] : M_{ii} = \langle \psi | \hat{Z}_i | \psi \rangle$,
 $\forall (i, j) \in [n] \times [n] \text{ s.t. } J_{ij} \neq 0 : M_{ij} = \langle \psi | \hat{Z}_i \hat{Z}_j | \psi \rangle$.
 - 3: $P', S' \leftarrow \text{Simplification}(P, S, M)$
 \triangleright Problem-specific simplification is informed by the correlations stored in M .
 - 4: **return** P', S'
-
-

value. Generalizations of this deterministic scheme to randomized (e.g., evolutionary) schemes will be studied in the future.

As we lay out next, these update steps are carried out in the original formulation of the respective optimization problems, facilitating the adoption of techniques from the extensive literature on classical combinatorial optimization. We call this update step `Reduce` and describe it in Algorithm 1.

The problem-specific updates employed for the optimization problems considered in this work constitute the `Simplification` function in Algorithm 1, and are described next. However, we emphasize that the design choices we make are by no means the only option, with many modifications and extensions possible.

1. MIS simplification

At the beginning of the `Simplification` routine for MIS, we find the entry of M with the largest absolute value. Depending on the sign of the correlation with the largest absolute value and on whether it lies on the diagonal (corresponding to a one-point correlation) or not (corresponding to a two-point correlation), we perform a different simplification. The simplifications are designed such that the correlation-informed reductions are consistent with the independence constraint. The four cases are visualized in Fig. 2 and described next.

- (a) If $M_{ii} \geq 0$ was selected, we set the i th vertex to be in the independent set (IS). We then remove all vertices connected to the i th node from the graph, as including them would violate the independence constraint.
- (b) If $M_{ii} < 0$ was selected, we remove the i th vertex from the graph.
- (c) If $M_{ij} > 0$ [for $(i, j) \in E$] was selected, we remove both nodes from the graph—this is because the only positively correlated assignment of *connected* vertices consistent with the independence constraint is not to include either node in the IS.
- (d) If $M_{ij} < 0$ [for $(i, j) \in E$] was selected, we remove nodes that are connected to both node i and node

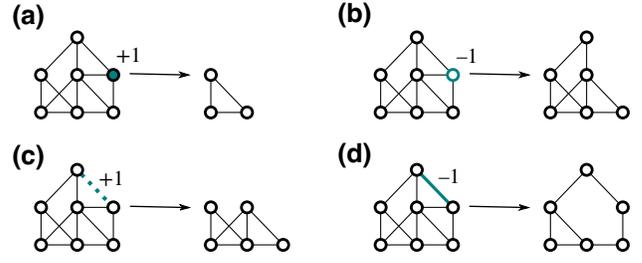


FIG. 2. MIS update rules described in the main text carried out on an example graph. Panels (a) and (b) show reduction rules for the case of a positive and negative *one-point* correlation, respectively. Analogously, panels (c) and (d) show the respective reduction rules for the case of a positive and negative *two-point* correlation.

j simultaneously, i.e., we remove every node k for which $(k, i) \in E$ and $(k, j) \in E$. Intuitively, we know that if the variables are negatively correlated, one of them will be assigned to be in the IS. Thus, we only remove variables that are connected to both of them simultaneously.

If no nodes were removed from the graph using the described simplifications [only possible in case (d)], we repeat the procedure using the next largest correlation in terms of its absolute value. Finally, we identify all connected components (i.e., connected subgraphs that are not part of any other connected subgraph) of the graph. If the connected component contains less than a certain number of vertices n_c (for concreteness, we use $n_c = 15$ here), we compute its maximum independent set by a brute-force search, and remove it from the graph. This step is mostly performed in an effort to save resources when the algorithm is deployed on real quantum hardware.

Crucially, the rules delineated above explicitly ensure that the independence constraint is obeyed. Hence, the solutions obtained by repeated applications of `Simplification` are definitely feasible. This holds irrespective of the quality of the (quantum) information used to perform the simplification, thus bestowing the scheme with the resilience required in the NISQ era. However, although the update rules improve the performance of the algorithm by ensuring solution feasibility, the quality of the obtained solutions ultimately hinges on the quality of the supplied correlations.

Interestingly, in contrast to the RQAOA update rules the QIRO-MIS update rules proposed here are *local* according to the definition of locality in Sec. I, as no new connections between variables are created. However, the RQAOA update rules lead to intermediate Hamiltonians that generically do not take the form of an MIS Hamiltonian as in Eq. (1). Conversely, at each stage of the QIRO algorithm the reduced problem is a valid MIS problem (see also Fig. 2). Thus, if one can encode the original MIS problem onto

an analog quantum device (e.g., a neutral atom quantum processor), the same can be done with intermediate, simplified problems. As such, one might be able to harness the nonlocal effects generated by the quantum many-body time evolution on a neutral atom quantum processor, circumventing the issues of local quantum algorithms. Previously, such adiabatic protocols have been shown to generate solutions surpassing those obtained by running QAOA, even at depths higher than $p = 1$ used in this manuscript [14].

2. MAX-2-SAT simplification

In the MAX-2-SAT case the `Simplification` step likewise begins by identifying the entry of M with the largest absolute value. Depending on whether the largest value corresponds to a one-point or a two-point correlation, we apply a different reduction.

- (a) If M_{ii} was selected, we set the variable $x_i \rightarrow [\text{sign}(M_{ii}) + 1]/2$, i.e., a positive (negative) one-point correlation corresponds to assigning x_i to TRUE (FALSE).
- (b) If M_{ij} (for $i < j$) was selected, we replace $x_i \rightarrow x_j$ if $\text{sign}(M_{ij}) = 1$, and $x_i \rightarrow \bar{x}_j$ if $\text{sign}(M_{ij}) = -1$. Intuitively, this update step captures the relationship between variables x_i and x_j as inferred from the correlation. If there is a positive (negative) correlation, we assign x_i and x_j the same (opposite) values.

We note that only the two-point correlation reduction step described in (b) constitutes a nonlocal update because it potentially introduces a new link between previously unconnected variables. Because we only consider formulae in conjunctive normal form (i.e., conjunctions of disjunctions), setting a certain variable to the value required by the clause is sufficient to satisfy the clause (see the example update step in Fig. 1). Hence, we can remove the clause from consideration, thus simplifying the problem. Specifically, all clauses in which a given literal evaluates to TRUE can be removed. Conversely, if the literal evaluates to FALSE, only the literal itself can be removed from the clause. If a certain clause becomes empty (i.e., if both literals were removed from it), it corresponds to an *unsatisfied* clause.

Next, we apply inference rules, which is a standard technique in the field of satisfiability solvers [30,31]. Inference rules deduce information from the structure of the current Boolean formula, simplifying it by assigning values to certain variables and thus speeding up further computation. The rules guarantee that the optimal solutions of the current and simplified problems are of equal quality. The chosen inference rules (e.g., the pure literal rule [32]) are conceptually simple rules that incur a small computational overhead. Further details on the specific inference rules we employ are provided in Appendix C. Finally, if the number of remaining variables falls below n_c (we use $n_c = 10$ for

ALGORITHM 2. QIRO.

Input: Problem P .

Output: Complete solution S .

- 1: Initialize empty solution S .
 - 2: **while** $\text{size}(P) > 0$ **do**
 - 3: $P, S \leftarrow \text{Reduce}(P, S)$ ▷ See Algorithm 1.
 - 4: **end while**
 - 5: **return** S ▷ Return the solution.
-
-

the MAX-2-SAT problem), the solution of this simplified problem is found by a brute-force search of the remaining solution space.

C. Quantum-informed recursive optimization algorithm

We now possess all of the ingredients for the QIRO algorithm. At each step of QIRO, we begin by preparing a quantum state (as described in Sec. III A). We proceed by using the correlation information by following the problem-specific prescriptions from Sec. III B. The simple end-to-end QIRO procedure is summarized in Algorithm 2.

D. Backtracking

We now turn our attention to backtracking. QIRO—is specified in Algorithm 2—is a polynomial-time algorithm. Therefore, under common complexity theoretical assumptions, there exist instances of NP-hard problems that QIRO will not be able to solve to optimality. This means that some of the reduction steps performed by QIRO to arrive at a candidate solution might be wrong. A similar observation was made for RQAOA in Ref. [33], prompting the authors to consider reinforcement learning to steer RQAOA update steps. Here, we propose backtracking to derive improved solutions starting from an initial candidate solution obtained by QIRO. Through backtracking, we obtain a method to refine solutions by using the same quantum resources more frequently. This approach stands in stark contrast to most other quantum optimization algorithms, which usually require enhanced quantum resources (e.g., increasing the depth p in QAOA) to deliver improved solutions.

In this work, we only employ this strategy for the MAX-2-SAT problem. However, extending this approach to other optimization problems such as the MIS problem is straightforward.

As visualized in Fig. 3 the simplification process of QIRO can be represented by means of a search tree. Decisions (i.e., update steps) made by the algorithm correspond to arrows in the search tree, which lead to simplified problems represented by nodes. Following the update rules in Sec. III B (i.e., successively applying `Reduce`) traces out a path through the search tree (see the red arrows and

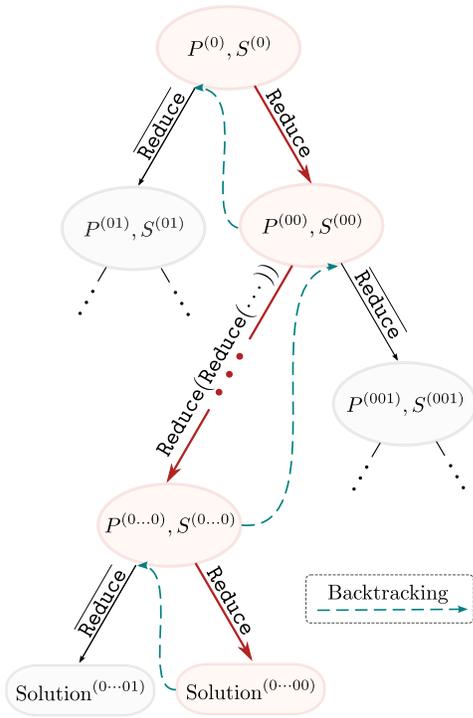


FIG. 3. Schematic visualization of the backtracking procedure used to obtain improved solutions. Starting from the original problem $P^{(0)}$ and an empty solution $S^{(0)}$, we perform a series of reductions (red arrows and nodes) that recursively simplify the problem, until the first candidate solution ($\text{Solution}^{(0\cdots 00)}$) is reached. This corresponds to applying QIRO as outlined in Algorithm 2. Backtracking, as visualized by the dashed teal arrows, then amounts to revisiting reduction steps made along the path. At the chosen ancestor node we have backtracked to, we can then make an alternative reduction, here denoted by $\overline{\text{Reduce}}$ (see the black arrows). Concretely, $\overline{\text{Reduce}}$ amounts to reversing the initially made decision, not needing extra calls to the quantum device. In this work, we limit ourselves to making a binary decision at each node. As such, we can label the reduced problem and partial solution at each node with a binary string encoding the path connecting the node to the original problem node. Each 0 (1) in the superscript labeling P and S corresponds to one application of Reduce ($\overline{\text{Reduce}}$) starting from the original problem.

nodes in Fig. 3), which eventually leads to a candidate solution (see $\text{Solution}^{(0\cdots 00)}$ in Fig. 3). Here, the superscript $(0\cdots 00)$ indicates that this is the solution obtained in the initial passage through the search tree.

We would like to consider alternative solutions, obtainable by deviating from the initial path through the search tree that led to the initial candidate solution $\text{Solution}^{(0\cdots 00)}$. To this end, we define the $\overline{\text{Reduce}}$ function, which reverses the decision previously made at the revisited node (i.e., partial problem and solution). We emphasize that this initial reversal of the previously made decision requires no additional measurements from the quantum device—we

ALGORITHM 3. QIRO + backtracking (QIRO + BT).

```

Input: Problem  $P$ .
Output: Complete solution  $S$ .
1: Initialize empty solution  $S$ .
2:  $L \leftarrow \{\}$   $\triangleright$  List to store partial problems & solutions.
3: while  $\text{size}(P) > 0$  do
4:   Append  $(P, S)$  to  $L$ .
5:    $P, S \leftarrow \text{Reduce}(P, S)$   $\triangleright$  See Algorithm 1.
6: end while
7: for  $P', S'$  in  $L$  do  $\triangleright$  Backtracking.
8:    $P', S' \leftarrow \overline{\text{Reduce}}(P', S'; \text{Solution}^{(0\cdots 00)})$ 
    $\triangleright$  Make alternative decision.
9:   while  $\text{size}(P') > 0$  do
10:     $P', S' \leftarrow \text{Reduce}(P', S')$ 
11:   end while
12:   Store the obtained solution.
13: end for
14: return Best solution found.

```

simply perform the opposite reduction to that we performed initially. For example, if at first the decision was made to assign the value of a variable $x_i = \text{TRUE}$, we reverse that decision and set $x_i = \text{FALSE}$ instead. Thus, we diverge from the original path through the search tree as visualized in Fig. 3. We note that $\overline{\text{Reduce}}$ requires information about the solution it is deviating from in order to be able to perform the reversed decision.

In principle, exploring the entire search tree essentially corresponds to a brute-force search of the complete solution space. Because we are interested in designing heuristic algorithms that run in polynomial time, we limit the permitted exploration. We find that a simple rule for which the number of QAOA circuit executions scales quadratically in the problem size is sufficient to produce satisfactory results for the considered problem sizes of MAX-2-SAT. However, more (or less) sophisticated approaches can be designed and tailored to the particular needs and resources of the end user. Next, we describe the specific backtracking strategy used in our simulations and refer the reader to Algorithm 3 for the pseudocode.

First, we generate the initial candidate solution $\text{Solution}^{(0\cdots 00)}$ by successively applying Reduce , as visualized in Fig. 3. Put differently, we first apply QIRO *without* backtracking. Next, we backtrack to each of the ancestor nodes of the initial candidate solution (as visualized by the dashed teal lines in Fig. 3). Each of the revisited parent nodes corresponds to a different simplified problem and a corresponding partial solution. The decision made in the initial passage through the search tree is then reversed (by applying $\overline{\text{Reduce}}$), whereby a new node in the search tree (e.g., $P^{(01)}, S^{(01)}$ in Fig. 3) is created. Here, each 0 (1) in the superscript labeling P and S corresponds to one application of Reduce ($\overline{\text{Reduce}}$) starting from the original problem. Subsequently, if the problem has not

been fully reduced yet, from this new node a new candidate solution is generated by repeatedly applying `Reduce`. In contrast to the initial application of `Reduce` the subsequent applications of `Reduce` require additional calls to the quantum device—see the definition of `Reduce` in Algorithm 1. In other words, each revisited node leads to a new candidate solution that is identical to the initial candidate solution in the values of the variables fixed before the revisited node. From the revisited node onward, the new candidate solution differs from the initial solution. Finally, the output of the algorithm is the best candidate solution found in this way. We note that the order in which the parent nodes are revisited is not important.

IV. RESULTS

Next, we present the results obtained by applying QIRO to random instances of the MIS and MAX-2-SAT problems. In numerical simulations we compare the performance of quantum algorithms (QIRO, RQAOA) and two prominent classical heuristics (simulated annealing, greedy algorithms). Of course, more complex tailored classical approaches can be used, such as the parallel tempering [34–36], borealis [37], and extremal optimization [38] heuristics, as well as a plethora of exact solvers benchmarked at the annual (MAX-)SAT competition [39].

Efficient classical simulations of $p = 1$ QAOA allow us to study the system size scaling on problem instances with various connectivities of the underlying interaction graph. Results using the lowest depth of QAOA should mainly be interpreted as a lower bound on the performance of QIRO. Furthermore, we investigate the robustness of QIRO and RQAOA against perturbations of the optimal parameters of the QAOA circuits generating the correlations. This allows us to emulate failures of the classical optimizer to find the optimal parameters. Notably, we also deploy QIRO on the QuEra Aquila neutral atom device (as available on Amazon Braket) to (approximately) solve the MIS problem on unit disk graphs with more than a hundred nodes.

A. MIS

We begin by presenting the results of QIRO and RQAOA as applied to the MIS problem. First, we consider the algorithmic performance as a function of the system size on the well-studied family of Erdős-Rényi graphs, with fixed average degrees [40]. As a classical comparison, we use two variants of a greedy algorithm. This choice is motivated by recent findings that information from quantum devices can be used to enhance the performance of greedy solvers [41]. First, we consider a simple *random greedy* algorithm, which iteratively assigns random nodes to the independent set, and removes nodes adjacent to the selected node such that no violations are possible. As the second comparison, we use the *minimal degree greedy*

algorithm, which iteratively assigns a node with the minimal degree to the independent set and likewise removes adjacent nodes to ensure the validity of the final solution [42]. We note that the minimal degree greedy algorithm was designed chiefly for sparse graphs (i.e., of low average degree). This is indicated by the performance guarantee of the greedy algorithm, which is inversely proportional to the graph’s average degree [42]. However, we note that, even for the sparse graphs considered throughout this section, stronger algorithms exist in the literature—see Ref. [43] for benchmarks.

Importantly, because all of the algorithms considered here operate by iteratively simplifying the problem, their performance is indicative of the fitness of the information used to perform these simplifications.

For each system size and average degree, we generated a set of 50 random Erdős-Rényi graphs and solved these with the `ReduMis` heuristic solver from the `KaMIS` project [44,45]. We use the solution obtained by `ReduMis` as (a proxy for) the optimal solution, and report the percentage of graphs for which each of the algorithms considered performs on par with the `ReduMis` reference.

The results for graph sizes with $n = 40$ –200 nodes and average degrees of 3, 5, and 12 are shown in Fig. 4. We used $\lambda = 1.1$ in the MIS Hamiltonian [see Eq. (1)] for RQAOA and QIRO. The performance of the two quantum algorithms (QIRO and RQAOA) is found to lie between the random and minimal degree greedy algorithms. This highlights the limitations of the quantum correlations obtained from QAOA at $p = 1$, which only “sees” the immediate neighbors of every node. As such, it is unsurprising that the minimal degree greedy algorithm—which makes the optimal greedy move based on information from nearest neighbors—outperforms QIRO and RQAOA informed by $p = 1$ QAOA. However, our results in Appendix E indicate that (at least on small graphs), QIRO with QAOA at $p > 1$ outperforms the greedy benchmark.

Furthermore, RQAOA (slightly, but consistently) outperforms QIRO. This is likely due to the fact that the QIRO simplification rules as defined in Sec. III B 1 are manifestly local. The fact that the gap between QIRO and RQAOA seems to diminish for denser graphs (i.e., for increasing average degree) corroborates this intuitive explanation, as we expect locality to be a weaker limitation if the underlying interaction graph has a higher connectivity.

A major advantage of QIRO (as compared to, e.g., RQAOA) is that the problem-specific classical subroutines can be used to enforce feasibility of the solutions generated (i.e., all obtained candidate solutions are guaranteed to be independent sets). To showcase this feature, we performed numerical simulations of QIRO and RQAOA with $p = 1$ QAOA with suboptimal parameters. The goal of these experiments is to consider the robustness of the algorithms against lower-quality quantum information. The parameters are generated by sampling a uniform 30×30 grid in

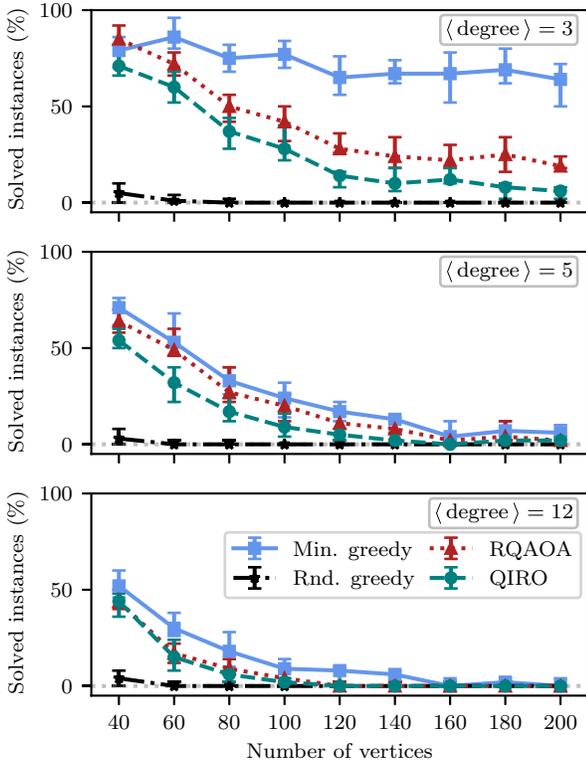


FIG. 4. System size scaling for the success rate of the random greedy (“Rnd. greedy”), minimal degree greedy (“Min. greedy”), RQAOA, and QIRO for Erdős-Rényi graphs with different expectation values of the average degree. At each system size and average degree 50 graphs were generated, and ten runs of each algorithm were performed. We report the median percentage of graphs for which we find independent sets of the size found by the classical `ReduMis` benchmark algorithm. The error bars denote the minimum and maximum performance observed across the ten runs. Note that the results for QIRO and RQAOA are obtained using simulations of QAOA at depth $p = 1$.

the (β, γ) parameter plane, and choosing the parameters corresponding to the q th quantile in the energy values, among the 30×30 energies obtained. Given their relevance for neutral atom quantum devices [13,14], we use UDGs (see Appendix D) for these experiments.

The results for an ensemble of 50 UDGs with 137 nodes and two different values of the penalty term λ are shown in Fig. 5(b). For every graph and parameter quality q , ten runs of each algorithm were performed. We show the mean approximation ratio of valid solutions produced by RQAOA and QIRO as plotted against the quality quantile q of the parameters used to generate the correlations by the QAOA. The parameters at $q = 0$ correspond to the fully optimized parameters. As expected, QIRO only produces valid solutions. Remarkably, if perfect (i.e., optimized) parameters are supplied to RQAOA, it likewise produces exclusively valid solutions. However, for both values of λ , there is a sharp drop-off in the number

of valid solutions for RQAOA. Because finding optimal parameters in an actual experimental setting is challenging [46,47], the robustness of QIRO is a valuable feature, in particular for realistic experimental scenarios. Additionally, results in Fig. 5(b) suggest that increasing the quality of the quantum correlations has a positive impact on the performance of both QIRO as well as RQAOA. This is in agreement with the outcome of our experiments at $p > 1$ in Appendix E, where a clear improved QIRO performance with increasing p is observed.

Surprisingly, results shown in Fig. 5(b) indicate that the number of valid solutions produced by RQAOA falls off at a comparatively better parameter quality (lower q) for the larger value of λ , defying the expectations that a larger penalty term λ should favor valid solutions. In general, we observe complex behaviour in the (λ, q) parameter plane (see Fig. 10 and the associated discussion in Appendix F). We note that determining the magnitude of penalty terms is a well-known issue when incorporating hard constraints into an unconstrained optimization setup [10,11].

Importantly, at each step of the QIRO algorithm, what remains to be solved is a valid MIS problem. Moreover, if the starting graph is a UDG [as in Fig. 5(a)] then so are all successive (simplified) graphs—see also Fig. 2. Because UDGs can be naturally embedded on neutral atom quantum processors [13,14], we can use such devices to produce the required quantum correlations. In what follows, we present results from the QuEra Aquila neutral atom quantum processor, as accessed through Amazon Braket. The details on the (quasi)adiabatic protocols used to generate the correlations on the QuEra Aquila device can be found in Appendix D. Because of the limited accessibility of the device, we report the results on a per-instance basis. We sort the instances in increasing order of their classical hardness [see Eq. (D1)], as proposed by Ebadi *et al.* [14]. We provide details about the hardness of the considered graph instances in Appendix D.

In Fig. 6 we display the approximation ratios $|\text{IS}|/|\text{MIS}|$ obtained by QIRO with correlations from the QuEra device (QIRO QuEra) and classical simulations of $p = 1$ QAOA (QIRO QAOA). On the QuEra device, we generate the correlations using the protocols outlined in Appendix D. We compare these to the sizes of the independent sets found by the minimal degree greedy algorithm. The reference maximum independent set is computed using the tensor network algorithm from Ref. [48]. For clarity, results of RQAOA are omitted from the plot, because the approximation ratios found for UDGs were generally comparable to those found by QIRO QAOA [as seen in Fig. 5(b)].

Our results indicate that the quantum correlations from the QuEra machine are better at guiding the QIRO procedure than numerical simulations of QAOA at $p = 1$. This is in line with previous results for solving the MIS problem using analog protocols on such hardware [14,49]. A larger sample size would be required to solidify such claims.

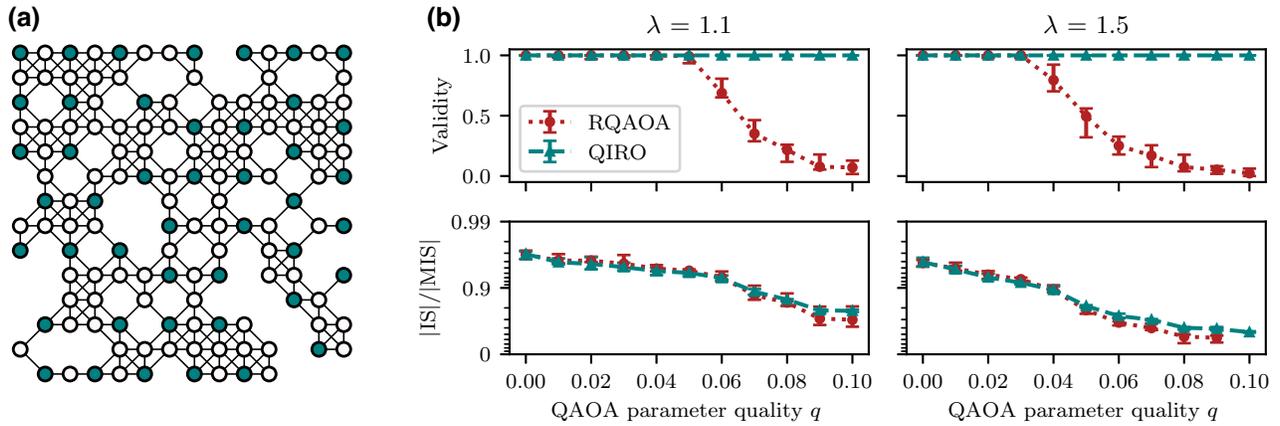


FIG. 5. (a) An example unit disk graph with 137 nodes. Teal nodes indicate a maximum independent set of size $|MIS| = 45$. (b) Data for 500 random unit disk graphs with 137 nodes, comparing the robustness of RQAOA and QIRO, for two different magnitudes of the penalty term λ [see Eq. (1)]. The upper row shows the fraction of valid solutions obtained by each algorithm. The lower row displays the ratio between the size of the largest independent set (IS) found and the maximum independent set (MIS) size. Both the validity and $|IS|/|MIS|$ are plotted with respect to the quality of the $p = 1$ QAOA parameters, with parameters at $q = 0.0$ corresponding to the optimal parameters, and higher values of q corresponding to worse parameters (see the main text). We note that we only report the approximation ratio achieved by RQAOA if the solution produced by it is valid (i.e., corresponds to an independent set).

However, initial results appear to suggest that QIRO QuEra performs comparably to the minimal degree greedy algorithm, reinforcing our belief that improved quantum correlations boost the performance of QIRO. Strikingly, the minimal degree greedy algorithm fails dramatically on some graphs (e.g., instance 1), where it achieves an approximation ratio of around 0.6, while QIRO QuEra is able to solve the same problems (nearly) optimally. This preliminary evidence showcases the potential utility of quantum correlations for guiding optimization algorithms.

B. MAX-2-SAT

Next, we analyze the performance of QIRO and QIRO with backtracking (QIRO + BT) on random MAX-2-SAT

instances with clauses of length two. We compare our results to those obtained with RQAOA and simulated annealing (SA) [50], an established and relatively simple classical heuristic. More detailed simulations with other, more powerful heuristics are left for future studies. Details about our implementation of SA can be found in Appendix G.

We first analyze the system size scaling of the algorithms' performance for random MAX-2-SAT instances with up to 160 variables, at three different values of the clause-to-variable ratio α . We use the MAX-SAT solver RC2 [51] to compute the optimal solutions, and use parallel tempering (PT) [34] to compute (approximately) optimal solutions for instances where the RC2 solver timed out. Details of our implementation of PT can be found

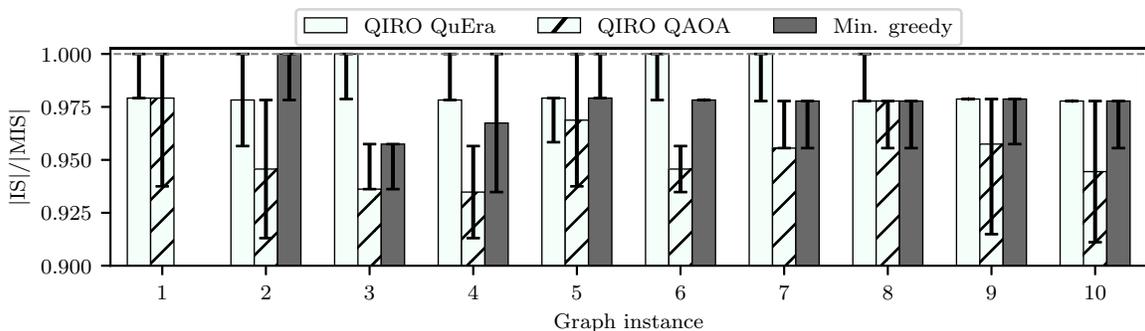


FIG. 6. Approximation ratio comparing the size of the independent set $|IS|$ found by QIRO with correlations from the QuEra neutral atom quantum processor (QIRO QuEra), and correlations from numerical simulations of $p = 1$ QAOA (QIRO QAOA), with the size of the maximum independent set $|MIS|$. For comparison, the ratios obtained by the minimal degree greedy algorithm (“Min. greedy”) are shown. In instance 1 the obtained approximation ratio of the greedy algorithm is outside of the range of the displayed values with $|IS|/|MIS| \approx 0.63$. The median value over five runs is plotted, with error bars showing the best and worst solutions found.

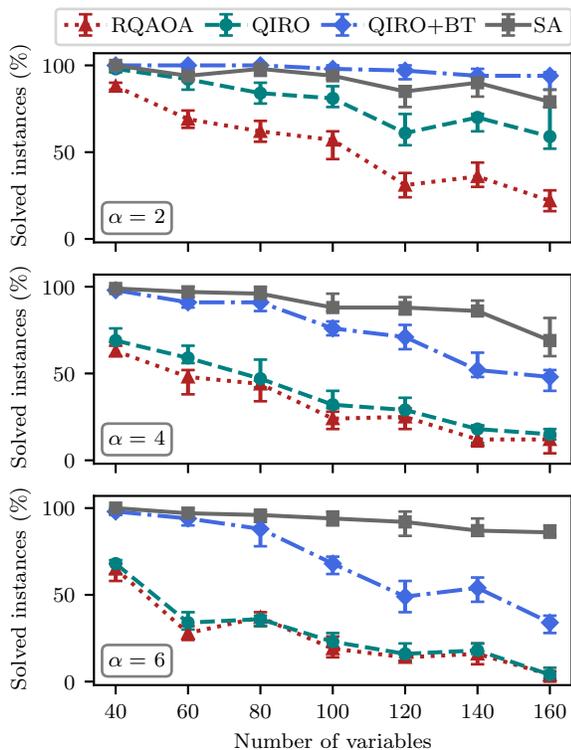


FIG. 7. System size scaling of the performance of the different algorithms, at three different values of α . At each system size and α , 50 random instances were generated with two variables per clause, each variable and its polarity chosen uniformly at random. Parallel tempering (PT) was used to find nearly optimal solutions; we then report the median percentage of instances on which an algorithm performs at least as well as PT across ten independent runs. The error bars indicate the minimal and maximal performance across these ten runs. Simulation details about SA and PT can be found in Appendix G.

in Appendix G. We then report the number of times the algorithms perform at least as well as PT or the RC2 solver.

Figure 7 showcases the results of our experiments. QIRO performs better than RQAOA for all clause-to-variable ratios α and over the entire range of the considered problem sizes. This shows that performance gains can be made by including very simple problem-specific update rules that require only modest classical resources. We note, however, that the improvement is especially prominent for lower values of α , and diminishes for problems with higher α (connectivity). Moreover, due to the inclusion of inference rules, the Boolean formula is simplified in fewer update steps, thus requiring fewer calls to the quantum device (and fewer optimizations of the QAOA circuit) than in RQAOA.

As expected, we note a decrease in performance with increasing problem sizes, with both QIRO and RQAOA often struggling to find optimal solutions at the largest problem sizes considered. However, with the addition of backtracking, QIRO + BT is able to optimally solve more

than half of the instances for almost all system sizes and all values of α considered here. For $\alpha = 2$, it even slightly outperforms our SA implementation. It should be noted that while the number of calls to the quantum device increases significantly (i.e., from linear to quadratic in the problem size) if backtracking is included, the requirements on the quantum hardware in terms of the gate and qubit count are kept constant. This fundamentally differs from QAOA, where an improved performance can only be obtained by increasing the circuit depth p , and thus the quantum resources required.

V. DISCUSSION

In this work we have extended previous efforts on hybrid quantum-classical algorithms for combinatorial optimization, by introducing a family of QIRO algorithms. We propose the use of quantum-informed classical update steps to recursively simplify the original optimization problem, allowing for the incorporation of problem-specific classical subroutines (e.g., satisfiability inference rules) from the rich literature on classical combinatorial optimization [25]. We have shown via numerical simulations for random instances of two NP-hard combinatorial optimization problems that such classical subroutines can significantly strengthen the algorithm’s performance. Additionally, we compare the performance of QIRO to RQAOA, providing, to the best of our knowledge, the first systematic numerical study of RQAOA on problems that are not \mathbb{Z}_2 symmetric.

Importantly, we think that QIRO offers a promising approach for solving optimization problems with *hard* constraints, which have previously proven challenging for quantum algorithms [10,11,52]. Using the MIS problem as a guiding example, we have shown how classical update steps can be used to enforce hard constraints and only produce feasible solutions. This is true regardless of the quality of the quantum resources, endowing QIRO with the robustness essential for algorithms within the NISQ era.

Interestingly, we observe that keeping the quantum resources equal, RQAOA outperforms QIRO on the MIS problem. We believe that this is a consequence of the local nature of the QIRO update rules for MIS. This is corroborated by our finding that improved quantum correlations from the neutral atom device enhance QIRO’s performance compared to when QAOA at $p = 1$ is used to generate the correlations. The intuitive explanation is that the quantum many-body evolution on a neutral atom quantum processor provides sufficient nonlocal effects to counteract the locality of the update steps. In fact, adiabatic protocols analogous to those used to generate correlations here have previously been shown to produce candidate solutions surpassing those produced by QAOA at depths higher than $p = 1$ [14]. Taken together, our results may serve as yet another indication that locality is a limiting factor for (quantum) optimization algorithms [53].

Furthermore, we have confirmed the intuitive expectation that higher-quality quantum correlations should enhance the performance of the algorithm, as evidenced by means of extensive numerical simulations, as well as experiments on an actual neutral atom quantum device. This result suggests that advances in quantum hardware should lead to further performance gains.

We note that we performed no optimization of the protocols used to generate the correlations in our experiments on the QuEra Aquila neutral atom device. Previously, dramatic improvements in the quality of candidate solutions have been obtained by first optimizing the schedules [14,49]. Moreover, the (unoptimized) protocols and device used in this study (and specified in Appendix D) have been shown to produce candidate solutions typically corresponding to either very small independent sets, or possessing multiple violations of the independence constraint [49]. Thus, it is reasonable to expect that such schedule optimization schemes could greatly benefit the quantum state preparation step in QIRO. Moreover, the ability of QIRO to find good (feasible) solutions despite the limitations of the quantum hardware serves as yet another demonstration of its robustness.

In addition, we have introduced the use of backtracking to obtain improved results via additional calls to existing quantum hardware, rather than requiring deeper circuits, as is usually the case in quantum (optimization) algorithms. This scheme has led to promising results for the MAX-2-SAT problem, where QIRO enhanced with backtracking could optimally solve most random problem instances, even at comparatively large instance sizes. Because large instances are accessible in numerical simulations (at $p = 1$ of QAOA), we were able to sensibly compare the performance of QIRO(+BT) with widely used classical heuristic solvers such as simulated annealing.

Finally, we would like to stress that our work should be seen as a first demonstration of a general template for a larger class of hybrid quantum-classical algorithms for combinatorial optimization. We firmly believe that hybrid techniques will be necessary to unlock the advantages of quantum devices. The modular nature of QIRO offers a plethora of ways in which the presented toolbox can be modified; we touch upon some potential extensions in the next section.

VI. OUTLOOK

Several modifications to the QIRO algorithm as presented here can be envisioned. Exploring update rules for optimization problems beyond those examined in this manuscript would be a compelling avenue for future research. Likewise, it would be intriguing to devise the update rules for weighted versions of the problems addressed here. Moreover, one could attempt to extend the information extracted from the quantum state beyond

just the maximal correlation. This could be accomplished by a randomized strategy, with the selection probabilities determined as some function of the correlation matrix. Furthermore, devising update rules incorporating correlations beyond those between nearest neighbors in the interaction graph might be useful, especially in combination with deeper QAOA circuits or on analog devices with long-range interactions (e.g., Rydberg atom arrays). Additionally, one could exploit the locality of QAOA, and simultaneously round multiple correlations at each update step, if the corresponding variables are outside of each other's sphere of influence. Thus, one could save resources by requiring fewer calls to the quantum routines, and thus mitigate one of the limitations of the presented scheme. On the other hand, if deeper quantum circuits (e.g., QAOA at $p > 1$) are available, extracting correlations beyond those between variables adjacent in the interaction graph might be beneficial. Finally, considering problem-tailored alternatives to a brute-force search of the remaining solution space at a threshold problem size n_c may significantly improve the performance and resource efficiency of QIRO.

It is unreasonable to expect that the simple backtracking strategy presented here would be equally successful for larger problem instances, or instances from different problem classes. Hence, one might need to devise improved backtracking strategies, both in terms of the resources required, as well as in terms of the performance. One possibility is the use of reinforcement learning to guide the reduction process, as proposed by Patel *et al.* [33]. Resources might be saved by efficient pruning of the search tree using existing methods [30], identifying and ignoring unfavorable parts of the solution space. Therefore, fewer quantum state preparations would be required. Finally, one could use backtracking to enforce constraints, by identifying steps where a constraint was violated, backtracking to it, and altering the decision made at that step.

It would likewise be interesting to consider alternative methods of generating the quantum correlations. This could include efficient numerical methods to obtain correlations with the QAOA at $p > 1$, by means of, e.g., tensor network simulations [54,55]. Naturally, it would be interesting to acquire correlations from higher-depth QAOA circuits executed on real digital quantum computers, or by generating them using quantum annealers based on superconducting flux qubits [56,57]. One could consider optimizing analog protocols (using, e.g., Bayesian optimization [49]) to obtain better correlations, and enhance the effectiveness of the algorithm. Moreover, modifying the cost function with respect to which the quantum resources are optimized might be helpful, similarly to the approach employed by Caha *et al.* [58]. Finally, one could attempt to generate correlations using classical approximation algorithms, as proposed by Wagner *et al.* [59]. Long story short, there is a lot to be done, but this paper is already long enough.

Conveniently, the modular nature of QIRO allows us to analyze the promise of quantum computers for combinatorial optimization, e.g., by comparing QIRO with different quantum and classical methods of computing the correlations. Should shortcomings of quantum approaches (e.g., locality) be identified in this manner, the classical sub-routines in QIRO can be designed to rectify them. As such, we hope that our work can bring us closer to determining whether quantum devices can provide value for combinatorial optimization problems.

ACKNOWLEDGMENTS

J.R.F. and A.K. thank Libor Caha and Alexander Klichev for insightful discussions. The authors thank Lilly Palackal, Maximilian Passek, Carlos Riofrío, and Gili Rosenberg for thorough reviews of the manuscript, and the Amazon Braket, BMW, and QuEra teams for their support. C.M. thanks the Munich Quantum Valley initiative, which is supported by the Bavarian State Government with funds from the Hightech Agenda Bayern Plus. H.G.K. would like to thank Am Platzl 1A for providing the necessary environment for creative thinking. An open-source implementation of QIRO is available online [60].

APPENDIX A: LOW-DEPTH QAOA SIMULATIONS

Here, we provide details of our implementation of the numerical simulations of QAOA at $p = 1$. If we limit ourselves to problems with general Ising-like (i.e., quadratic) cost Hamiltonians of the form

$$\hat{H}_c = \sum_i h_i \hat{Z}_i + \sum_{i < j} J_{ij} \hat{Z}_i \hat{Z}_j, \quad (\text{A1})$$

it was shown that the expectation values of \hat{Z}_i and $\hat{Z}_i \hat{Z}_j$ according to the depth- $(p = 1)$ QAOA state $|\psi(\beta_1, \gamma_1)\rangle = e^{-i\beta_1 \hat{H}_{\text{mix}} e^{-i\gamma_1 \hat{H}_c} |+\rangle^{\otimes n}$ can be calculated in $\mathcal{O}(n)$ time via simple analytical expressions [3,28]. Because the cost Hamiltonian comprises $\mathcal{O}(n^2)$ terms, the complexity of simulating the calculation of the correlation matrix M in the update step of Algorithm 1 is $\mathcal{O}(n^3)$. Therefore, assuming $\mathcal{O}(n)$ update steps, this yields an overall complexity of $\mathcal{O}(n^4)$ for the full QIRO simulation and $\mathcal{O}(n^5)$ for the QIRO + BT simulation, when $\mathcal{O}(n)$ backtracking steps are performed, as is the case following the prescription of Algorithm 3.

We comment here on the difference in the measurement processes between QIRO and vanilla QAOA. If the cost Hamiltonian is k local, the process of finding the optimal parameters according to the expectation value of the cost Hamiltonian in principle requires only measurements with k qubits involved for both algorithms. However, after the optimal parameters are found, QAOA samples classical bit strings from the quantum state for which measurements of

all qubits involved are required. This makes QAOA exponentially difficult to simulate classically for any depth p [61]. In contrast, the sampling step is replaced by a classical update step in QIRO. Hence, using the depth- $(p = 1)$ QAOA as state preparation for QIRO allows us to simulate the algorithm efficiently.

One may wonder why this should be a quantum algorithm of interest, if it is amenable to efficient classical simulations. The answer to this lies in the fact that the situation changes dramatically at higher depths of QAOA with $p > 1$ where no method for analytically computing the one-point and two-point correlations classically in polynomial time is known (to us). The same holds at $p = 1$ when calculating the expectation values of k -body interaction terms, with $k \geq 3$. Therefore, these cases could lead to an algorithm that provides quantum advantage. In general, we expect an increasing performance of QIRO for higher depths of QAOA, as indicated by the MIS experiments in Sec. IV A. We note that similar observations were previously made by Bravyi *et al.* [3].

APPENDIX B: RQAOA

To overcome the limitations of the QAOA due to its local structure [3,53], Bravyi *et al.* [3,7] introduced RQAOA. The algorithm optimizes a QAOA state based on the current problem Hamiltonian and uses this information to reduce the problem via the elimination of a single variable. Recursive application of this step reduces the problem until a specified threshold of remaining variables n_c is reached. The remaining problem can be solved by a classical solver (e.g., by a brute-force search). The elimination step is motivated by the rounding of fractional correlations in relaxations of linear programs [3]. While hardly a fair comparison in terms of resources required, RQAOA has been shown to significantly outperform QAOA [7,62].

It was first applied to the \mathbb{Z}_2 -symmetric Max-Cut problem, which *exclusively* comprises two-body terms. Here, we (trivially) extend the approach to also include single-body terms. We additionally note that a generalization to Hamiltonians with an arbitrary degree of interaction can be envisaged. However, for Hamiltonians with K -body terms with $K \geq 3$, one faces an additional obstacle, as the degree of the interactions can increase when an elimination step is performed. This comes with additional requirements on the quantum hardware; as we are mostly interested in NISQ-era algorithms, we do not consider such extensions here.

Assuming a quadratic Hamiltonian of form (A1), the pseudocode for the RQAOA, as proposed in Refs. [3,7], is given in Algorithm 4.

APPENDIX C: MAX-2-SAT SIMPLIFICATION DETAILS

Here we provide the details on the update rules employed in the QIRO update step for MAX-2-SAT.

ALGORITHM 4. RQAOA.

Input: Problem Hamiltonian \hat{H} , threshold n_c .
Output: Solution S .

- 1: $L \leftarrow \{\}$ \triangleright Initialize list to store frozen correlations.
- 2: **for** $k = 1$ **to** $n - n_c$ **do**
- 3: Find optimal variational parameters of QAOA:
 $\beta^*, \gamma^* = \arg \min_{\beta, \gamma} \langle \psi(\beta, \gamma) | \hat{H} | \psi(\beta, \gamma) \rangle$
- 4: $|\psi\rangle \leftarrow |\psi(\beta^*, \gamma^*)\rangle$
- 5: Store correlations in $M \in \mathbb{R}^{n \times n}$, initialize $M_{ij} = 0$:
 $\forall i \in [n] : M_{ii} = \langle \psi | \hat{Z}_i | \psi \rangle,$
 $\forall (i, j) \in [n] \times [n] \text{ s.t. } J_{ij} \neq 0 : M_{ij} = \langle \psi | \hat{Z}_i \hat{Z}_j | \psi \rangle.$
- 6: $(i, j) = \arg \max_{(i, j) \in [n] \times [n]} |M_{ij}|$
- 7: Append $\{(i, j), \text{sign}(M_{ij})\}$ to L .
- 8: **if** $i = j$ **then** \triangleright One-point correlation.
- 9: $\hat{H} \leftarrow \text{Repl. } \hat{Z}_i \text{ in } \hat{H} \text{ with } \text{sign}(M_{ii}) \cdot 1$
- 10: **else** \triangleright Two-point correlation.
- 11: $\hat{H} \leftarrow \text{Repl. } \hat{Z}_i \text{ in } \hat{H} \text{ with } \text{sign}(M_{ij}) \cdot \hat{Z}_j$
- 12: **end if**
- 13: **end for**
- 14: $S_{n_c} \leftarrow \text{Bruteforce}(\hat{H})$ \triangleright Solve the remaining problem.
- 15: $S \leftarrow \text{Reconstruct}(S_{n_c}, L)$ \triangleright Extract the solution from L .

Inference rules are a commonly used building block of many classical satisfiability solvers [31]. They are used to simplify the formula by fixing the assignments of specific variables and therefore accelerate the remaining computation. Here, we use only very simple inference rules that have low computational (and conceptual) complexity. The selected rules ensure that the optimal solutions of both the original and the simplified problems are equally good in quality. Despite their simplicity, the inference rules improve the performance of QIRO, both in terms of quality (see Sec. IV B) and in terms of resource efficiency. This is because fewer calls to the quantum device are required, due to the simplifications performed by the inference rules. The selected rules, listed next, were proposed in Refs. [30,31].

- (a) The *pure literal rule* [32]: if a literal exclusively occurs with positive (negative) polarity then the value of the corresponding variable is set to TRUE (FALSE).
- (b) The *almost common clause rule* [32]: if a MAX-2-SAT instance includes the clauses $x_i \vee x_j$ and $\bar{x}_i \vee x_j$ then both clauses can be replaced with a single-literal clause (i.e., unit clause) x_j . If neither x_i nor \bar{x}_i appears in the formula anymore, the variable can be arbitrarily set to either TRUE or FALSE.
- (c) The *complementary unit clause rule* [63]: if the formula contains single-literal clauses x_i and \bar{x}_i then these two clauses are removed. Again, if neither x_i nor \bar{x}_i appears in the formula thereafter, the assignment of the variable can be arbitrarily set to either TRUE or FALSE.

ALGORITHM 5. Reduce (MAX-2-SAT).

Input: Formula ϕ , partial solution S .
Output: Simpl. formula ϕ , updated partial solution S .

- 1: Prepare low-energy quantum state $|\psi\rangle$.
- 2: Store correlations in $M \in \mathbb{R}^{n \times n}$:
 $\forall i \in [n] : M_{ii} = \langle \psi | \hat{Z}_i | \psi \rangle,$
 $\forall (i, j) \in [n] \times [n] \text{ s.t. } J_{ij} \neq 0 : M_{ij} = \langle \psi | \hat{Z}_i \hat{Z}_j | \psi \rangle.$
- 3: **if** $i = j$ **then** \triangleright One-point correlation.
- 4: Assign variable $x_i \rightarrow \frac{1}{2}(\text{sign}(M_{ii}) + 1)$.
- 5: **else** \triangleright Two-point correlation.
- 6: **if** $\text{sign}(M_{ij}) = 1$ **then**
- 7: Replace $x_i \rightarrow x_j$ in ϕ and extend S accordingly.
- 8: **else**
- 9: Replace $x_i \rightarrow \bar{x}_j$ in ϕ and extend S accordingly.
- 10: **end if**
- 11: **end if**
- 12: $\phi' \leftarrow \{\}$ \triangleright Initialize empty formula for comparison.
- 13: **while** $\phi \neq \phi'$ and $\text{size}(\phi) > n_c$ **do** \triangleright Inference rules.
- 14: $\phi' \leftarrow \phi$
- 15: $\phi, S \leftarrow \text{PureLiteral}(\phi, S)$
- 16: $\phi, S \leftarrow \text{DominatingUnitClause}(\phi, S)$
- 17: $\phi, S \leftarrow \text{AlmostCommonClause}(\phi, S)$
- 18: $\phi, S \leftarrow \text{ComplementaryUnitClause}(\phi, S)$
- 19: **end while**
- 20: **if** $\text{size}(\phi) \leq n_c$ **then**
- 21: $\phi, S \leftarrow \text{BruteForce}(\phi, S)$
- 22: **end if**
- 23: **return** ϕ, S

- (d) The *dominating unit clause rule* [63]: if the total count of clauses, regardless of their length, that contain a variable x_i (\bar{x}_i) is not greater than the count of unit clauses that contain \bar{x}_i (x_i), then the variable x_i is set to FALSE (TRUE).

The complete QIRO update step Reduce for MAX-2-SAT is shown in Algorithm 5.

Importantly, the MAX-2-SAT QIRO update rules presented here preserve the problem structure in that only clauses with two or less variables appear in the formula. Consequently, the Hamiltonian remains in the same form as in Eq. (A1).

APPENDIX D: MIS ON A NEUTRAL ATOM QUANTUM DEVICE

Here, we provide details about the adiabatic protocols on a neutral atom quantum processor that we used to prepare the quantum states used to guide the QIRO algorithm for the MIS problem. The MIS problem on UDGs naturally emerges in the context of neutral atom quantum processors based on Rydberg atom arrays [13,14]. UDGs are a family of graphs, where two nodes are connected only if the (Euclidean) distance between them is smaller than some threshold radius R_d .

Here, we generate such graphs by positioning nodes on an underlying square lattice with lattice constant a .

TABLE I. Values of the hardness parameter \mathcal{HP} for the graph instances presented in Fig. 6.

Graph instance	1	2	3	4	5	6	7	8	9	10
\mathcal{HP}	3	8	14	39	49	95	197	383	863	1435

We set $\sqrt{2}a < R_d < 2a$, such that we get nearest-neighbor and diagonal connectivity. We chose a lattice constant $a = 5.3 \mu\text{m}$, inspired by a previous study [49]. The graphs presented here were generated by positioning 137 nodes on such an underlying lattice with 14×14 sites. This corresponds to a filling factor of roughly 70%. An example graph is shown in Fig. 5(a).

Furthermore, we classified the generated instances in terms of their classical hardness parameter

$$\mathcal{HP} := \frac{N_{|\text{MIS}|-1}}{|\text{MIS}| N_{|\text{MIS}|}}, \quad (\text{D1})$$

where N_M denotes the number of independent sets of size M . The hardness parameter was shown to be related to the performance of quantum and classical algorithms alike [14]. However, we note that evidence has recently been put forward that \mathcal{HP} only influences the performance of Markov chain Monte Carlo-based algorithms (e.g., SA) [64].

Here, we consider graphs spanning several orders of magnitude in \mathcal{HP} , which we computed using the tropical tensor network algorithm provided within the `GenericTensorNetworks` library [48,65]. Specifically, the graphs considered in Fig. 6 have been sorted in order of increasing hardness, with the exact values of \mathcal{HP} given in Table I.

We now briefly summarize how neutral atom arrays can be used to solve the MIS problem on UDGs [13,14]. The dynamics of neutral atom quantum devices are governed by the Hamiltonian $\hat{H}(t) = \hat{H}_{\text{dr}} + \hat{H}_{\text{cost}}$, where the individual terms read

$$\begin{aligned} \hat{H}_{\text{dr}} &= \frac{\hbar}{2} \Omega(t) \sum_i \hat{X}_i, \\ \hat{H}_{\text{cost}} &= -\hbar \Delta(t) \sum_i \hat{n}_i + \sum_{i<j} V_{ij} \hat{n}_i \hat{n}_j. \end{aligned} \quad (\text{D2})$$

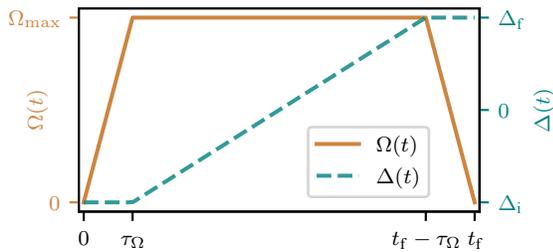


FIG. 8. Time dependency of the Rabi frequency $\Omega(t)$ and the laser detuning $\Delta(t)$. We use $t_f = 4 \mu\text{s}$, $\Omega_{\text{max}} = 15.8 \text{ MHz}$, $\tau_{\Omega} = 0.1 t_f$, $\Delta_i = -30 \text{ MHz}$, and $\Delta_f = 60 \text{ MHz}$.

Here, Ω is the Rabi frequency, Δ is the laser detuning, and the interaction term $V_{ij} = C_6 \|\mathbf{x}_i - \mathbf{x}_j\|_2^{-6}$, with C_6 the van der Waals coefficient. The particular value C_6 depends on the atom species used in the experimental setup. The number operator $\hat{n}_i := \frac{1}{2}(\mathbb{1}_2 - \hat{Z}_i)$ counts the number of Rydberg excitations on the i th site.

An essential property of the Hamiltonian in Eq. (D2) is the so-called Rydberg blockade phenomenon, in which two atoms cannot simultaneously be in the (excited) state $|1\rangle$ if the distance between them is smaller than the Rydberg blockade radius $R_b \equiv (C_6/\hbar\Omega)^{1/6}$ [14]. Therefore, if one sets the UDG radius to be equal to the Rydberg blockade radius ($R_d = R_b$), ground states of \hat{H}_{cost} obey the independence constraint on UDGs if one assigns atoms that are in state $|1\rangle$ to the MIS [13]. Additionally, if we set $0 < \Delta < V_{ij}$, the ground state of H_{cost} maximizes the number of excitations, while not violating independence, hence corresponding to solutions to the original MIS problem. We thus start the adiabatic protocol with $\Delta_i < 0$, such that the ground state is $|0\rangle^{\otimes |V|}$ and then (slowly) increase the detuning until some final positive value, as shown in Fig. 8, such that the ground states of the final Hamiltonian correspond to the solutions of the MIS problem.

We ran such protocols on the QuEra Aquila device, which is available through the Amazon Braket service. The detailed parameters can be found in the caption of Fig. 8, and are partially motivated by a previous study [49]. At each stage of the QIRO algorithm we estimate the correlations using a hundred measurements of the quantum state prepared by the adiabatic protocol described in Fig. 8.

APPENDIX E: HIGHER-DEPTH QAOA EXPERIMENTS

To provide additional support to the claim that a higher quality of correlations yields better performance of QIRO, we performed numerical simulations of QIRO using correlations from QAOA at $p > 1$. In contrast to $p = 1$, no analytical formulae for the required correlators are known at $p > 1$, forcing us to resort to state-vector simulations of QAOA circuits. Consequently, we had to limit our experiments to solving the MIS problem on smaller graph instances with 12 nodes, randomly chosen from the Erdős-Rényi ensemble as in Sec. IV A. For each fixed expected degree, we generated a hundred random instances, and found their MIS using a brute-force search of the solution space. We then ran QIRO informed by correlations from the QAOA implementation in the PennyLane library

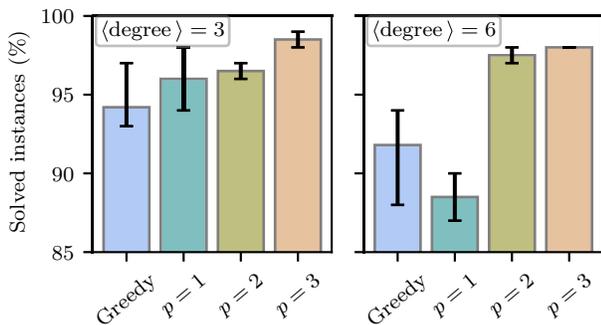


FIG. 9. The fraction of optimally solved instances by QIRO using QAOA at depths $p \in \{1, 2, 3\}$, and by the minimal degree greedy algorithm. The height of the bars show the mean fraction across ten runs of the algorithms, with the error bars indicating the best and worst performance across those runs.

[66], using the gradient descent optimizer with 300 iterations and 15 restarts of the optimizer at different random initial values of the variational parameters. For these experiments, we set $n_c = 1$ because of the comparatively smaller problem instance sizes.

The results in Fig. 9 show the fraction of optimal independent sets found by QIRO with correlations from QAOA at depths $p \in \{1, 2, 3\}$, and by the minimal degree greedy algorithm. The performance of QIRO clearly improves with increased QAOA depth, for both considered graph densities. Together with the results in Sec. IV A, this provides yet another indication that enhanced correlations lead to better performance of QIRO. Moreover, these experiments indicate that correlations from deeper circuits enable QIRO to surpass the performance of the greedy benchmark. Although this matches the intuitive expectations, further experiments on larger instances and quantum hardware are required to strengthen these claims.

APPENDIX F: FURTHER RQAOA VALIDITY EXPERIMENTS

In Sec. IV A we assessed the performance of QIRO and RQAOA at different qualities of QAOA parameters. As the results presented in Fig. 5(b) exhibited interesting behavior, we here report an extended set of experiments, where RQAOA validities for additional magnitudes of the penalty term λ [cf. Eq. (1)] are shown.

Interestingly, the results partially defy the naïve expectation that a larger penalty term should favor enforcing constraints. This is especially prominent in the region between $\lambda = 1.2$ and $\lambda = 5$, where the robustness of RQAOA seems to decrease. However, the trend reverses again for $\lambda = 10$.

While a thorough analysis of the artefacts observed in Fig. 10 is beyond the scope of this paper, we offer some insights into what might underpin the peculiar behavior at intermediate values of λ . We suspect that this is caused

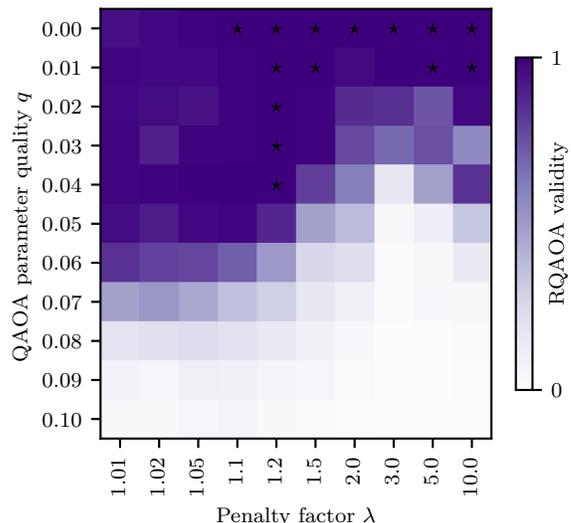


FIG. 10. Ratio of valid solutions obtained by RQAOA for different qualities of the QAOA parameters q (at $p = 1$) and different penalty factors λ . As in the results in Fig. 5(b), at each (q, λ) the same 50 UDGs of size 137 were generated and ten runs of RQAOA were performed. We here report the ratio valid (i.e., independent) sets found by RQAOA. Black stars indicate values of (q, λ) where all found solutions were feasible.

by an intricate interplay between the deformations of the spectrum of the cost Hamiltonian (1) when λ is varied, with the particularities of QAOA states at $p = 1$. As λ is varied, only the energy levels of states that correspond to configurations with violations of the independence constraint are changed. Thus, the nature of the excited states can change dramatically, i.e., states corresponding to configurations with violations might shift relative to states corresponding to independent sets. It is difficult to predict the amplitudes of such excited states in the optimal QAOA states at $p = 1$, let alone at perturbed parameters. As such, a pattern similar to that seen is not inconceivable. However, more experiments are required to pin down the exact underlying mechanism.

APPENDIX G: PARALLEL TEMPERING AND THE SIMULATED ANNEALING SETUP

For a MAX-2-SAT instance with n variables, starting from a random initial configuration, we perform simulated annealing with $600n$ attempted variable flips. The annealing schedule consists of evenly spaced inverse temperatures $\beta = T^{-1}$, with an initial value $\beta = 0$ and a final inverse temperature of $\beta = 6$.

For PT [34–36], we initialize 12 replicas of the problem in random configurations with different temperatures. We chose the fine-tuned temperatures given in Table II. These were chosen in accordance with the prescription that replica exchanges should be accepted with probabilities between 0.2 and 0.8 [67]. We tested this empirically

TABLE II. Temperatures of the replicas in our implementation of PT.

Replica index	1	2	3	4	5	6	7	8	9	10	11	12
Temperature	0.10	0.20	0.29	0.39	0.50	0.62	0.75	0.90	1.09	1.33	1.67	2.20

for a collection of representative MAX-2-SAT instances. We then used the following PT scheme. In each run of PT we carried out 15 000 cycles, with one simulation cycle consisting of one Monte Carlo sweep (n single spin flips) per replica followed by an attempted replica exchange between all replica pairs with neighboring temperatures, starting from the lowest temperature. The configuration with the lowest energy found in any of the replicas during the course of the algorithm was finally returned as the solution.

- [1] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [ArXiv:1411.4028](#) (2014).
- [2] D. Gamarnik, The overlap gap property: A topological barrier to optimizing over random structures, [Proc. Natl. Acad. Sci.](#) **118**, e2108492118 (2021).
- [3] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to variational quantum optimization from symmetry protection, [Phys. Rev. Lett.](#) **125**, 260505 (2020).
- [4] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case, [ArXiv:2004.09002](#) (2020).
- [5] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples, [ArXiv:2005.08747](#) (2020).
- [6] C.-N. Chou, P. J. Love, J. S. Sandhu, and J. Shi, in *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 229, edited by M. Bojańczyk, E. Merelli, and D. P. Woodruff (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022), p. 41:1.
- [7] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Hybrid quantum-classical algorithms for approximate graph coloring, [Quantum](#) **6**, 678 (2022).
- [8] J. R. McClean, M. P. Harrigan, M. Mohseni, N. C. Rubin, Z. Jiang, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Low-depth mechanisms for quantum optimization, [PRX Quantum](#) **2**, 030312 (2021).
- [9] H. Karimi, G. Rosenber, and H. G. Katzgraber, Effective optimization using sample persistence: A case study on quantum annealers and various Monte Carlo optimization methods, [Phys. Rev. E](#) **96**, 043312 (2017).
- [10] A. Lucas, Ising formulations of many NP problems, [Front. Phys.](#) **2**, 1 (2014).
- [11] F. Glover, G. Kochenberger, and Y. Du, A tutorial on formulating and using QUBO models, [ArXiv:1811.11538](#) (2019).
- [12] S. Butenko and P. M. Pardalos, Ph.D. thesis, USA, 2003, aAI3120100.
- [13] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, Quantum optimization for maximum independent set using Rydberg atom arrays, [ArXiv:1808.10816](#) (2018).
- [14] S. Ebadi, *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, [Science](#) **376**, 1209 (2022).
- [15] J. Wurtz, P. Lopes, N. Gemelke, A. Keesling, and S. Wang, Industry applications of neutral-atom quantum computing solving independent set problems, [ArXiv:2205.08500](#) (2022).
- [16] S. A. Cook, in *Proceedings of the Third Annual ACM Symposium on Theory of Computing—STOC '71* (ACM Press, Shaker Heights, Ohio, United States, 1971), p. 151.
- [17] P. Cheeseman, B. Kanefsky, and W. M. Taylor, in *Proceedings of the 12th International Joint Conference on Artificial Intelligence—Volume 1, IJCAI'91* (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1991), p. 331.
- [18] L. Zeng, J. Zhang, and M. Sarovar, Schedule path optimization for adiabatic quantum computing and optimization, [J. Phys. A: Math. Theor.](#) **49**, 165305 (2016).
- [19] M. Mezard, G. Parisi, and R. Zecchina, Analytic and algorithmic solution of random satisfiability problems, [Science](#) **297**, 812 (2002).
- [20] J. Slegers, R. Olij, G. van Horn, and D. van den Berg, Where the really hard problems aren't, [Oper. Res. Perspect.](#) **7**, 100160 (2020).
- [21] G. Ochoa, F. Chicano, and M. Tomassini, in *Parallel Problem Solving from Nature—PPSN XVI* (Springer International Publishing, Leiden, The Netherlands, 2020), p. 125.
- [22] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (W. H. Freeman & Co., USA, 1990).
- [23] B. Aspvall, M. F. Plass, and R. E. Tarjan, A linear-time algorithm for testing the truth of certain quantified Boolean formulas, [Inf. Process. Lett.](#) **8**, 121 (1979).
- [24] V. Chvatal and B. Reed, in *Proceedings 33rd Annual Symposium on Foundations of Computer Science (SFCS '92)* (IEEE Computer Society, Pittsburgh, PA, USA, 1992), p. 620.
- [25] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization*, Dover Books on Computer Science (Dover Publications, Mineola, NY, 1998).
- [26] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [ArXiv:quant-ph/0001106](#) (2000).
- [27] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, [Phys. Rev. E](#) **58**, 5355 (1998).
- [28] 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).

- [29] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, *Rep. Progr. Phys.* **83**, 054401 (2020).
- [30] A. Abramé and D. Habet, ahmaxsat: Description and evaluation of a branch and bound Max-SAT solver, *J. Satisf. Boolean Model. Comput.* **9**, 89 (2015).
- [31] C. M. Li and F. Many, MaxSAT, hard and soft constraints, *Front. Artif. Intell. Appl.* **185**, 613 (2021).
- [32] B. Borchers and J. Furman, A two-phase exact algorithm for MAX-SAT and weighted MAX-SAT problems, *J. Comb. Optim.* **2**, 299 (1998).
- [33] Y. J. Patel, S. Jerbi, T. Bäck, and V. Dunjko, Reinforcement learning assisted recursive QAOA, *ArXiv:2207.06294* (2022).
- [34] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [35] D. J. Earl and M. W. Deem, Parallel tempering: Theory, applications, and new perspectives, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- [36] F. Romá, S. Risau-Gusman, A. Ramirez-Pastor, F. Nieto, and E. Vogel, The ground state energy of the Edwards–Anderson spin glass model with a parallel tempering Monte Carlo algorithm, *Physica A* **388**, 2821 (2009).
- [37] Z. Zhu, C. Fang, and H. G. Katzgraber, Borealis—A generalized global update algorithm for Boolean optimization problems, *Optim. Lett.* **14**, 2495 (2020).
- [38] S. Boettcher and A. G. Percus, in *Computational Modeling and Problem Solving in the Networked World* (Springer US, Chandler, Arizona, 2003), p. 61.
- [39] M. Mahajan and F. Slivovsky, in *26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023)* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Alghero, Italy, 2023).
- [40] P. Erdős and A. Rényi, On random graphs I, *Publ. Math. Debrecen* **6**, 290 (1959).
- [41] M. Dupont, B. Evert, M. J. Hodson, B. Sundar, S. Jeffrey, Y. Yamaguchi, D. Feng, F. B. Maciejewski, S. Hadfield, M. S. Alam, Z. Wang, S. Grabbe, P. A. Lott, E. G. Rieffel, D. Venturelli, and M. J. Reagor, Quantum-enhanced greedy combinatorial optimization solver, *Sci. Adv.* **9**, eadi0487 (2023).
- [42] M. M. Halldórsson and J. Radhakrishnan, Greed is good: Approximating independent sets in sparse and bounded-degree graphs, *Algorithmica* **18**, 145 (1997).
- [43] M. Böther, O. Kißig, M. Taraz, S. Cohen, K. Seidel, and T. Friedrich, in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25–29, 2022* (OpenReview.net, 2022), <https://openreview.net/forum?id=mk0HzdqY7i1>.
- [44] S. Lamm, P. Sanders, C. Schulz, D. Strash, and R. F. Werneck, Finding near-optimal independent sets at scale, *J. Heurist.* **23**, 207 (2017).
- [45] D. Hespe, C. Schulz, and D. Strash, Scalable kernelization for maximum independent sets, *ACM J. Exp. Algorithmics* **24**, 1 (2019).
- [46] J. Lee, A. B. Magann, H. A. Rabitz, and C. Arenz, Progress toward favorable landscapes in quantum combinatorial optimization, *Phys. Rev. A* **104**, 032401 (2021).
- [47] L. Bittel and M. Kliesch, Training variational quantum algorithms is NP-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [48] J.-G. Liu, X. Gao, M. Cain, M. D. Lukin, and S.-T. Wang, Computing solution space properties of combinatorial optimization problems via generic tensor networks, *SIAM J. Sci. Comput.* **45**, A1239 (2023).
- [49] J. R. Finžgar, M. J. A. Schuetz, J. K. Brubaker, H. Nishimori, and H. G. Katzgraber, Designing quantum annealing schedules using Bayesian optimization, *ArXiv:2305.13365* (2023).
- [50] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [51] A. Ignatiev, A. Morgado, and J. Marques-Silva, RC2: An efficient MaxSAT solver, *J. Satisf. Boolean Model. Comput.* **11**, 53 (2019).
- [52] S. Hadfield, Z. Wang, B. O’Gorman, E. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 34 (2019).
- [53] M. B. Hastings, Classical and quantum bounded depth approximation algorithms, *ArXiv:1905.07047* (2019).
- [54] C. Huang, *et al.*, Efficient parallelization of tensor network contraction for simulating quantum computation, *Nat. Comput. Sci.* **1**, 578 (2021).
- [55] J. Gray and S. Kourtis, Hyper-optimized tensor network contraction, *Quantum* **5**, 410 (2021).
- [56] M. W. Johnson, *et al.*, Quantum annealing with manufactured spins, *Nature* **473**, 194 (2011).
- [57] A. D. King, *et al.*, Coherent quantum annealing in a programmable 2000 qubit Ising chain, *Nat. Phys.* **18**, 1324 (2022).
- [58] L. Caha, A. Kliesch, and R. Koenig, Twisted hybrid algorithms for combinatorial optimization, *Quantum Sci. Technol.* **7**, 045013 (2022).
- [59] F. Wagner, J. Nüßlein, and F. Liers, Enhancing quantum algorithms for quadratic unconstrained binary optimization via integer programming, *ArXiv:2302.05493* (2023).
- [60] <https://github.com/jernejrudifinzigar/qiro>.
- [61] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, *ArXiv:1602.07674* (2019).
- [62] E. Bae and S. Lee, Recursive QAOA outperforms the original QAOA for the MAX-CUT problem on complete graphs, *ArXiv:2211.15832* (2023).
- [63] R. Niedermeier and P. Rossmanith, New upper bounds for maximum satisfiability, *J. Algorithms* **36**, 63 (2000).
- [64] R. S. Andrist, M. J. A. Schuetz, P. Minssen, R. Yalovetzky, S. Chakrabarti, D. Herman, N. Kumar, G. Salton, R. Shaydulin, Y. Sun, M. Pistoia, and H. G. Katzgraber, Hardness of the maximum independent set problem on unit-disk graphs and prospects for quantum speedups, *ArXiv:2307.09442* (2023).
- [65] J.-G. Liu, L. Wang, and P. Zhang, Tropical tensor network for ground states of spin glasses, *Phys. Rev. Lett.* **126**, 090506 (2021).
- [66] V. Bergholm, *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, *ArXiv:1811.04968* (2022).
- [67] H. G. Katzgraber, Introduction to Monte Carlo methods, *ArXiv:0905.1629* (2011).