

Received 2 May 2024; revised 24 July 2024; accepted 31 July 2024; date of publication 14 August 2024;
date of current version 16 October 2024.

Digital Object Identifier 10.1109/TQE.2024.3443660

Solving Nonnative Combinatorial Optimization Problems Using Hybrid Quantum–Classical Algorithms

JONATHAN WURTZ¹ , STEFAN H. SACK^{1,2} , AND SHENG-TAO WANG¹ 

¹QuEra Computing Inc., Boston, MA 02135 USA

²Institute of Science and Technology Austria, 3400 Klosterneuburg, Austria

Corresponding author: Jonathan Wurtz (e-mail: jonathan.wurtz@gmail.com)

This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) Optimization with Noisy Intermediate-Scale Quantum Devices program under Grant W911NF2010021 and in part by the DARPA Small Business Technology Transfer program under Award 140D0422C0035.

ABSTRACT Combinatorial optimization is a challenging problem applicable in a wide range of fields from logistics to finance. Recently, quantum computing has been used to attempt to solve these problems using a range of algorithms, including parameterized quantum circuits, adiabatic protocols, and quantum annealing. These solutions typically have several challenges: 1) there is little to no performance gain over classical methods; 2) not all constraints and objectives may be efficiently encoded in the quantum ansatz; and 3) the solution domain of the objective function may not be the same as the bit strings of measurement outcomes. This work presents “nonnative hybrid algorithms”: a framework to overcome these challenges by integrating quantum and classical resources with a hybrid approach. By designing nonnative quantum variational ansatzes that inherit some but not all problem structure, measurement outcomes from the quantum computer can act as a resource to be used by classical routines to indirectly compute optimal solutions, partially overcoming the challenges of contemporary quantum optimization approaches. These methods are demonstrated using a publicly available neutral-atom quantum computer on two simple problems of Max k -Cut and maximum independent set. We find improvements in solution quality when comparing the hybrid algorithm to its “no quantum” version, a demonstration of a “comparative advantage.”

INDEX TERMS Combinatorial mathematics, optimization, quantum advantage, quantum algorithm, quantum computing.

I. INTRODUCTION

Combinatorial optimization is a challenging task that is widely applicable across a range of domains, from logistics, finance, operations, chemistry, computer science, and more. The underlying problem is deceptively simple: given some problem, such as protein folding or vehicle routing, find a solution in the domain of some objective function that extremizes that function, such as a protein conformation that minimizes energy or delivery route that maximizes utility. Typically, the objective is some low-degree polynomial or other efficient-to-compute function $C(\chi) : X \rightarrow \mathbb{R}$ that maps some domain of valid solutions $\chi \in X$ to a value capturing the quality of that solution.

These problems are combinatorial, in the sense that the domain X of possible solutions is a discrete set given by combinations of subunits, typically growing exponentially

in the problem size. This deceptively simple formulation hides extreme complexity; often, general NP-complete problems can be reduced to these optimization problems, which implies that efficient solvers for these problems may collapse the polynomial hierarchy and imply complexity class equivalence $P = NP$. Under the widely accepted assumption that $P \neq NP$, classical optimizers must resort to heuristic “guess and check” algorithms, which search through a large range of candidate solutions and guarantee optimal results only in exponential time.

A promising alternative to classical optimizers is to use quantum computing [1], which leverages quantum phenomena to generate approximate or optimal solutions faster or of better quality than existing classical methods. These algorithms fall into two main categories. The first category of parameterized quantum circuits [2] uses a digital circuit

model parameterized by variational controls $\vec{\alpha}$ to prepare a state whose probability distribution over measurement outcomes $P(z) = |\langle z|\vec{\alpha}\rangle|^2$ is biased toward good bit string solutions $z \in X$, maximizing, for example, the average objective function value $\langle C(z) \rangle = \sum_z P(z)C(z)$. Typically, the structure of the problem instance is directly inherited in the parameterized circuit; for example, the quantum approximate optimization algorithm (QAOA) [3] and its generalizations [4] alternate between some global mixing unitary and a diagonal target unitary generated by accumulating phase proportional to the objective function value $\hat{U} = \sum_z \exp(-i\gamma C(z))|z\rangle\langle z|$. If the objective function has an efficient low-degree polynomial representation with few hard constraints, such as a quadratic unconstrained binary optimization (QUBO), each clause may be represented by some k -local Z-phase gate. Alternatively, hardware-native ansatzes [5] inherit little or no structure from the underlying problem, instead targeting particular hardware-constrained connectivity and optimizing over a larger set of variational parameters.

The second category of adiabatic quantum computation [6] uses coherent analog dynamics or quantum annealing to prepare the ground state of some physical quantum Hamiltonian. By designing the interaction terms, the ground state directly or indirectly encodes the maxima of the objective function $C(z)$.

Both approaches run into three challenges that may otherwise block commercial-scale adaptation, whose problems typically do not fall into clear problem classes. First is that quantum algorithm performance over contemporary classical optimization may be small or nonexistent [7], with challenging performance guarantees or sampling rates required to beat best-in-class routines [8].

Second is that it may be difficult to encode every constraint and objective directly into the parameterized circuit or Hamiltonian; for example, a problem may have long-range connectivity that is not native to the hardware implementation [9], include some high-order polynomial term that cannot be efficiently reproduced by a hardware native k -local gate set [10], or require hard constraints that may be difficult to encode [11], [12]. These incomplete mappings may reduce the performance of the quantum algorithm, as the Hamiltonian ground state may not encode the combinatorial solution, or the parameterized circuit may lose some important context in the problem structure.

Third and most crucially, the domain X of problem solutions may not be the domain of bit strings $\{0, 1\}^N$, so even the implementation of a direct encoding [13] may be unclear or suffer from large qubit overhead [14]. For example, a vehicle routing problem corresponds to an ordered list of locations for a vehicle to travel, while encodings exist [15]. There is a qubit overhead in mapping from binary outcomes to valid solutions. This challenge is a particular blocker for the utility-scale use of quantum computers solving optimization problems, as realistic instances rarely directly have a binary solution domain.

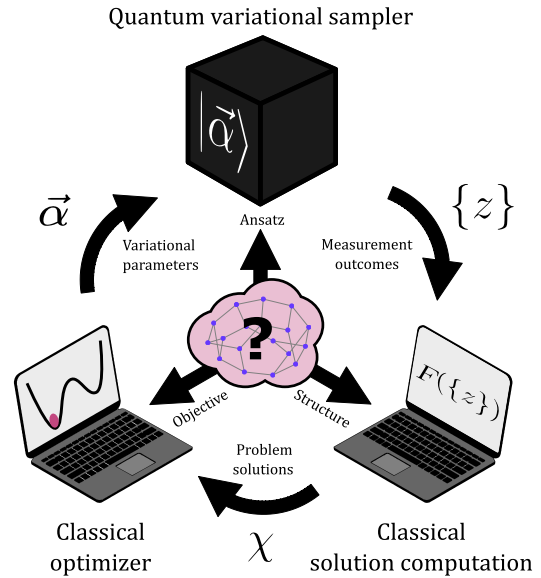


FIGURE 1. Outline of nonnative hybrid algorithms. A given problem (center) motivates some problem-dependent ansatz $|\vec{\alpha}\rangle$. The state is prepared and measured on a quantum computer (top), generating samples $\{z\}$ over the probability distribution $P(z) = |\langle z|\vec{\alpha}\rangle|^2$. These samples are sent to a classical algorithm F (bottom right), which uses the bit strings to compute a problem solution $F(\{z\}) = \chi$, by postprocessing individual bit strings, using expectation values as a “quantum hint,” or using the distribution. These classically computed solutions, which do not have to be bit strings, are fed into a classical variational optimizer (bottom left), which optimizes variational parameters $\vec{\alpha}$ to maximize the objective function $\langle C(\chi) \rangle$. The hybrid algorithm then continues in this loop until convergence.

These challenges are particularly acute when considering real-world optimization problems. These problems are complicated, typically having nontrivial mixed-integer solution spaces with many constraints and multiple objectives that must be optimized in a dynamic or time-constrained manner. Furthermore, these problems typically do not have a natural reduction to specific combinatorial optimization problems, such as MaxCut or MaxSat, and require some reduction, typically to QUBO at a high overhead cost [16].

This article proposes a framework called “nonnative hybrid algorithms” (NNHA) that attempts to circumvent these challenges by maximally leveraging both quantum and classical computational resources to solve generic potentially nonnative combinatorial optimization problems, as well as proposing a new metric of “comparative advantage” over classical algorithms. The key idea, as illustrated in Fig. 1, is to use the quantum device as a variational sampler [17], generating bit strings $\{z\}$ from measurement outcomes from some hardware native parameterized ansatz that inherits some but not all structure of the nonnative target problem. These bit strings are then used as a resource for a classical algorithm, which uses direct measurement outcomes, correlation functions, or probability distribution to compute some approximate or exact solution $\chi = F(\{z\})$ to the target problem. This contrasts with contemporary quantum optimization, where bit strings directly encode solutions and classical resources only decode bit strings to solutions. In this way, the

quantum resources act as a coprocessor or hardware accelerator [18] for the overall computation, solving subproblems or generating biased probability distributions instead of directly solving the target problem.

While the terminology of “hybrid computing” is well accepted [19], this article proposes an extension of the hybrid architectures beyond the paradigm of classical resources simply seeking good variational parameters for some ansatz state [20]. Instead, the key idea is that classical resources contribute actively to both parameter finding and directly constructing candidate solutions to the target problem. This work augments other ideas, such as greedy [21], [22], recursive [23], [24], [25], multilevel [26], [27], branch-and-bound [28], genetic [29], preprocessing [30], [31], or linear programming [32] approaches, and shares similarity to modern reservoir [33], extreme learning [34], and generative [35] machine learning approaches.

The rest of this article is organized as follows. Section II outlines the general structure of the hybrid optimization scheme and various forms (types 1–4) of classical coprocessing. Section III demonstrates a type-1 algorithm for the MaxCut problem, which inherits performance guarantees of the classical subroutine as a partial solution to challenge (1) of contemporary quantum optimization. Section IV demonstrates a type-2 algorithm for Max k -Cut, where the solution domain is not bit strings, a solution to challenge (3). Finally, Section V demonstrates a type-1 and type-3 algorithm for the maximum independent set (MIS) problem, where hard constraints cannot be completely imposed by quantum ansatzes, a solution to challenge (2). Examples IV and V also include implementation on QuEra’s publicly available neutral-atom quantum computer, Aquila [36], and are reproducible via code notebooks here [37]. All examples are intended as prototype demonstrations of the framework instead of general solutions to specific problems.

II. METHODS

An NNHA algorithm leverages two key components. The first is the quantum variational sampler, which efficiently generates bit strings $\{z\}$ from some probability distribution P parameterized by variational parameters $\vec{\alpha}$ and problem instance I

$$\{z\} \leftarrow P(z|\vec{\alpha}) = \langle z|\hat{\rho}(\vec{\alpha})|z\rangle \quad (1)$$

where $\hat{\rho}$ is the state generated by the quantum ansatz $\hat{\rho} \approx |\vec{\alpha}\rangle\langle\vec{\alpha}|$ parameterized by problem instance I given a particular algorithm design. Note that there are no strict requirements on coherence or even that the underlying probability distribution be drawn from a quantum device. A quantum device serves as an efficient generator of measurements and may be replaced by, e.g., emulation [38], graph neural networks [39], classical Boltzmann machines [40], and so forth [41]. A “no-quantum” or classical version of a hybrid algorithm is represented by a state $\hat{\rho}(\vec{\alpha}_c)$ that is efficient to sample classically with a set of parameters $\vec{\alpha}_c$. A natural extension samples from multiple parameterizations $\vec{\alpha}$, so that each bit string is

the concatenation of each measurement outcome in the set of parameterizations. For example, an ansatz may sample individual bit strings from a range of nonequilibrium evolution times or from several hardware-native best fits of the problem onto a quantum circuit [42].

A. CLASSICAL COMPUTATION

The second key component of NNHA is classical computation, which serves two main functions. The first is to use bit strings generated from k measurement outcomes into valid problem solutions, through some function

$$F : \{\{0, 1\}^n\}^k \mapsto X \quad (2)$$

which maps ensembles of bit strings to solutions $F(\{z\}) = \chi$. The particulars of such a function should be problem dependent but fall into four main types.

Type 1: Computation that “postprocesses” individual candidate bit strings to generate individual solutions. For example, the function could use a measurement outcome to directly “warm start” a classical heuristic optimizer, which refines the initial solution into a locally better solution [43]. This class of hybrid computation is exemplified in Section III, which uses a local greedy heuristic to find a local maximum in the energy landscape.

Type 2: Computation that uses expectation values $\langle\hat{z}\rangle$ and correlation functions $\langle\hat{z}_i\hat{z}_j\rangle$, including higher order correlation functions, e.g., $\langle\hat{z}_i\hat{z}_j\hat{z}_k\hat{z}_l\rangle$, as a “quantum hint.” The values of these expectation values are problem-structure dependent, and shallow circuits may be used to infer the local structure of the problem instance. For example, the connected correlation function can be shown to be strictly local [44] for certain ansatzes, and intermediate-depth evolutions can generate correlations within some lightcone that may be larger than a k -local classical algorithm can handle. This class of hybrid computation is exemplified in Section IV, which uses the low-energy eigenvectors of the connected correlation matrix to solve Max k -Cut.

Type 3: Computation that directly uses the distribution of measurement outcomes $\{z\} \leftarrow P(z|\vec{\alpha})$ as a “reservoir.” Certain quantum probability distributions can be hard to sample, which suggests that samples from this distribution may be a natural resource for optimization. For example, if the distribution is known to be approximately thermal, one could use parallel tempering and replica exchange to uniformly sample low-temperature distributions with simulated annealing. An optimization-focused example is exemplified in Section V, where an MIS is optimized using cluster update simulated annealing.

Type 4: While not covered in this article, the notion of “iterative refinement” or recursive schemes can be considered as an NNHA and is the subject of recent research [21], [22], [23], [24], [25]. Under a recursive method, measurement outcomes or expectation values are used to fix a subset of problem variables, reducing the problem size. Then, the reduced problem is fed back recursively until the problem size is reduced to zero. While these schemes are “hybrid,”

they differ from types 1–3 as they require iterative calls to quantum resources with different ansatzes instead of static distributions $P(z|\vec{\alpha})$.

B. PARAMETER OPTIMIZATION

The second function of classical computation is variational optimization. Typically, the computational problem can be formulated in terms of a maximization problem for a cost or objective function $C(\{z\})$ [45]. However, because bit strings are not necessarily valid solutions, the optimizer must search not over bit strings, but instead over classically computed solutions $C(F(\{z\}))$, so the optimizer finds parameters $\vec{\alpha}$ that extremize

$$\text{MAX}_{\vec{\alpha}} \langle C(F(\{z\})) \rangle \quad \text{with} \quad \{z\} \leftarrow P(z | \vec{\alpha}) \quad (3)$$

where the brackets $\langle \dots \rangle$ denote the average over many samples. The goal of the classical processor is thus to optimize the variational parameters $\vec{\alpha}$ such that the objective function is maximized. This task is typically carried out by a classical optimization algorithm.

The timescale for operations on a quantum processor is typically orders of magnitude slower than on a classical processor, which motivates the careful choice of an optimization algorithm that is not only resilient to noise but also resourceful and requires few calls to the quantum processor for convergence. We found the constrained optimization by linear approximations (COBYLA) method to be a suitable choice for this task [46], though other optimization methods may be appropriate [47]. COBYLA does not require the computation of a gradient, which makes it more resilient to hardware noise. The method works by constructing an approximation model of the objective function and constraints in the neighborhood of the current point, which is maximized and updated with each iteration until convergence.

For each iteration, the cost function value is estimated from a set of M measurement outcomes $\{z\}$. While the variance of the expectation values typically scales as $\sim 1/\sqrt{M}$, we find that optimization with fewer shots per cost estimation can lead to similar or better performance than optimization with many shots for the noise present on neutral-atom quantum hardware. The main driver for good performance is the number of iterations of the optimization algorithm rather than an accurate estimation of the objective function.

While on superconducting hardware, 10^3 – 10^4 shots are typically used for the estimation of the objective function [48], we found that as few as ~ 10 shots to work well for optimization on neutral-atom quantum hardware. In the data presented in this article, we typically use 100 shots. For the variational ansatzes used in this work, we find good convergence within 30 iterations and 3000 total shots on the quantum processor. The few shots for convergence may be a feature of analog quantum computation where the number of variational parameters is limited, and the ansatz structure is simple. We expect that for a more general digital

ansatz, such as the hardware-efficient ansatz [49] or unitary coupled cluster ansatz [20], more iterations are required for convergence.

C. COMPARATIVE ADVANTAGE

A crucial caveat of hybrid algorithms comes in comparing performance between hybrid and classical algorithms. If the hybrid optimization performance is better than some other classical method, is this due to the specific problem subclass, the classical part of the hybrid algorithm, or a clear “quantum advantage” [50]? This work uses the insights of designing algorithms with a classical or no-quantum limit [30]. Given some set of variational parameters, the quantum sampling subroutine may be replaced by an efficient and equivalent classical subroutine, which generates a classical-only variant of the algorithm. For example, the ansatz could be a simple-to-sample product state or matrix product state. In this way, the performance guarantee of the hybrid algorithm is inherited by that of the classical subroutine. The performance of the algorithm is then compared between the classical-only limit, and the full hybrid implementation that does have quantum resources. If the presence of quantum resources increases performance, then there is a “quantum boost” [22], signaling a comparative advantage in including these resources in computation.

III. MAXCUT WITH QAOA AND GREEDY FLIP

A simple demonstration of a type-1 NNHA can be shown as applied to the MaxCut problem using an enhancement of the QAOA. For this example, the choice of problem, ansatz, and the classical algorithm is explicitly simple to pedagogically introduce the concept as well as demonstrate the notion of “comparative advantage.” This example can be expanded to more complicated examples as shown in later sections, as well as extended to designing bespoke solvers for application-scale solutions and is intended to serve as an initial inspiration for in-depth bespoke solutions.

MaxCut is an unconstrained binary optimization on graphs. Given a graph G of vertices V and edges E , MaxCut strives to find the bipartition of vertices $\{V_+, V_-\}$ such that a maximum number of edges have one vertex in each partition (“cut”). MaxCut is a prototypical NP-hard optimization problem, with classical optimization limited performance guarantees ≤ 0.878 in the worst case [51].

The QAOA is the prototypical variational quantum optimization algorithm [3], which generates an ansatz that alternates p times between some mixer B and target C generator with $2p$ variational parameters $\vec{\alpha} = \{\vec{\beta}, \vec{\gamma}\}$

$$|\vec{\alpha}\rangle = \prod_{j=1}^p e^{-i\beta_j \hat{B}} e^{-i\gamma_j \hat{C}} |+\rangle. \quad (4)$$

The target generator \hat{C} has the property of sharing its eigen-spectra with the objective function $\hat{C} = \sum_z C(z)|z\rangle\langle z|$, and the initial state $|+\rangle$ is the maximal eigenstate of the mixing




FIGURE 2. Structure and performance for type-1 NNHA using postprocessed QAOA, solving MaxCut on a random ensemble of 256 3-regular 16-vertex graphs. (Top) Action of the greedy flip algorithm. The greedy algorithm iterates through vertices based on some initial bipartition based on quantum measurement (left) and tries to flip partitions based on the 1-local environment, accepting if the objective increases. The process repeats until no further flips can increase the energy, finding a local maximum (right). (Bottom) Comparative performance of hybrid, quantum, and classical algorithms. The line is the average expected approximation ratio of each graph over the ensemble, while the violin plots show the distribution of expectation values over the ensemble. Observe that the hybrid greedy postprocessing method (purple) outperforms base QAOA (red) and begins to outperform the $p = 0$ classical-only limit (black dashed) at $p \sim 4$, indicating an ensemble performance advantage.

operator, which is typically taken to be a uniform sum of X terms $\hat{B} = \sum_j \hat{\sigma}_x^j$.

MaxCut and its generalization of QUBO is a natural target problem for the QAOA [52], [53], due to the simple translation from measurement outcomes to problem solutions. The space of bit string measurement outcomes $\{0, 1\}^n$ is in one-to-one correspondence with MaxCut bipartitions, by simply assigning each vertex to the bipartition given by the measurement outcome of the corresponding qubit $0 \mapsto V_+$ and $1 \mapsto V_-$. This fact will not be true for later problems.

While QAOA has limited performance guarantees in the small p limit [44] and converges to exact solutions in the large p counterdiabatic limit [54]. The algorithm is typically heuristic and struggles to beat contemporary classical optimizers on a range of problems. A local classical algorithm can strictly outperform $p = 2$ QAOA [7], and a simple 1-local greedy heuristic can outperform even $p \sim 6$ base QAOA [8], as seen in Fig. 2.

QAOA may be enhanced by merging it with a local greedy algorithm using a type-1 NNHA using “postprocessing” [43]. The algorithm includes three parts: the ansatz, the variational optimizer, and the classical solution generation. The

ansatz is the base QAOA of (4), where the objective generator is the sum of quadratic terms $\hat{C} = \sum_{(ij) \in G} \hat{\sigma}_z^i \hat{\sigma}_z^j$ over all edges in the graph G , and the mixing generator is the uniform sum over Pauli X terms, $\hat{B} = \sum_i \hat{\sigma}_x^i$. To simplify the analysis, this demonstration uses fixed angles optimized for ν -regular graphs [55], bypassing the variational optimization step.

The classical postprocessing algorithm is chosen to be a 1-local greedy flip heuristic [56], shown in Fig. 2. The algorithm is initialized by a single measurement outcome, which is inferred to be an initial candidate solution by mapping a bit string to a bipartition $0 \mapsto V_+$ and $1 \mapsto V_-$. Next, the solution is postprocessed with the greedy algorithm. Each vertex is trialed to be “flipped” to be in the opposite partition. Each of these trials increases or decreases the objective value of the solution, depending on the local solution from neighboring vertices. The greedy solver simply picks the vertex flip with the largest increase in objective value; tied vertices are selected at random. The process iterates until no local flip increases the objective value and the algorithm returns a solution. If the algorithm instead chooses a random vertex instead of trialing all vertices, this update rule is the same as zero-temperature simulated annealing; a generalization beyond the scope of this work could include implementation of finite-temperature simulated annealing (e.g., replica exchange); see Section V for more details.

The greedy flip heuristic has a performance guarantee of $\lceil \nu/2 \rceil / \nu$ and converges to a solution in a time of order N^2 and memory of order N (with a simple generalization to time of order N using zero-temperature simulated annealing [57]). The hybrid algorithm can recover the classical-only limit by choosing all $\beta_i = 0$ or doing $p = 0$ layers, in which case the sampling probability is uniform $P(z) = 2^{-N}$, which is trivial to sample classically. Note that while it may be efficient to simulate expectation values of QAOA using tensor networks or other methods that leverage the strict locality of QAOA, sampling from certain QAOA distributions can be shown to be in the class #P [58], though the reduction does not work for specific instances such as 3-regular MaxCut.

A comparative demonstration of this hybrid algorithm is shown in Fig. 2 using emulated state-vector simulation. Results are characterized using an ensemble of 256 random 3-regular graphs with 16 vertices. There are a few key takeaways from these results. First, the classical $p = 0$ limit of the algorithm outperforms the base QAOA even for depths $p \sim 6$, highlighting the poor performance of quantum-only algorithms against even the simplest classical heuristics.

Second, the shot-to-shot objective value of the hybrid algorithm is strictly greater than or equal to that generated from base QAOA measurements due to the greedy action of postprocessing. In this way, any advantage in terms of performance improvements over base QAOA is self-evident. This performance gain is clear when comparing the red and purple lines of Fig. 2. While this may be a trivial fact, this notion of maximally using classical resources is the key insight of this work.

Third, for advancing depths of QAOA and thus more quantum resources, the hybrid algorithm heuristically appears to have better performance than both the classical or quantum algorithms individually, as shown by the rise above the black dashed classical-only limit. Unlike the performance increase over quantum-only, this increase is not self-evident and is only a heuristic observation. In principle, the probability distribution over measurement outcomes for the chosen fixed angles may maliciously sample from regions in solution space that have many low-quality local maxima, reducing the performance of the greedy heuristic below the classical-only limit. This phenomenon occurs and is further explored in Section V-A. It is interesting to explore how optimal parameters may change when optimizing over postprocessed objective values, but such a study is beyond the scope of this work.

This simple first example serves only as an initial demonstration of NNHA. There may be many interesting future directions in investigating different problem classes, postprocessing methods, approximate or hardware-native ansatzes, variational optimizations, benchmarking against state-of-the-art optimizers, as well as applying this postprocessing to real-world optimization problems. In addition, type-1 postprocessing can be used as a subroutine for other NNHA, as will be used in every proceeding example.

IV. MAX k -CUT WITH QUENCH DYNAMICS AND SPECTRAL CLUSTERING

An extension of MaxCut is the Max k -Cut problem [24], [27], [59], [60], which seeks to find k disjoint subsets of a graph such that a maximum number of edges have vertices of different subsets. For $k = 2$, the problem reduces to the binary MaxCut problem, and if $k \geq \chi(G)$, the chromatic number of the graph, the optimal solutions are also k -coloring. Due to MaxCut and graph coloring reducing to Max k -Cut, which are NP-complete, the decision variant of Max k -Cut is also NP-complete.

As discussed in Section III, $k = 2$ -MaxCut is a natural problem for quantum optimization since the domain of solutions is directly mapped to bit strings, and the objective function is naturally mapped in ZZ or parity gates. $k \geq 3$ faces the challenge that the direct mapping from bit strings to solutions does not exist. While encoding strategies, such as 1-hot exist [14], they may become unwieldy for larger k and suffer from limited qubit numbers and coherence.

For these reasons, NNHA may be a viable solution to efficiently extract approximate or exact Max k -Cut solutions using nonnative quantum hardware. In this section, we propose an example or prototype type-2 NNHA using expectation values and correlation functions that integrate two classical routines: spectral clustering and type-1 greedy postprocessing. Keeping with the design intent of NNHA, this implementation will have a natural no-quantum limit and includes both classical solution generation and variational optimization steps. This example implementation should be seen as an example type-2 NNHA, with a specific method

and target problem, and should be considered as inspiration instead of generalization to more complicated applications.

A common graph analysis method is spectral graph theory [61]. This example implementation takes strong inspiration from spectral clustering, which is a method to partition nodes based on eigenanalysis-based features [62]. The classical algorithm works in the following steps.

- 1) Given a graph G , compute the eigenvalues E_λ and eigenvectors $V_{i\lambda}$ of the graph Laplacian $J(G)$ or adjacency matrix. The eigenvectors can be interpreted as the “spring modes” of the graph if every edge can be considered to be a harmonic potential between vertices. The lowest energy eigenvectors correspond to long-wavelength spring modes with many vertices moving together and the highest energy eigenvectors correspond to short-wavelength spring modes with adjacent vertices moving opposite.
- 2) Identify the largest k eigenvectors V of the Laplacian as a feature vector F_i for each vertex i

$$\vec{F}_i = (V_{i,N}, V_{i,N-1}, \dots, V_{i,N-k}). \quad (5)$$

If the smallest k eigenvectors are selected, this method will construct an approximate k -clustering.

- 3) Given the feature vector \vec{F}_i of each vertex i , implement a k -clustering algorithm to each vertex to identify each vertex with a color. This implementation uses a simple k -means clustering. The output is a vector S_i of colors as a candidate Max k -Cut solution.

This solution can be enhanced by type-1 greedy postprocessing of the classical output, similar to Section III. The algorithm iteratively tries every single-vertex color update, selects the one that decreases the cost function the most, and repeats until no further updates are possible. This algorithm is a generalization over Section III to multiple colors.

This algorithm can be modified and augmented using a quantum computer with a type-2 NNHA that uses connected correlation functions as a resource for clustering. Instead of using the eigenvectors of the connectivity matrix as the feature vector, we propose to use the eigenvectors of the weighted sum of connected correlation functions

$$C_{ij}(\alpha, \lambda) = \sum_{i=1}^n \lambda_i (\langle \vec{\alpha}_i | \hat{n}_i \hat{n}_j | \vec{\alpha}_i \rangle - \langle \vec{\alpha}_i | \hat{n}_i | \vec{\alpha}_i \rangle \langle \vec{\alpha}_i | \hat{n}_j | \vec{\alpha}_i \rangle) \quad (6)$$

where $|\vec{\alpha}_i\rangle$ are n wave functions of quantum variational parameters $\vec{\alpha}$ and classical variational weights λ_i . This is similar to [63], where the direct sign-based partition used for a $k = 2$ MaxCut is replaced by a more general k -means clustering as required for $k > 2$ colors. There is of course flexibility in the design of ansatz $|\vec{\alpha}\rangle$; in this case, we use a nonequilibrium quench of the Rydberg neutral-atom Hamiltonian, where each atom is placed at a position scaled from vertex positions on the original graph. For more details, see code notebooks here [37]. This ansatz requires a constraint on the viable problem instances solved with this algorithm: they

must be 2-D geometric with constrained local connectivity, and each vertex must have a position assigned in the 2-D space.

This work selects problem instances that are King’s subgraphs with a 30% random dropout similar to the graph ensembles of [64]. We select a $k = 3$ solution space, as $k = 2$ is the prototypical MaxCut problem and any Kings graph can be trivially $k = 4$ -colored with a tiling.

We choose the variational ansatz $|\vec{\alpha}_i\rangle$ to be three different quenches from the ground state of the analog neutral-atom Hamiltonian. Each wave function is parameterized by $\alpha_i = (t_i, \Delta_i)$, where $t \in [0, 4]\mu\text{s}$ is the total evolution time and $\Delta \in [0, 15] \text{ rad}/\mu\text{s}$ is the detuning. We select a lattice spacing of $4.8\mu\text{m}$ between sites, and a Rabi drive of $\Omega = 15.0 \text{ rad}/\mu\text{s}$, corresponding to a dynamical blockade radius of $R_b \approx 6.7\mu\text{m}$. We find that $M = 100$ measurements are adequate to recover statistics. Computations are implemented using the QuEra computer “Aquila” through Amazon Braket. For more details, see [36] and example notebook [37].

Each of the six quantum and three classical variational parameters can be simultaneously optimized. Due to the cost of quantum optimization, this is done dimensionwise: for every set of quantum parameters α_i , the classical parameters λ_i are optimized using Bayesian methods. Then, the parameter space of quantum parameters is explored using a bounded stochastic search.

Such an ansatz is well physically motivated. At short times and short-range connected quantum systems, correlations and entanglement spread ballistically [65] or diffusively [66] within a small distance. Thus, at times of order $t \sim 1/\Omega$, the connected correlation function between vertex i and others will only be nonzero over a region of radius ~ 1 unit cell away. Thus, the eigenvectors should recover some of the local structure of the original graph. By varying the times t_i of each quench, the ansatz probes various local length scales. By varying the weights λ_i of the different quenches, one can “target” a particular length scale. Finally, by varying the detunings Δ_i of the different quenches, one can infer the local connectivity of each vertex due to the Rydberg blockade effect. In totality, the variational ansatz exploits local correlations and entanglement in the system to extract quasi-local graph structure in a way that may be difficult to reproduce with an equivalent classical analysis.

To recap, the type-2 hybrid spectral clustering algorithm works given the following steps, as shown in Fig. 3. Given some Max k -Cut instance on a unit disk graph G with vertex positions \vec{x}_i , M bit strings are measured from three variational ansatzes of two parameters each α_i , which is a quench of the Rydberg neutral-atom Hamiltonian. Given nN bit string measurements, the connected correlation function weighted by n classical parameters λ_i is computed. The k largest eigenvectors of $C(\alpha, \lambda)$ are used as a feature vector \vec{F}_i for a k -means clustering to identify each vertex with a color, and finally, the candidate result is postprocessed to be locally optimal with a zero-local greedy flip algorithm. While the expectation

values are deterministic (in the $N \rightarrow \infty$ measurement limit), candidate solutions are nondeterministic due to the stochastic nature of k -means clustering, flip, and finite-sample averaging. Finally, the variational parameters are optimized using Bayesian and stochastic optimizers.

This algorithm, as all NNHA must, has a no-quantum limit. If $t = 0$, then the connected correlation function is trivially $C(\alpha, \lambda) = 0$, and thus, the eigenspectra are degenerate, eigenvectors are Haar random, and the clustering generates a random guess for each vertex color.

Results for this implementation as implemented on QuEra’s cloud-accessible quantum computer “Aquila” [36] are shown in Fig. 3 for a representative graph and $k = 3$. It is clear that the hybrid version of the algorithm outperforms the classical-only limit. This is a clear demonstration of comparative advantage: the presence of quantum resources increases the performance of the algorithm. However, this is not “quantum advantage,” as other structure-aware classical algorithms may be able to beat this heuristic performance.

Finally, it should be noted that this example should not be used directly for solving Max k -Cut optimization problems. Instead, it should serve as a *template* for other implementations of type-2 NNHA. Other problems with more complicated objectives and geometries may require other variational ansatzes or classical steps, for which the insight of their design may only come from an understanding of the problem and the action of existing classical optimization routines. Such exploration is beyond the scope of this work.

V. MIS WITH QUANTUM ANNEALING AND CLUSTER SIMULATED ANNEALING

For the final example, we solve the MIS problem using a type-3 NNHA. For completeness, we also discuss the recent results of Ebadi22 [64], which can be retroactively interpreted as a hardware demonstration of a type-1 postprocessing NNHA, though the encoding itself is very close to native. This example extends the results of that work in a type-3 NNHA by using measurement outcomes as a reservoir for a cluster simulated annealing step.

The MIS problem is a combinatorial graph problem that seeks to find the largest subset of vertices of some graph G such that no two vertices share an edge. The MIS is an NP-complete optimization problem [67] and has worst-case classical performance guarantees that scale logarithmically due to the “overlap gap” property [68]. MIS is a good example of constraint mismatch in quantum algorithms: it is natural to map a bit string z to a set $S = \{i \text{ if } z_i = 1\}$. However, not every bit string is a valid independent set (for example, the all-ones bit string), and generating an ansatz with such hard constraints may be difficult. For this reason, some amount of classical processing must happen to get valid solutions. A subset of independent sets are maximal independent sets, which are independent sets such that adding any vertex to the set makes the set no longer an independent set. By definition, the MIS must also be a maximal independent set. Certain bit

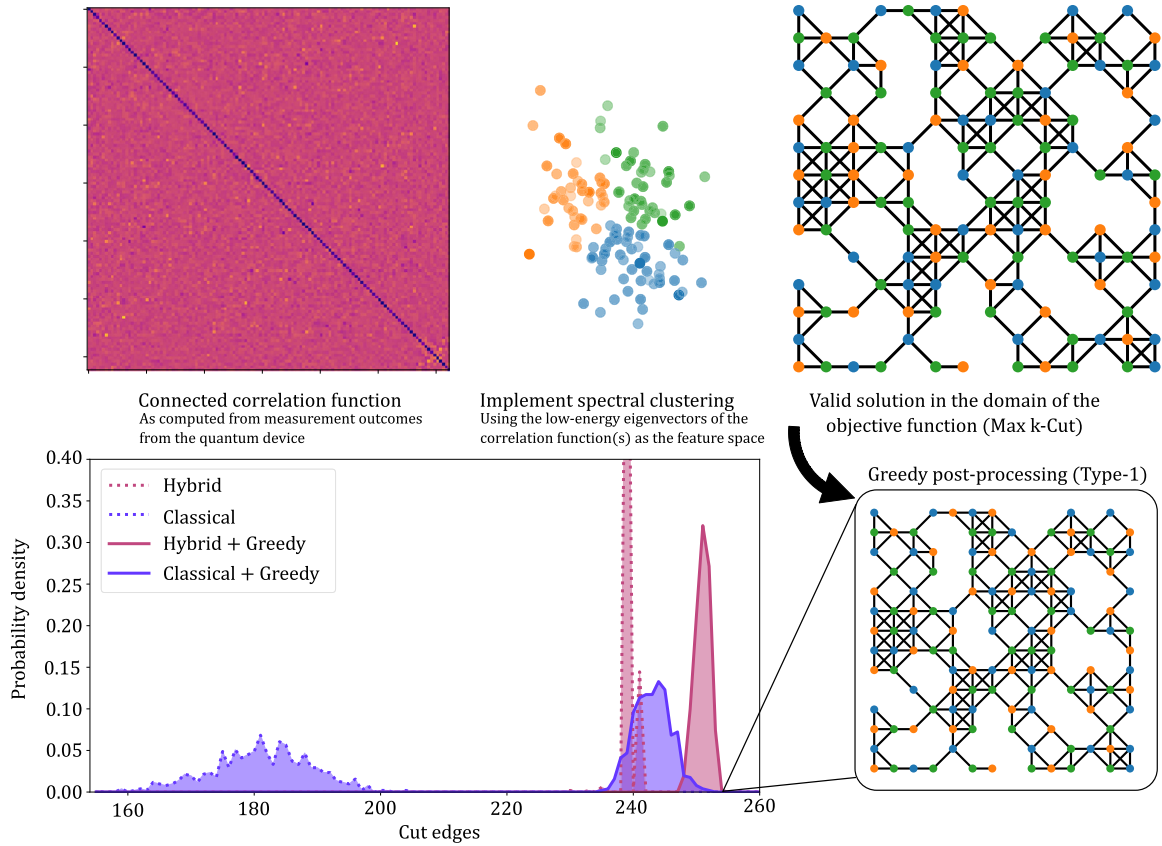


FIGURE 3. Structure and performance of spectral clustering solving the Max k -Cut problem on unit disk graphs, as implemented on QuEra’s cloud-accessible quantum computer “Aquila.” Top left plots the variationally optimized connected weighted correlation function (6) for the particular problem instance. The three lowest eigenvectors of this matrix serve as a feature vector for k -means clustering (6) which generates a candidate Max k -Cut solution (top right). Each candidate solution can be further postprocessed with a zero-local greedy flip algorithm to find a locally optimal solution. The bottom left plots the objective distribution of solutions. The quantum version outperforms the classical-only limit (bottom left) of a greedy postprocessed random choice in finding more optimal solutions, which is an indicator of comparative advantage.

strings (for example, the all-zeros bit string) do not map to valid maximal independent sets. A natural solution space X for MIS is thus the set of all maximal independent sets.

A. EBADI22 SOLUTION

Recent seminal results by Ebadi22 [64] solve MIS on a subset of graphs, called unit disk graphs, using neutral-atom quantum computers. The key insight is that the strong van der Waals interaction between neutral atoms excited to a Rydberg state causes a strong energy shift preventing nearby atoms from both being in the doubly excited states, in a phenomenon called the “Rydberg blockade.” Given some choice of detuning (e.g., Z field) and placing atoms at scaled positions of the vertices of the target graph, the Rydberg atom Hamiltonian can encode the MIS of a subset of graphs, called “unit disk graphs” in its ground state, by matching the Rydberg blockade radius to the unit disk radius [69]. The ground state can be prepared adiabatically using protocols such as those of Fig. 4; for more details, see [36] and [64].

Although the ground state theoretically encodes the solution, the prepared state may not be the ground state, due to decoherence, SPAM errors, diabatic effects from the ansatz, and more. This is a constraint or objective mismatch between

the target problem and the underlying ansatz. For this reason, the measurement outcomes $\{z\}$ may not directly map onto maximal independent sets and require some level of type-1 postprocessing. Ebadi22 [64] used a zero-local greedy heuristic, which works similarly to what follows.

Given some graph G and current independent set I , a random vertex is added from the subset of vertices of G that are not in the independent set or adjacent to a member of the independent set. The addition is iterated until no further vertices can be added and the independent set is maximal. This heuristic, which is zero-local as it is not aware of any edge structure when selecting vertices, has a performance guarantee of $2/(\nu + 1)$ on ν -regular graphs and is classically optimal up to constant factors [70].

The greedy heuristic is used in a two-step process to generate a valid maximal independent set. First, independent set violations are greedily removed as follows. Given some graph G and candidate solution z , the greedy add heuristic is implemented starting with the empty set on the subgraph $G[z]$ induced by z , generating some independent set z' . If z is an independent set, then $G[z]$ has no edges, and thus, the independent set z' is the same as z . If z is not an independent set, then $z' \subset z$ is an independent set through the greedy add


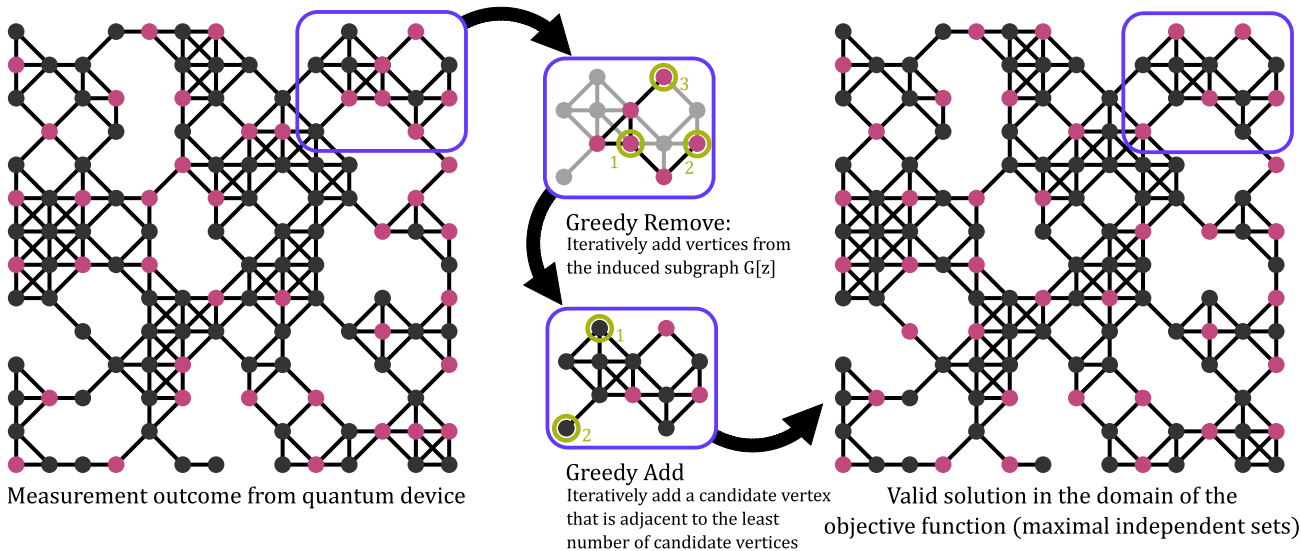


FIGURE 4. Structure and performance of greedy postprocessing solving the MIS on unit disk graphs, as implemented on QuEra’s cloud-accessible quantum computer “Aquila.” A measurement outcome from the quantum device (top left) is used as a “warm start” for a heuristic greedy classical postprocessing. First, independent set violations are removed by doing a zero-local greedy add on the induced subgraph $G[z]$, and then, the solution is made maximal by randomly adding vertices until no more can be added. The measurement outcomes are generated by a piecewise linear variational adiabatic state preparation (bottom left). Results of the protocol for the example problem instance are shown on the bottom right; for these optimal parameters, the hybrid algorithm is outperformed by the classical algorithm in average solution and maximum solution, indicating no comparative advantage for the particular problem instance. Given a more optimized preparation protocol, it is possible to improve performance and possibly get a quadratic speedup over classical methods [64].

procedure. Likewise, the candidate solution z' may not be a maximal independent set and thus not a valid solution. By greedy adding warm started by z' , the generated solution χ will be both maximal and independent and thus be in the solution space. These steps are shown in Fig. 4 (top).

This algorithm has a natural no-quantum limit, by replacing the measurement outcome z with the all-zeros or all-ones bit string. The greedy algorithm will act equivalently on both to add or remove vertices until maximal, and the performance serves as a comparative benchmark against the quantum-enhanced version.

Results for this algorithm as implemented on QuEra’s cloud-accessible quantum computer “Aquila” [36] are shown in Fig. 4 (bottom right), which plots the probability distribution of maximal independent set sizes for the particular graph instance shown; we find that the hybrid (classical) algorithm generates independent sets of average size 44.282

(45.178) with the hybrid algorithm having a performance ratio of 0.9801, indicating that the hybrid algorithm is outperformed on average by the classical-only limit. Similarly, the classical-only algorithm finds a better single-shot solution than quantum with an 81.2% probability and finds an MIS of size 47 with a 1.80% probability, which the hybrid algorithm never finds. These results suggest that, in this case, the quantum elements of the algorithm in fact *hinder*, not *help* the performance.

This is a crucial fact: including quantum resources in an NNHA does not result in a self-evident improvement in performance over classical-only algorithms. Indeed, any algorithm must be carefully designed to avoid reducing performance. In this example, the $1/R^6$ Rydberg interaction causes a “repulsion” of Rydberg states outside of the Rydberg radius, which favors extra space between adjacent excitations. In an extreme case, it may cause excitations to typically be

distance 3 apart instead of distance 2, generating a small but maximal independent set. It is interesting to consider how these patterns can be leveraged; for example, generating such distributions might be used as a resource for finding minimum dominating sets. Increasing the complexity and length of the adiabatic protocol may lead to improved performance and avoid these pitfalls, which is the case for [64]. Such an algorithm and hardware-specific tuning is out of the scope of this work.

B. CLUSTER UPDATE: TYPE-3 SOLUTION

A natural example of type-3 NNHA is cluster update using parallel tempering [71] and simulated annealing [57], [72], where subsystems are exchanged between two samples from Boltzmann distributions of solutions at different temperatures. By coupling the low-temperature ensemble to the high-temperature one, a Markov update Monte Carlo can be kicked out of local maxima to more efficiently explore the low-energy state space of the distribution. In this implementation, instead of two classically generated distributions, the high-temperature distribution is replaced by samples drawn from the quantum device. The distribution may not be thermal with respect to the objective function, though recent work on QAOA finds that Boltzmann distributions are surprisingly common [73]. If the ansatz adiabatically prepares almost-ground states of a Hamiltonian that almost encodes the objective function, as is the case for MIS and neutral atoms, the direct samples from the reservoir may nonetheless be close to the maximum, and thus, the ensemble may generate relatively low-temperature states [74].

Instead of a direct parallel tempering implementation, which occasionally directly exchanges samples from different reservoirs, a more advanced version instead exchanges subsystems of samples from each reservoir, e.g., a “cluster update.” In this case, only a small portion of the current Monte Carlo solution step in the low-temperature reservoir is replaced with the high-temperature distribution. Cluster update methods with scale-free sizes have been found to converge to better results faster than single-vertex updates [75]. Note that because many individual quantum samples are aggregated into a single solution, the resulting output may be better than any individual sample.

The cluster update NNHA works as follows, as shown in Fig. 5. Given some solution χ (initialized as the empty independent set):

- 1) Pick some subsystem to update. There are several strategies on how to select a cluster. For instance, one could select a pair of adjacent vertices for which one is in the independent set, in which case the method reduces to a diffusive search over Hamming distance 2-adjacent maximal independent sets [72]. Alternatively, one could select vertices within a fixed radius of a central vertex, in which case the method diffusively explores across larger Hamming distance moves. In this example, the insights of self-organized criticality

for optimization are used by selecting clusters based on a sandpile model [75]. A sandpile model works by iteratively adding “sand grains” to random vertices (slow process) until a vertex reaches a collapse threshold given by its adjacency; then, the vertex sheds a grain to each adjacent vertex (fast process), which may trigger further collapses. This fast process is iterated until no further vertices collapse; the cluster is generated by all vertices that underwent collapse. A non-scale-free dissipation term is added by removing each grain from a collapsing vertex with a 5% probability. This process is known to generate a scale-free size distribution of clusters. It has been found on similar objectives that such a cluster update outperforms fixed-cluster optimization techniques [75].

- 2) Sample an approximate solution χ from the high-temperature reservoir by calling the greedy algorithm of Section V-A as a subroutine. The no-quantum limit of that subroutine serves as the no-quantum limit of this algorithm.
- 3) Merge the solution χ with the reservoir solution χ' , as $z = \chi \cup \chi'$. In the vertices of the cluster selected by step 1, update the solution χ to the solution of χ' ; outside, keep the solution χ .
- 4) Enforce that the new candidate solution z is a valid maximal independent set with a type-1 postprocessing, by a mapping $F(z) \rightarrow \chi$ using the greedy heuristic of Section V-A to guarantee the new solution is a valid maximal independent set. Even if both χ and χ' are valid maximal independent sets, the merged candidate may have independent set violations along the boundary.
- 5) Accept or reject the update based on the change in objective value using the Metropolis–Hastings update rule [76] $P_{\text{accept}} = \text{MIN}[1, \exp(-\beta\Delta)]$, where β is the inverse temperature and $\Delta = C(\chi') - C(\chi)$ is the change in objective value.
- 6) Finally, go to 1) for some fixed number of steps or break after the solution χ becomes stationary given $\beta \rightarrow \infty$. Update steps equal to the number of vertices in the graph is one epoch.

Results for this algorithm as implemented on QuEra’s cloud-accessible quantum computer “Aquila” [36] are shown in Fig. 5, which is implemented for a particular graph instance. Parameters are variationally optimized to an objective function that averages the objective over three epochs of annealing steps. A bounded stochastic optimizer finds the best parameters to be $t_{\text{max}} = 3.80\mu\text{s}$, $\Delta_{\text{min}} = -13.47\text{rad}/\mu\text{s}$, and $\Delta_{\text{max}} = 41.95\text{rad}/\mu\text{s}$. We find that the hybrid (classical) algorithm generates independent sets of size 45.72 (45.48) with the hybrid algorithm having a performance ratio of 1.005, indicating that the NNHA slightly outperforms the classical-only limit. Similarly, the NNHA samples an equal-or-better solution with an 83% probability. For more details, see an example code notebook here [37].

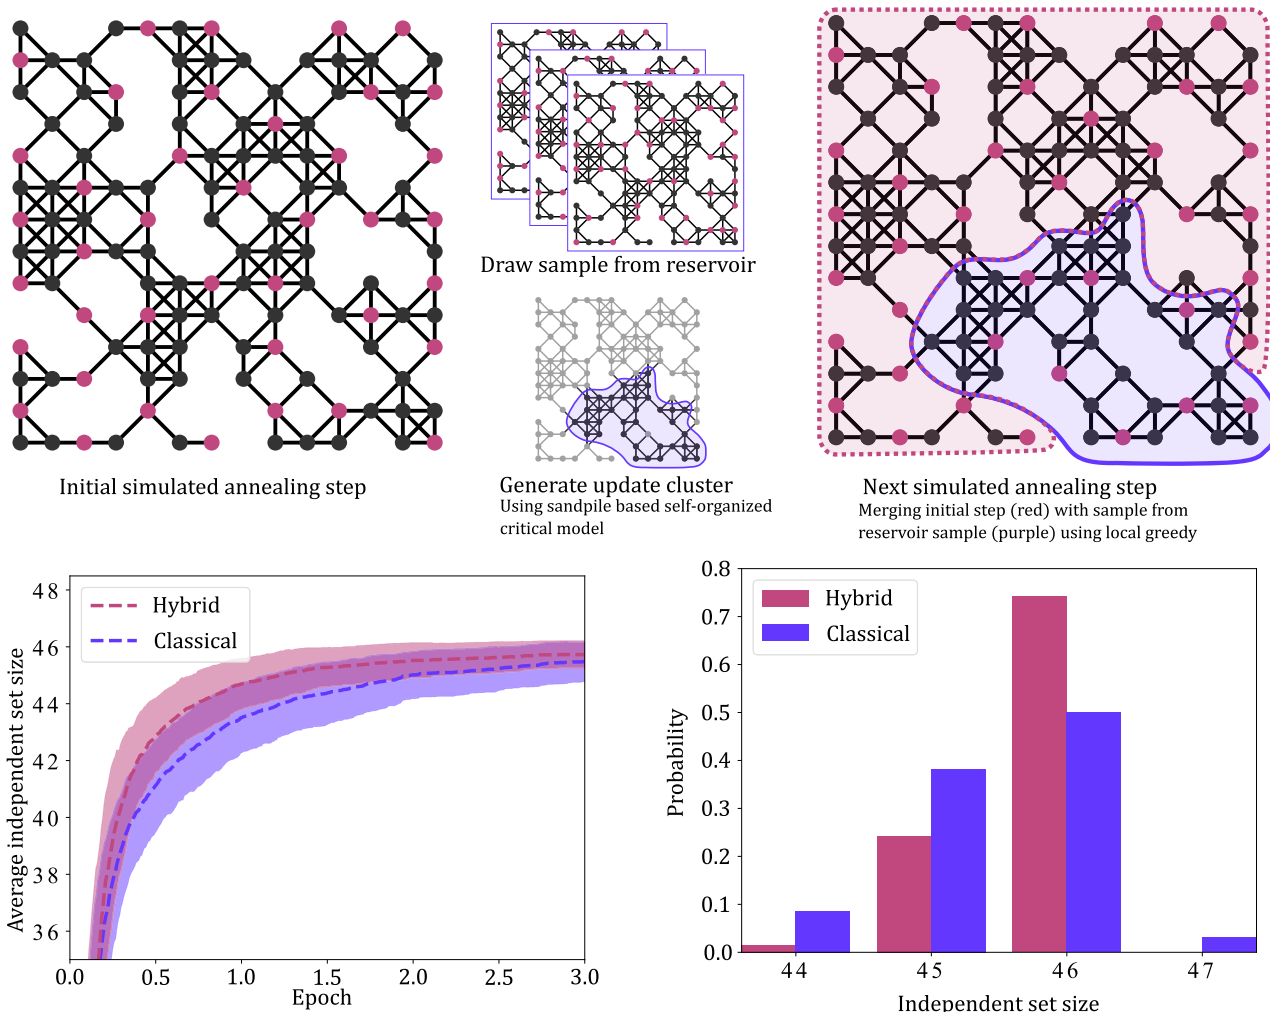


FIGURE 5. Structure and performance of cluster update simulated annealing implementation using type-3 NNHA, as implemented on QuEra’s cloud-accessible quantum computer “Aquila.” Top sketches the structure of the algorithm, which iteratively updates and improves a candidate solution (top left). First, a sample is drawn based on a measurement outcome from a neutral-atom quantum annealing protocol, and a cluster is generated based on a self-organized critical sandpile model (top middle). Then, the candidate solution is merged with the new sample (top right) by replacing the solution within the cluster with the reservoir measurement. The process is iterated many times until converged (bottom left), where each epoch is N update steps. Given variationally optimized parameters, it is clear that the NNHA converges to optimal solutions faster than the classical-only limit (purple). This is an indicator of comparative advantage, as the algorithm performs better when quantum resources are included. The bottom right plots the distribution of solutions after three epochs of zero-temperature updates. While the NNHA has a better average performance, the classical variant can occasionally outperform and find the actual MIS of size 47.

For this particular instance, the hybrid annealing algorithm converges to finding larger independent sets faster than the equivalent classical algorithm, which simply greedily adds vertices within the reservoir. This may be as expected, as the independent set samples “know” more about the global structure of the graph than the zero-local greedy algorithm. We observe that performance is highly dependent on the variational parameters. For example, if the final detuning is too small, the effective blockade radius is too large, and the Rydberg states will typically be further apart. This causes measurement samples to accidentally block good maximal solutions, and the NNHA converges to solutions *slower* than the classical-only limit. Such behavior may be useful if, for example, the solution space is over minimum dominating sets, but is otherwise detrimental.

Furthermore, even though the quantum variant finds a larger expected independent set size on average, the classical-only limit may occasionally find a larger independent set size of 47. While this result is seen for a specific problem instance, this still suggests that classical algorithms may outperform NNHA depending on the particulars of the objective. Improvements to this prototype implementation may also see an enhanced performance, though such a study is beyond the scope of this work.

Finally, it should be emphasized that this particular implementation is a template for other type-3 postprocessing implementations. For example, other kinds of cluster update strategies may be used or optimized over different objectives such as dominating sets. In particular, cluster updates may be used to solve MIS problem instances much larger than

can fit on quantum hardware. By implementing quantum annealing dynamics on each (subextensive) cluster plus some fixed radius around a cluster, a large graph can be broken into many small subgraphs and stitched together using annealing. In this way, arbitrary large graph problems may be solved using a reasonably sized neutral-atom quantum computer, at the cost that correlations in each sample will be small in extent. Such an implementation is beyond the scope of this work. Alternatively, the samples could be used to generate clusters instead of the updates themselves in a multilevel algorithm [27].

VI. CONCLUSION AND FUTURE DIRECTIONS

There are three key challenges to contemporary quantum optimization. First, the performance of heuristic quantum algorithms is typically quite poor in comparison to even simple contemporary classical optimization algorithms. Second, the constraints and objectives may not be native to the underlying hardware, requiring expensive overhead and encoding. Finally, the solution domain of relevant real-world problems is not typically bit strings, typically requiring some encoding overhead.

These challenges are highlighted when attempting to solve real-world optimization problems with quantum or hybrid methods. Typical problems are large, complicated, multiobjective, constrained, and do not naturally or natively map onto quantum hardware.

This work has presented a framework of NNHA as a method to overcome these challenges, by incorporating quantum resources to assist classical optimization so that the quantum execution does not need to directly encode solutions. In this way, the quantum computer acts as a coprocessor for classical computation, offloading specific subparts of the optimization while the classical components refine and include more complex real-world components of the problem instance. Such an architecture is particularly natural for high-performance computing applications, where classical computing resources may be much larger than any quantum resource in the near future. Notably, this version of “hybrid” computing goes beyond the paradigm of only using classical resources for variational optimization by also directly assisting with computing solutions.

The major contribution of this work is to propose the idea of measurement outcomes from the quantum device $\{z\}$ not needing to directly map to solutions. Instead, NNHA uses bit strings indirectly as a resource for a classical step to compute a solution. In this way, classical resources contribute actively to both parameter finding and directly constructing candidate solutions to the target problem. This contrasts with contemporary quantum optimization, where bit strings directly encode solutions and classical postprocessing only decodes bit strings into solutions [4]. We enumerate with examples three kinds of NNHA: type 1 uses individual bit strings as a “warm start” to postprocess and refine solutions; type 2 uses expectation values and correlation functions as a “quantum hint” to glean hard-to-find structure of problem instances;

and type 3 directly uses the distribution of bit strings as a “reservoir” for computation.

While this work presents prototype solutions to demonstrate each type of hybrid optimization, they are by no means complete. Indeed, these implementations should be seen as templates and inspiration for bespoke solutions to real-world optimization applications. Instead, these methods should serve the key objective of solving optimization problems beyond simple problem classes and including the full complexity of real-world objectives. By maximally leveraging classical resources and designing NNHA that closely synergizes both quantum and classical subroutines, we come closer to a goal of the field: quantum practicality, where quantum resources improve outcomes and solve problems in a practical, demonstrable, and real-world setting.

ACKNOWLEDGMENT

The authors would like to thank Alexander Keesling, Maddie Cain, Nate Gemelke, and Phillip Weinberg for helpful discussions and Danylo Lykov who had early contributions to this work.

REFERENCES

- [1] A. Abbas et al., “Quantum optimization: Potential, challenges, and the path forward,” 2023, *arXiv:2312.02279*, doi: [10.48550/arXiv.2312.02279](https://doi.org/10.48550/arXiv.2312.02279).
- [2] M. Cerezo et al., “Variational quantum algorithms,” *Nature Rev. Phys.*, vol. 3, no. 9, pp. 625–644, 2021, doi: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9).
- [3] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” 2014, *arXiv:1411.4028*, doi: [10.48550/arXiv.1411.4028](https://doi.org/10.48550/arXiv.1411.4028).
- [4] K. Blekos et al., “A review on quantum approximate optimization algorithm and its variants,” *Phys. Rep.*, vol. 1068, pp. 1–66, 2024, doi: [10.1016/j.physrep.2024.03.002](https://doi.org/10.1016/j.physrep.2024.03.002).
- [5] E. Farhi, J. Goldstone, S. Gutmann, and H. Neven, “Quantum algorithms for fixed qubit architectures,” 2017, *arXiv:1703.06199*, doi: [10.48550/arXiv.1703.06199](https://doi.org/10.48550/arXiv.1703.06199).
- [6] T. Albash and D. A. Lidar, “Adiabatic quantum computation,” *Rev. Modern Phys.*, vol. 90, no. 1, 2018, Art. no. 015002, doi: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002).
- [7] K. Marwaha, “Local classical MAX-CUT algorithm outperforms $p = 2$ QAOA on high-girth regular graphs,” *Quantum*, vol. 5, 2021, Art. no. 437, doi: [10.22331/q-2021-04-20-437](https://doi.org/10.22331/q-2021-04-20-437).
- [8] D. Lykov, J. Wurtz, C. Poole, M. Saffman, T. Noel, and Y. Alexeev, “Sampling frequency thresholds for the quantum advantage of the quantum approximate optimization algorithm,” *npj Quantum Inf.*, vol. 9, no. 1, 2023, Art. no. 73, doi: [10.1038/s41534-023-00718-4](https://doi.org/10.1038/s41534-023-00718-4).
- [9] J. Weidenfeller et al., “Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware,” *Quantum*, vol. 6, Dec. 2022, Art. no. 870, doi: [10.22331/q-2022-12-07-870](https://doi.org/10.22331/q-2022-12-07-870).
- [10] S. Hadfield, “On the representation of Boolean and real functions as Hamiltonians for quantum computing,” *ACM Trans. Quantum Comput.*, vol. 2, no. 4, pp. 1–21, 2021, doi: [10.1145/3478519](https://doi.org/10.1145/3478519).
- [11] 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*, vol. 12, no. 2, 2019, Art. no. 34, doi: [10.3390/a12020034](https://doi.org/10.3390/a12020034).
- [12] R. LaRose, E. Rieffel, and D. Venturelli, “Mixer-phaser ansätze for quantum optimization with hard constraints,” *Quantum Mach. Intell.*, vol. 4, no. 2, Jun. 2022, Art. no. 17, doi: [10.1007/s42484-022-00069-x](https://doi.org/10.1007/s42484-022-00069-x).
- [13] A. Lucas, “Ising formulations of many NP problems,” *Front. Phys.*, vol. 2, 2014, Art. no. 5, doi: [10.3389/fphy.2014.00005](https://doi.org/10.3389/fphy.2014.00005).
- [14] S. Okada, M. Ohzeki, and S. Taguchi, “Efficient partition of integer optimization problems with one-hot encoding,” *Sci. Rep.*, vol. 9, no. 1, 2019, Art. no. 13036, doi: [10.1038/s41598-019-49539-6](https://doi.org/10.1038/s41598-019-49539-6).

- [15] U. Azad, B. K. Behera, E. A. Ahmed, P. K. Panigrahi, and A. Farouk, "Solving vehicle routing problem using quantum approximate optimization algorithm," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 7, pp. 7564–7573, Jul. 2023, doi: [10.1109/TITS.2022.3172241](https://doi.org/10.1109/TITS.2022.3172241).
- [16] R. Biswas et al., "A NASA perspective on quantum computing: Opportunities and challenges," *Parallel Comput.*, vol. 64, pp. 81–98, 2017, doi: [10.1016/j.parco.2016.11.002](https://doi.org/10.1016/j.parco.2016.11.002).
- [17] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchitsky, and R. Melko, "Quantum Boltzmann machine," *Phys. Rev. X*, vol. 8, no. 2, 2018, Art. no. 021050, doi: [10.1103/PhysRevX.8.021050](https://doi.org/10.1103/PhysRevX.8.021050).
- [18] "Cuda quantum—Nvidia cuda quantum documentation," Accessed: Apr. 2022. [Online]. Available: <https://nvidia.github.io/cuda-quantum/latest/index.html>
- [19] "Three truths and the advent of hybrid quantum computing," Accessed: Apr. 2022. [Online]. Available: <https://medium.com/d-wave/three-truths-and-the-advent-of-hybrid-quantum-computing-1941ba46ff8c>
- [20] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New J. Phys.*, vol. 18, no. 2, 2016, Art. no. 023023, doi: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023).
- [21] R. Ayanzadeh, J. E. Dorband, M. Halem, and T. Finin, "Quantum-assisted greedy algorithms," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2022, pp. 4911–4914, doi: [10.1109/IGARSS46834.2022.9884795](https://doi.org/10.1109/IGARSS46834.2022.9884795).
- [22] M. Dupont et al., "Quantum-enhanced greedy combinatorial optimization solver," *Sci. Adv.*, vol. 9, no. 45, 2023, Art. no. eadi0487, doi: [10.1126/sciadv.adi0487](https://doi.org/10.1126/sciadv.adi0487).
- [23] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, "Obstacles to variational quantum optimization from symmetry protection," *Phys. Rev. Lett.*, vol. 125, no. 26, Dec. 2020, Art. no. 260505, doi: [10.1103/PhysRevLett.125.260505](https://doi.org/10.1103/PhysRevLett.125.260505).
- [24] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, "Hybrid quantum-classical algorithms for approximate graph coloring," *Quantum*, vol. 6, Mar. 2022, Art. no. 678, doi: [10.22331/q-2022-03-30-678](https://doi.org/10.22331/q-2022-03-30-678).
- [25] J. R. Finzgar, A. Kerschbaumer, M. J. A. Schuetz, C. B. Mendl, and H. G. Katzgraber, "Quantum-informed recursive optimization algorithms," *PRX Quantum*, vol. 5, no. 2, 2024, Art. no. 020327, doi: [10.1103/PRXQuantum.5.020327](https://doi.org/10.1103/PRXQuantum.5.020327).
- [26] H. Ushijima-Mwesigwa, R. Shayduln, C. F. A. Negre, S. M. Mniszewski, Y. Alexeev, and I. Safro, "Multilevel combinatorial optimization across quantum architectures," *ACM Trans. Quantum Comput.*, vol. 2, no. 1, pp. 1–29, 2021, doi: [10.1145/3425607](https://doi.org/10.1145/3425607).
- [27] A. Angone, X. Liu, R. Shayduln, and I. Safro, "Hybrid quantum-classical multilevel approach for maximum cuts on graphs," in *Proc. IEEE High Perform. Extreme Comput. Conf.*, 2023, pp. 1–7, doi: [10.1109/HPEC58863.2023](https://doi.org/10.1109/HPEC58863.2023).
- [28] S. Chakrabarti, P. Minssen, R. Yalovetzky, and M. Pistoia, "Universal quantum speedup for branch-and-bound, branch-and-cut, and tree-search algorithms," 2022, *arXiv:2210.03210*, doi: [10.48550/arXiv.2210.03210](https://doi.org/10.48550/arXiv.2210.03210).
- [29] M. J. Schuetz et al., "Optimization of robot-trajectory planning with nature-inspired and hybrid quantum algorithms," *Phys. Rev. Appl.*, vol. 18, no. 5, 2022, Art. no. 054045, doi: [10.1103/PhysRevApplied.18.054045](https://doi.org/10.1103/PhysRevApplied.18.054045).
- [30] J. Wurtz and P. J. Love, "Classically optimal variational quantum algorithms," *IEEE Trans. Quantum Eng.*, vol. 2, 2021, Art. no. 3104107, doi: [10.1109/TQE.2021.3122568](https://doi.org/10.1109/TQE.2021.3122568).
- [31] M. Ponce et al., "Graph decomposition techniques for solving combinatorial optimization problems with variational quantum algorithms," 2023, *arXiv:2306.00494*, doi: [10.48550/arXiv.2306.00494](https://doi.org/10.48550/arXiv.2306.00494).
- [32] F. Wagner, J. Nüßlein, and F. Liers, "Enhancing quantum algorithms for quadratic unconstrained binary optimization via integer programming," 2023, *arXiv:2302.05493*, doi: [10.48550/arXiv.2302.05493](https://doi.org/10.48550/arXiv.2302.05493).
- [33] M. Kornjača, "Large-scale quantum reservoir learning with an analog quantum computer," 2024, *arXiv:2407.02553*, doi: [10.48550/arXiv.2407.02553](https://doi.org/10.48550/arXiv.2407.02553).
- [34] J. Wang, S. Lu, S.-H. Wang, and Y.-D. Zhang, "A review on extreme learning machine," *Multimedia Tools Appl.*, vol. 81, no. 29, pp. 41611–41660, 2021, doi: [10.1007/s11042-021-11007-7](https://doi.org/10.1007/s11042-021-11007-7).
- [35] B. A. Wilson et al., "Non-native quantum generative optimization with adversarial autoencoders," 2024, *arXiv:2407.13830*, doi: [10.48550/arXiv.2407.13830](https://doi.org/10.48550/arXiv.2407.13830).
- [36] J. Wurtz et al., "Aquila: Quera's 256-qubit neutral-atom quantum computer," 2023, *arXiv:2306.11727*, doi: [10.48550/arXiv.2306.11727](https://doi.org/10.48550/arXiv.2306.11727).
- [37] "Bloqade-python-MIS," 2024. [Online]. Available: https://github.com/QuEraComputing/bloqade-python-mis/tree/main/postprocessing_implementations
- [38] G. Vidal, "Efficient simulation of one-dimensional quantum many-body systems," *Phys. Rev. Lett.*, vol. 93, 2004, Art. no. 040502, doi: [10.1103/PhysRevLett.93.040502](https://doi.org/10.1103/PhysRevLett.93.040502).
- [39] J. Zhou et al., "Graph neural networks: A review of methods and applications," 2021, *arXiv:1812.08434*, doi: [10.48550/arXiv.1812.08434](https://doi.org/10.48550/arXiv.1812.08434).
- [40] R. Salakhutdinov and G. Hinton, "Deep Boltzmann machines," in *Proc. 12th Int. Conf. Artif. Intell. Statist.*, 2009, pp. 448–455. [Online]. Available: <https://proceedings.mlr.press/v5/salakhutdinov9a>
- [41] S. Czischek, M. S. Moss, M. Radzihevsky, E. Merali, and R. G. Melko, "Data-enhanced variational Monte Carlo simulations for Rydberg atom arrays," *Phys. Rev. B*, vol. 105, 2022, Art. no. 205108, doi: [10.1103/PhysRevB.105.205108](https://doi.org/10.1103/PhysRevB.105.205108).
- [42] A. Hashim et al., "Optimized swap networks with equivalent circuit averaging for QAOA," *Phys. Rev. Res.*, vol. 4, 2022, Art. no. 033028, doi: [10.1103/PhysRevResearch.4.033028](https://doi.org/10.1103/PhysRevResearch.4.033028).
- [43] N. Sachdeva et al., "Quantum optimization using a 127-qubit gate-model IBM quantum computer can outperform quantum annealers for nontrivial binary optimization problems," 2024, *arXiv:2406.01743*, doi: [10.48550/arXiv.2406.01743](https://doi.org/10.48550/arXiv.2406.01743).
- [44] J. Wurtz and P. Love, "MaxCut quantum approximate optimization algorithm performance guarantees for $p > 1$," *Phys. Rev. A*, vol. 103, 2021, Art. no. 042612, doi: [10.1103/PhysRevA.103.042612](https://doi.org/10.1103/PhysRevA.103.042612).
- [45] G. Nannicini, "Performance of hybrid quantum-classical variational heuristics for combinatorial optimization," *Phys. Rev. E*, vol. 99, no. 1, 2019, Art. no. 013304, doi: [10.1103/PhysRevE.99.013304](https://doi.org/10.1103/PhysRevE.99.013304).
- [46] M. J. D. Powell, "A direct search optimization method that models the objective and constraint functions by linear interpolation," in *Advances in Optimization and Numerical Analysis. Mathematics and Its Applications*. Dordrecht, The Netherlands: Springer, 1994, pp. 51–67, doi: [10.1007/978-94-015-8330-5_4](https://doi.org/10.1007/978-94-015-8330-5_4).
- [47] X. Bonet-Monroig et al., "Performance comparison of optimization methods on variational quantum algorithms," *Phys. Rev. A*, vol. 107, 2023, Art. no. 032407, doi: [10.1103/PhysRevA.107.032407](https://doi.org/10.1103/PhysRevA.107.032407).
- [48] S. H. Sack and D. J. Egger, "Large-scale quantum approximate optimization on nonplanar graphs with machine learning noise mitigation," *Phys. Rev. Res.*, vol. 6, no. 1, 2024, Art. no. 013223, doi: [10.1103/PhysRevResearch.6.013223](https://doi.org/10.1103/PhysRevResearch.6.013223).
- [49] A. Peruzzo et al., "A variational eigenvalue solver on a photonic quantum processor," *Nature Commun.*, vol. 5, 2014, Art. no. 4213, doi: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [50] N. Herrmann et al., "Quantum utility—definition and assessment of a practical quantum advantage," in *Proc. IEEE Int. Conf. Quantum Softw.*, 2023, pp. 162–174, doi: [10.1109/QSW59989.2023.00028](https://doi.org/10.1109/QSW59989.2023.00028).
- [51] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *J. ACM*, vol. 42, no. 6, pp. 1115–1145, 1995, doi: [10.1145/227683.227684](https://doi.org/10.1145/227683.227684).
- [52] G. E. Crooks, "Performance of the quantum approximate optimization algorithm on the maximum cut problem," 2018, *arXiv:1811.08419*, doi: [10.48550/arXiv.1811.08419](https://doi.org/10.48550/arXiv.1811.08419).
- [53] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, "Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices," *Phys. Rev. X*, vol. 10, 2020, Art. no. 021067, doi: [10.1103/PhysRevX.10.021067](https://doi.org/10.1103/PhysRevX.10.021067).
- [54] J. Wurtz and P. J. Love, "Counterdiabaticity and the quantum approximate optimization algorithm," *Quantum*, vol. 6, Jan. 2022, Art. no. 635, doi: [10.22331/q-2022-01-27-635](https://doi.org/10.22331/q-2022-01-27-635).
- [55] J. Wurtz and D. Lykov, "Fixed-angle conjectures for the quantum approximate optimization algorithm on regular MaxCut graphs," *Phys. Rev. A*, vol. 104, 2021, Art. no. 052419, doi: [10.1103/PhysRevA.104.052419](https://doi.org/10.1103/PhysRevA.104.052419).
- [56] J. Hirvonen, J. Rybicki, S. Schmid, and J. Suomela, "Large cuts with local algorithms on triangle-free graphs," 2014, *arXiv:1402.2543*, doi: [10.48550/arXiv.1402.2543](https://doi.org/10.48550/arXiv.1402.2543).
- [57] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983, doi: [10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671).
- [58] E. Farhi and A. W. Harrow, "Quantum supremacy through the quantum approximate optimization algorithm," 2019, *arXiv:1602.07674*, doi: [10.48550/arXiv.1602.07674](https://doi.org/10.48550/arXiv.1602.07674).

- [59] A. Coja-Oghlan, C. Moore, and V. Sanwalani, "MAX k-CUT and approximating the chromatic number of random graphs," in *Proc. 30th Int. Conf. Autom., Lang. Program.*, 2003, pp. 200–211, doi: [10.1002/rsa.20096](https://doi.org/10.1002/rsa.20096).
- [60] A. Buluc, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz, "Recent advances in graph partitioning," 2015, *arXiv:1311.3144*, doi: [10.48550/arXiv.1311.3144](https://doi.org/10.48550/arXiv.1311.3144).
- [61] M. D. Dragos, M. Cvetkovicc, and H. Sachs, *Spectra of Graphs: Theory and Applications*. Munich, Germany: Barth, 1995.
- [62] A. Ng, M. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Proc. 14th Int. Conf. Neural Inf. Process. Syst.: Natural Synthetic*, 2001, pp. 849–856. [Online]. Available: <https://proceedings.neurips.cc/paper/2001/hash/801272ee79cfd7fa5960571fee36b9b-Abstract.html>
- [63] M. Dupont and B. Sundar, "Extending relax-and-round combinatorial optimization solvers with quantum correlations," *Phys. Rev. A*, vol. 109, Jan. 2024, Art. no. 012429, doi: [10.1103/PhysRevA.109.012429](https://doi.org/10.1103/PhysRevA.109.012429).
- [64] S. Ebadi et al., "Quantum optimization of maximum independent set using Rydberg atom arrays," *Science*, vol. 376, no. 6598, pp. 1209–1215, 2022, doi: [10.1126/science.abo6587](https://doi.org/10.1126/science.abo6587).
- [65] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Commun. Math. Phys.*, vol. 28, no. 3, pp. 251–257, Sep. 1972, doi: [10.1007/BF01645779](https://doi.org/10.1007/BF01645779).
- [66] J. Lux, J. Müller, A. Mitra, and A. Rosch, "Hydrodynamic long-time tails after a quantum quench," *Phys. Rev. A*, vol. 89, 2014, Art. no. 053608, doi: [10.1103/PhysRevA.89.053608](https://doi.org/10.1103/PhysRevA.89.053608).
- [67] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, "Computational complexity of the Rydberg blockade in two dimensions," 2018, *arXiv:1809.04954*, doi: [10.48550/arXiv.1809.04954](https://doi.org/10.48550/arXiv.1809.04954).
- [68] D. Gamarnik, "The overlap gap property: A topological barrier to optimizing over random structures," *Proc. Nat. Acad. Sci.*, vol. 118, no. 41, Art. no. e2108492118, 2021, doi: [10.1073/pnas.2108492118](https://doi.org/10.1073/pnas.2108492118).
- [69] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, "Quantum optimization for maximum independent set using Rydberg atom arrays," 2018, *arXiv:1808.10816*, doi: [10.48550/arXiv.1808.10816](https://doi.org/10.48550/arXiv.1808.10816).
- [70] M. M. Halldórsson and J. Radhakrishnan, "Greed is good: Approximating independent sets in sparse and bounded-degree graphs," *Algorithmica*, vol. 18, no. 1, pp. 145–163, May 1997, doi: [10.1007/BF02523693](https://doi.org/10.1007/BF02523693).
- [71] C. J. Geyer, "Markov chain Monte Carlo maximum likelihood," 1991. [Online]. Available: <https://www.stat.umn.edu/geyer/f05/8931/c.pdf>
- [72] M. Cain, S. Chattopadhyay, J.-G. Liu, R. Samajdar, H. Pichler, and M. D. Lukin, "Quantum speedup for combinatorial optimization with flat energy landscapes," 2023, *arXiv:2306.13123*, doi: [10.48550/arXiv.2306.13123](https://doi.org/10.48550/arXiv.2306.13123).
- [73] P. C. Lotshaw et al., "Approximate Boltzmann distributions in quantum approximate optimization," *Phys. Rev. A*, vol. 108, 2023, Art. no. 042411, doi: [10.1103/PhysRevA.108.042411](https://doi.org/10.1103/PhysRevA.108.042411).
- [74] D. S. Wild, D. Sels, H. Pichler, C. Zanoci, and M. D. Lukin, "Quantum sampling algorithms for near-term devices," *Phys. Rev. Lett.*, vol. 127, 2021, Art. no. 100504, doi: [10.1103/PhysRevLett.127.100504](https://doi.org/10.1103/PhysRevLett.127.100504).
- [75] H. Hoffmann and D. W. Payton, "Optimization by self-organized criticality," *Sci. Rep.*, vol. 8, no. 1, 2018, Art. no. 2358, doi: [10.1038/s41598-018-20275-7](https://doi.org/10.1038/s41598-018-20275-7).
- [76] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970, doi: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97).