

# Network Community Detection on Small Quantum Computers

Ruslan Shaydulin,\* Hayato Ushijima-Mwesigwa, Ilya Safro, Susan Mniszewski, and Yuri Alexeev

In recent years, a number of quantum computing devices with small numbers of qubits have become available. A hybrid quantum local search (QLS) approach that combines a classical machine and a small quantum device to solve problems of practical size is presented. The proposed approach is applied to the network community detection problem. QLS is hardware-agnostic and easily extendable to new quantum computing devices as they become available. It is demonstrated to solve the 2-community detection problem on graphs of sizes of up to 410 vertices using the 16-qubit IBM quantum computer and D-Wave 2000Q, and compare their performance with the optimal solutions. The results herein demonstrate that QLS performs similarly in terms of quality of the solution and the number of iterations to convergence on both types of quantum computers and it is capable of achieving results comparable to state-of-the-art solvers in terms of quality of the solution including reaching the optimal solutions.

## 1. Introduction

The recent years saw rapid progress in the development of quantum computing (QC) devices. Multiple paradigms have been proposed and implemented in hardware introducing a variety of limitations that must be addressed prior to the wide application of QC. In particular, noisy intermediate scale quantum (NISQ) devices are widely expected to be limited to a few hundred, or perhaps a few thousand qubits,<sup>[1]</sup> severely restricting the size of the problems that can be tackled directly. As the potential of these NISQ-era quantum devices is becoming evident,<sup>[2]</sup> there is an increasing interest in developing algorithms that

leverage the small quantum devices that are becoming available. This requires the use of hybrid quantum-classical approaches where a problem is solved across a classical machine and a QC device.

The number of qubits in NISQ-era devices available at the time of writing is not nearly enough to demonstrate quantum advantage, which makes it especially hard to demonstrate the usefulness of quantum computers to solve real problems. For example, the possibility of quantum speedup using the hybrid quantum approximate optimization algorithm (QAOA) for a network problem similar to the one discussed in this paper (maximum cut or max-cut) is a subject of active discussion. On the one hand, there are theoretical results demonstrating that QAOA for max-cut problem improves upon best know classical

approximation algorithms for certain graphs.<sup>[3,4]</sup> At the same time, there are indications that achieving speedup using QAOA might require at least several hundred qubits.<sup>[5]</sup> Research and development of quantum algorithms is necessary as the number and quality of qubits is improving. These quantum algorithms can also be used to improve classical algorithms.<sup>[6]</sup> The need for development of new quantum algorithms was highlighted in the recent National Academy of Science report.<sup>[7]</sup> One of the important directions to make quantum computing feasible in the near future is to use various problem decomposition approaches to solve a large problems as a set of subproblems. This can be accomplished at various levels such as problem formulation or at the algorithmic level as demonstrated in this paper.

The decomposition approach might be the key method to achieve a quantum speedup on even modest-size NISQ devices in near-term future. To support this claim, there is an important and encouraging work,<sup>[8]</sup> where it was shown that large combinatorial optimization problems can be effectively decomposed into subproblems on quantum annealing hardware, while still obtaining high quality of the overall solution. It was demonstrated for solving embedding problems on D-Wave quantum computers, but we believe that the same technique can be used to improve dramatically the speed and performance of QAOA algorithms on universal quantum computers.

In this work, we introduce the quantum local search (QLS) algorithm for the network community detection problem that is based on the local search method.<sup>[9]</sup> Many different versions of the local search have been applied to numerous computationally

R. Shaydulin, Dr. H. Ushijima-Mwesigwa, Prof. I. Safro  
School of Computing  
Clemson University  
Clemson, SC 29634, USA  
E-mail: rshaydu@g.clemson.edu

Dr. S. Mniszewski  
Computer, Computational, & Statistical Sciences Division  
Los Alamos National Laboratory  
Los Alamos, NM 87545, USA

Dr. Y. Alexeev  
Computational Science Division  
Argonne National Laboratory  
Lemont, IL 60439, USA

 The ORCID identification number(s) for the author(s) of this article can be found under <https://doi.org/10.1002/qute.201900029>

DOI: 10.1002/qute.201900029

hard problems such as the satisfiability testing,<sup>[10]</sup> and the traveling salesman problem.<sup>[11,12]</sup> Local search is used for problems where a global solution cannot be computed directly but instead can be iteratively approximated in the space of candidate solutions (subproblems), until optimal (or sufficiently good). The important feature of QLS is that it is a hybrid hardware-agnostic algorithm that combines a classical machine with a small quantum device. In this method, QLS allows us to leverage available NISQ-era quantum devices to solve machine learning problems of practical size for the first time.

A version of the network community detection (also known as graph clustering) is an unsupervised machine learning problem used to identify substructure as communities in such networks as computer and information infrastructures, social activities, and biological interactions or co-occurrences. It is used to find non-trivial topological features, with patterns of connection between nodes that are neither exactly regular nor random. For example, in metabolic networks, communities correspond to a series of chemical reactions called metabolic pathways,<sup>[13]</sup> whereas in a protein interaction network, communities correspond to proteins with similar functionality inside a biological cell.<sup>[14]</sup> In this work we focus on using Newman's modularity-based community detection.<sup>[15]</sup>

QLS was applied to solving the 2-community detection problem on real networks of up to 410 nodes, while solving a 16 variable subproblem on a quantum device. To the best of our knowledge, this is the first attempt to tackle problems of this size using gate-model (universal) quantum computing. Also, QLS is shown to work with the D-Wave quantum annealer. We explore the potential of QLS as quantum devices become more and more capable and demonstrate its potential.

The small size of available quantum devices creates a challenge, since typical algorithms (both quantum and classical) look at a problem "as a whole," requiring large amounts of resources to store the description of the entire problem. While on classical computers storing the problem usually does not constitute a problem, it becomes a bottleneck when working with quantum computers that only have limited numbers of qubits and limited connectivity between qubits. The number of variables that can be represented in a quantum device is dependent on its underlying architecture.

A problem decomposition approach like local search presents a natural solution to this problem. A local search heuristic starts with some initial solution and searches its neighborhood iteratively, trying to find a better candidate solution with improved criterion (which is often an objective of the corresponding minimization or maximization of the problem). If a better solution is found, it replaces the current solution, and the search continues.<sup>[16]</sup> Searching the neighborhood is a local problem and its size can be restricted to fit on a small quantum device. In QLS for graph community detection, the neighborhood of the solution is searched by selecting a subset of vertices and collectively moving them between the communities with the goal of improving the global modularity metric.

The QLS approach provides an additional benefit of being fundamentally hardware-agnostic. Local neighborhood search can be encapsulated as a routine, allowing researchers to easily switch between different hardware implementations. This is especially

useful, since the landscape of quantum computing in the NISQ era is in a constant state of flux with many QC architectures available and new development happening constantly. It is not clear at this stage which architecture will become dominant in future. In this work, we demonstrate how the two most developed and popular current paradigms, universal quantum computing (UQC) and quantum annealing (QA), can be integrated into the QLS framework and utilized to solve problems of practical size. Both paradigms have demonstrated great potential on a number of important problems.<sup>[2,17–19]</sup>

In this paper, we do not aim to analyze performance of quantum optimization algorithms like quantum annealing or QAOA. Although we do present some performance results (see Figure 2), they by no means constitute an exhaustive comparison with classical state-of-the-art. Instead, they provide motivation for our work, demonstrating that the subproblems offloaded to quantum solvers are not trivial and that hybridization is needed. For benchmarking, analysis and exploration the reader is referred to one of a number of recent paper analyzing QAOA performance.<sup>[3,5,20]</sup> In other words, we do not focus on finding and quantifying quantum speedups. Instead, we focus on a different question: if these algorithms are indeed capable of providing speedups in the near term, how can we leverage them to solve practical problems?

It is important to point out that the introduction of a problem decomposition heuristic like QLS limits the possible quantum speedup. Since to the best of our knowledge, no asymptotic speedups have been shown so far for QAOA or QA, decomposition schemes limit the multiplicative speedup on the entire global problem by the multiplicative speedup on a small local subproblem. However, they still provide a way to take advantage of the small quantum devices that are becoming available.

The rest of the paper is organized as follows. We begin by introducing the community detection problem and hybrid local-search schemes. Then we describe the QC paradigms we utilize and the quantum algorithms used to perform local search. Finally, we provide the implementation details, present the results and discuss their significance.

## 2. The Community Detection Problem

The community detection problem (or modularity network clustering) is an NP-hard problem<sup>[21]</sup> with a variety of applications in complex systems.<sup>[22]</sup> Practical usefulness and complexity make community detection an interesting problem to tackle using QC. The goal of community detection in a network with an underlying simple undirected graph  $G = (V, E)$  is to split the set of vertices  $V$  into communities such that the modularity is maximized:<sup>[15]</sup>

$$H = \frac{1}{4|E|} \sum_{ij} \left( A_{ij} - \frac{k_i k_j}{2|E|} \right) s_i s_j = \frac{1}{4|E|} \sum_{ij} B_{ij} s_i s_j \quad (1)$$

where the variables  $s_i \in \{-1, +1\}$  indicate community assignment of vertex  $i$  ( $s_i = -1$  meaning vertex  $i$  is assigned to the first community and  $s_j = +1$  meaning that vertex  $j$  is assigned to the second community),  $k_i$  is a degree of  $i \in V$ , and  $A$  is the adjacency matrix of  $G$ . In this work, we focus on clustering the

network into two communities. There are several classical approaches to extend the problem to cases when the number of communities is greater than 2.<sup>[15,23]</sup>

Community detection using a hybrid quantum-classical approach targeted for specific quantum architectures has been demonstrated previously. The 2-community problem was solved using *qbsolv* and the D-Wave quantum annealer<sup>[24]</sup> and extended for *k*-communities.<sup>[25,26]</sup> Solving for 2-communities using QAOA and the IBM Q Experience was shown in ref. [25]. Solving for *k*-communities on signed graphs using block coordinate descent<sup>[27,28]</sup> and D-Wave quantum annealer was shown in ref. [29].

### 3. Quantum-Accelerated Decomposition Heuristics for Optimization

Central to the discipline of QC in the NISQ era is the problem of a limited number of available noisy qubits. For example, at the time of writing, the largest gate-model QC device available on the cloud was IBM Q 20 Tokyo<sup>[30]</sup> with twenty superconducting qubits. Twenty qubits translates into up to 20 variables due to connectivity constraints. This implies that the maximum number of nodes of a network we can cluster directly is 20. This example highlights the challenges of leveraging limited NISQ-era devices to solve practical problems and motivates our local-search approach. Note that same considerations apply for problems other than optimization. For example, similar hybrid approaches have been applied to blind quantum computation,<sup>[31–34]</sup> and distributed quantum machine learning.<sup>[35]</sup> Parallel quantum computation (PQC)<sup>[36]</sup> can be used to speed up Grover's search algorithm<sup>[37]</sup> by dividing a database on which the search is performed between an ensemble of quantum computers running in parallel.<sup>[38,39]</sup>

In response to the challenges of quantum computation in the NISQ era, a number of decomposition approaches have been explored. The methods described in this section use limited in size quantum optimization solvers to search a restricted neighborhood of a given solution with the goal of finding a better solution. Here, the given solution comes either from running a classical heuristic solver or from the previous iteration. These methods are inspired by the success of classical large-scale neighborhood local search methods (the reader is referred to [40] for a survey of local-search heuristics in general and to [41] for a survey of large-scale neighborhood methods in particular). It is important to note that unlike this paper, all the works described in this section focus exclusively on D-Wave quantum annealers.

The first family of methods builds on classical preprocessing methods for quadratic unconstrained binary optimization (QUBO) problems (see ref. [42] for a review). One such preprocessing technique is heuristically fixing variables. The variables are chosen by maintaining a set of elite solutions and fixing the variables that have the same value across many or all local optima, with the intuition being that they will have the same values for the global optimum.<sup>[43]</sup> Sample persistence variable reduction (SPVAR)<sup>[44]</sup> in its basic version uses a sample of solutions (obtained either from a quantum annealer or a classical heuristic) and fixes the variables that have the same value

across the entire sample. Then SPVAR uses a quantum annealer as the solver for the restricted QUBO. This method was later extended by introducing multistart (multiple samples) and was extensively benchmarked using both the D-Wave quantum annealer as well as state-of-the-art classical heuristics for Chimera Hamiltonians.<sup>[45]</sup>

The second family of methods extends iterative large-scale neighborhood local search methods. Local search commonly considers the neighborhood of bit strings that have Hamming distance one from the current solution at each step. The performance of local search methods can be improved by considering larger neighborhoods (Liu et al.<sup>[46]</sup> shows significant performance improvements for neighborhood of Hamming distance four, equivalent to fixing all but four variables). Quantum optimizers provide a potentially efficient way to explore these larger neighborhoods. This rather straightforward idea was introduced in ref. [47] and extended and rigorously tested in refs. [8,48,49]. A similar hybrid tree search method was presented in ref. [50]. These methods utilize the D-Wave quantum annealer as the quantum optimizer, enabling them to solve problems with thousands of variables. In this work, we limit the subproblem size to be small enough to fit on the IBM Q quantum computer, limiting the size of the problems we can tackle. D-Wave provides a set of utilities for problem decomposition, including a hybrid extension of the tabu search QSage.<sup>[51]</sup>

### 4. Quantum Local Search

To address the challenges outlined above, we introduce the QLS algorithm. QLS is a hybrid quantum-classical local-search approach, inspired by numerous existing local-search heuristics. QLS is motivated by the successful application of local-search heuristics to a variety of optimization problems. The novelty of QLS is that it can utilize both quantum annealers and universal quantum computers. In this work, we apply QLS to the problem of 2-community detection on graphs, but the success and versatility of local-search heuristics make us confident that QLS can be extended to other optimization problems.

In QLS for community detection, the local search starts with a random assignment of communities to vertices and attempts to iteratively optimize the current community assignment of a *subset* of vertices with the goal of increasing modularity. Here the space of potential community assignments of a subset of vertices plays the role of the neighborhood where the local search is performed. At each iteration, a subset  $X \subset V$  is populated by selecting vertices with the highest potential gain in modularity obtained when changing their community assignment. This can be done efficiently<sup>[15]</sup> since at each iteration we only need to update the gains of vertices in  $X$  and their neighbors. Then at each iteration, the community assignment of the vertices in the subset  $X$  (subproblem) is optimized using a routine that includes a call to a quantum device. The local search proceeds until it converges. We define convergence as three iterations with no improvement in modularity. Note that in general it is not necessary to consider all vertices before convergence: in the 2-community problem, random initial assignment would be correct for 50% of vertices on average. Our approach is outlined in Algorithm 1.

## Algorithm 1: QLS Community Detection

```

solution = initial_guess(G)
while not converged do
  X = populate_subset(G)
  // using IBM UQC or D-Wave QAOA
  candidate = solve_subproblem(G, X)
  if candidate > solution then
    solution = candidate
  
```

The subproblem of optimizing community assignment of the subset is formulated by fixing community assignment for all vertices not in the subset ( $i \notin X$ ) and encoding them into the optimization problem as boundary conditions. This is a commonly used technique in many heuristics.<sup>[52,53]</sup> Denoting fixed assignments by  $\tilde{s}_j$ , the subproblem can be formulated as:

$$\begin{aligned}
 Q_s &= \sum_{i>j|i,j \in X} 2B_{ij}s_i s_j + \sum_{i \in X} \sum_{j \notin X} 2B_{ij}s_i \tilde{s}_j \\
 &= \sum_{i>j|i,j \in X} 2B_{ij}s_i s_j + \sum_{i \in X} C_i s_i
 \end{aligned}$$

$$\text{where } C_i = \sum_{j \notin X} 2B_{ij}\tilde{s}_j \quad (2)$$

Clearly, maximizing (2) can only increase global modularity (1). The objective defined in Equation (2) can be optimized using a QC algorithm. The exact way the optimization is performed can vary between different QC implementations, making our approach extendable to new emerging QC platforms. We demonstrate this portability by implementing two subproblem optimizing routines that use IBM Q 16 Rueschlikon<sup>[30]</sup> and D-Wave 2000Q.<sup>[54]</sup> Additionally, we implement a subset optimization routine that uses the classical Gurobi solver<sup>[55]</sup> for quality comparison. The choice of Gurobi is not of importance, since for subproblems with 16 variables any classical integer programming solver is capable of finding the optimum.

## 5. Quantum Computing Paradigms

Quantum annealing (QA) is a form of adiabatic quantum computation (AQC).<sup>[56]</sup> QA solves an optimization problem by encoding it as an Ising model Hamiltonian, with the ground state of that Hamiltonian corresponding to the global solution of the optimization problem. The Ising Hamiltonian describes the energy of a collection of  $n$  spin variables, with each variable being in one of two spin states ( $\pm 1$ ). A spin configuration describes assignment of states to spin variables, with  $s_i$  denoting the state of spin variable  $i$  (note that the 2-community problem maps naturally to this system, with the resulting spin state,  $s_i$ , denoting community assignment). The energy of a configuration is then defined by

$$H(s) = \sum_{i>j} J_{ij}s_i s_j + \sum_i h_i s_i \quad (3)$$

where  $h_i$  correspond to external forces applied to spin variables, and  $J_{ij}$  to coupling strengths between pairwise spin interactions.<sup>[56]</sup>

QA finds the ground state of the objective Hamiltonian by performing a quantum evolution. As the initial Hamiltonian, QA uses a transverse field Hamiltonian. It introduces quantum fluctuations that help the annealing process to escape local minima by “tunneling through” hills in the energy landscape, enabling the evolution to move faster than adiabatic requirement would allow. As the evolution is performed, the transverse field Hamiltonian is slowly “turned off” (scaled with a coefficient decreasing to 0), such that the evolution finishes in a system described by the problem Hamiltonian.<sup>[56]</sup>

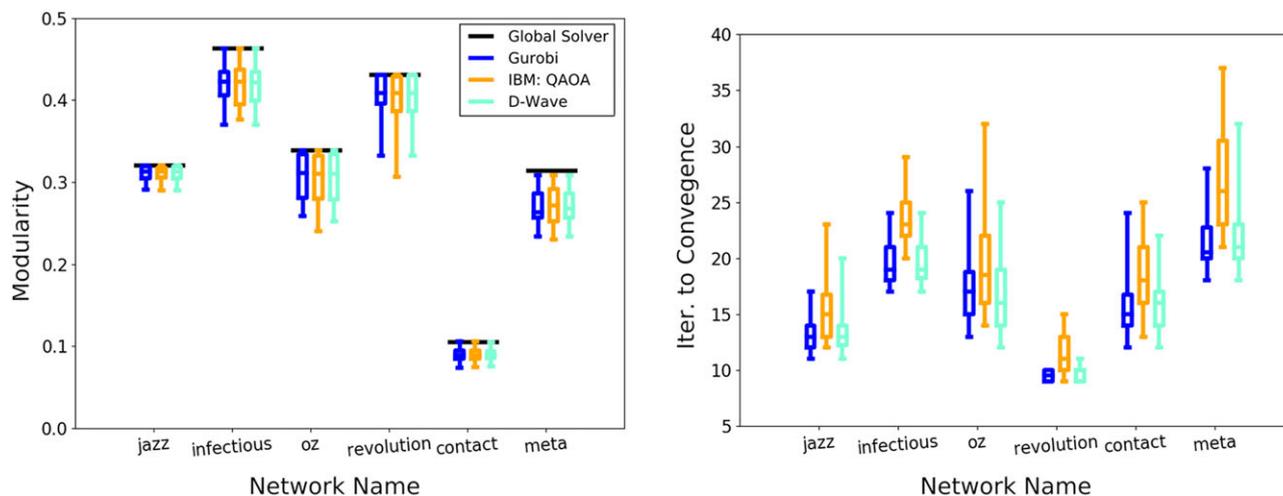
Since AQC was introduced in 2000 by Farhi et al.<sup>[57]</sup>, D-Wave Systems, Inc.,<sup>[54]</sup> IARPA’s Quantum Enhanced Optimization effort<sup>[58]</sup> and other researchers<sup>[59]</sup> have achieved a lot of progress in developing a system implementing QA<sup>[56]</sup> and applying it to a variety of problems, including optimization problems on graphs,<sup>[24]</sup> machine learning,<sup>[60]</sup> traffic flow optimization,<sup>[61]</sup> integer factoring,<sup>[62]</sup> and simulation problems.<sup>[63]</sup> Optimization problems can be solved by QA when formulated in the Ising form (2) or as a quadratic binary optimization (QUBO).

Universal (or gate-based) quantum computing was introduced in the 1980s<sup>[64]</sup> and has seen great theoretical advances since. Shor’s<sup>[65]</sup> and Grover’s<sup>[37]</sup> algorithms are two most famous examples of quantum algorithms with theoretically proven speedups over classical state-of-the-art. Universal quantum computing has been implemented in hardware by a number of companies, national laboratories, and universities.<sup>[66–70]</sup>

To optimize (2) on a universal quantum computer, we use a hybrid quantum-classical approach, QAOA.<sup>[71,72]</sup> Similar to QA, a problem is encoded as an objective Hamiltonian  $H$ . Then a quantum evolution is performed starting with some fixed initial easy-to-prepare state (traditionally, uniform superposition over computational basis states is used). The difference is that unlike QA, in which the evolution is analog, in QAOA the evolution is performed by applying a series of gates parameterized by a vector of variational parameters  $\theta$ . A hybrid approach, combining the quantum device performing the evolution and a classical optimizer, finds the optimal variational parameters. QAOA starts with an initial set of variational parameters  $\theta_0$ . At each step, a multiqubit state  $|\psi(\theta)\rangle$  parameterized by the variational parameters  $\theta$  is prepared on the quantum co-processor. Then a cost function  $E(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle$  is measured and the result is used by the classical optimizer to choose new parameters  $\theta$  with the goal of finding the ground-state energy  $E_G = \min_{\theta} \langle \psi(\theta) | H | \psi(\theta) \rangle$ . QAOA provides a viable path to quantum advantage,<sup>[73]</sup> making it a good algorithm to explore on near-term quantum computers.

## 6. Results and Discussion

We implement the classical part of QLS in Python 3.6, using NetworkX<sup>[74]</sup> for network operations. The subproblem solvers are implemented using QA (D-Wave SAPI), QAOA (IBM QISKit<sup>[75]</sup>) and the classical Gurobi solver.<sup>[55]</sup> Our framework is modular and easily extendable, allowing researchers to add new subproblem solvers as they become available. The framework is available on GitHub at <http://bit.ly/QLSCommunity>.



**Figure 1.** Box-plots showing the range of modularity scores for 2-community detection (left, greater is better) and number of solver calls (right, less is better), respectively, for the three different subproblem solvers. The results show that the proposed approach is capable of achieving results close to the state-of-the-art (Global Solver).

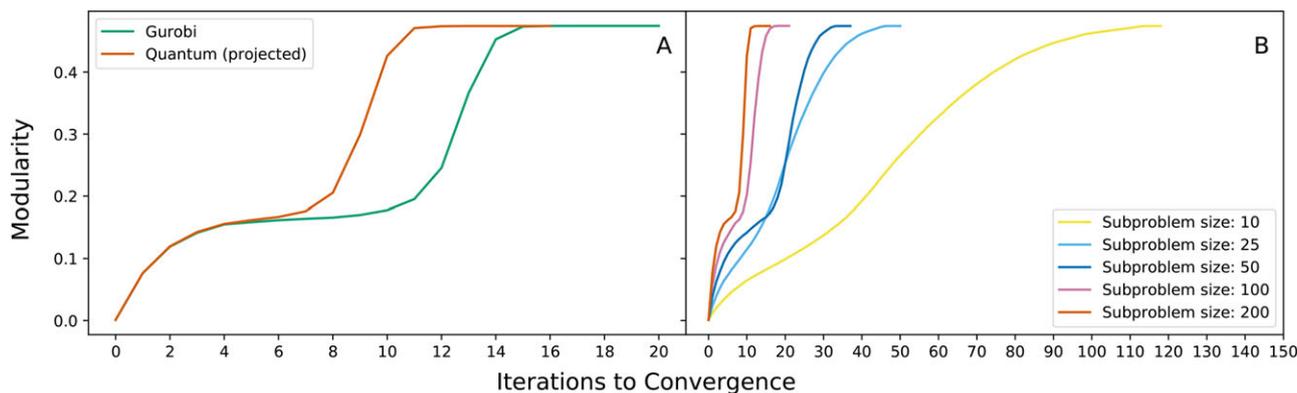
In order for a subproblem to be solved on the D-Wave system, the problem is embedded onto the physical layout (Chimera graph). The clique embedder<sup>[76]</sup> is used to calculate an embedding of a complete 16-variable problem once and is reused for each subproblem. In this work, we utilized D-Wave's Solver API (SAPI) which is implemented in Python 2.7, to interact with the system. We used the D-Wave 2000Q which has up to 2048 available qubits. Subproblems of approximately 64 variables can be solved on the the 2000Q, however, for a fair comparison, we limit ourselves to up to 16 variables. The D-Wave system is intrinsically a stochastic system, where solutions are sampled from a distribution corresponding to the lowest energy state. For each subproblem, the best solution out of 10 000 samples is returned.

The QAOA subproblem solver is implemented using the IBM QISKit framework. We ran QAOA with RYZ ansatz<sup>[77]</sup> on the IBM 16 Q Rueschlikon<sup>[30]</sup> with 16 qubits. For optimization of the variational parameters we used a SciPy<sup>[78]</sup> implementation of Constrained Optimization BY Linear Approximation (COBYLA) method.<sup>[79]</sup> For each subproblem, we performed optimization of the variational parameters  $\theta$  using a high-performance simulator<sup>[80]</sup> and ran QAOA with optimized parameters on a quantum device using the IBM Q Experience<sup>[30]</sup> cloud service. We allowed COBYLA 100 function evaluations (i.e., 100 QAOA runs on the simulator) to find optimal parameters  $\theta$ . We used this setup (training on a simulator and running on the quantum device) because of the limitations of the IBM Q Experience job queue at this time. In our experience, jobs submitted to the IBM quantum device can spend minutes to hours in queue, requiring days to complete a full variational parameter optimization loop. It is our understanding that this will be remedied in the future. The main downside of this setup is that the variational parameters trained on a simulator do not encode the noise profile of the device, decreasing the quality of the solution. This is one of the main factors contributing to slightly slower convergence for QAOA compared to other methods. In the future, as various QC devices become available, it will be straightforward to perform QAOA fully on a QC device. However, even using the cur-

rent setup we achieved very promising results, indicating great potential for applying variational quantum-classical methods to combinatorial optimization problems.

Our results are presented in **Figure 1**. We ran our algorithm on six real-world networks from the Koblenz Network Collection (KONECT) dataset<sup>[81]</sup> with up to 410 nodes as our benchmark. The networks come from different real world phenomena and include social and metabolic networks. For each network, we ran 30 experiments with different random seeds. The same set of seeds was used by the three subproblem solvers, with all solvers starting with the same initial guess and therefore making the results directly comparable. We fixed the subproblem size at 16 vertices. Our results demonstrate that QLS with both D-Wave QA and QAOA on IBM Q as quantum subproblem solvers perform similarly in terms of quality of the solution (modularity) and the number of iterations to convergence, and are capable of achieving results comparable to state-of-the-art. Our results are compared to results using the Gurobi Optimizer, which is a state-of-the-art solver for mathematical programming. We use the Gurobi Optimizer in two ways: first as a solver for solving the entire problem at once, which we report as the *Global Solver* and second as a solver for solving small size subproblems of fixed size within the local search framework. For solving the entire problem, the Gurobi Optimizer is unable to reach a provable global optima for most of the problems within the specified time frame. For the graph problems of up to approximately 400 variables, we run Gurobi (as a global solver) for up to 72 h and the results reported are within an optimality gap of up to 33%. For the smaller size subproblems of 16 variables, Gurobi was able to find the optimal solution within less than second. The networks and the set of seeds we used are available online at <http://bit.ly/QLSdata>.

The results demonstrate the promise of the proposed approach. We presented a framework that is able to find 2 communities in graphs of size up to 410 vertices using only NISQ-era devices. We explored the potential of our approach as new and better QC hardware becomes available in two ways. First, we used the classical Gurobi solver<sup>[55]</sup> to simulate the performance



**Figure 2.** A projection of QLS performance as the hardware size of quantum devices increases. A) Projected QLS performance as the quality of the local search solver solution is improved. The projection is performed by comparing the performance of classical solver Gurobi with time limit fixed at 0.25 s (D-Wave time to solution) and Gurobi with time limit 1000 s (projected good solution). The assumption is that the new quantum optimization algorithms would be able to scale and provide results of the same quality as Gurobi with time limit 1000 s while taking approximately the same time to solve the problem as they do today. B) Projected number of iterations for QLS to converge as larger devices become available (projection performed by using Gurobi as a subproblem solver).

improvements in QLS as the subproblem size is increased (see **Figure 2B**). We generate a 2000 node random graph with realistic community structure and known modularity.<sup>[82]</sup> Unsurprisingly, QLS finds the optimal solution faster (using fewer local search iterations) as the subproblem size increases. Second, we demonstrate the need for quantum acceleration by demonstrating the limitations of existing state-of-the-art solvers. We used Gurobi<sup>[55]</sup> as a subproblem solver with subproblem size of 200. **Figure 2A** shows that for the subproblem of this size, Gurobi cannot produce a good solution quickly. We compared Gurobi with time limit 0.25 s (the running time of QA on D-Wave) with Gurobi with time limit 1000 s, with the assumption that Gurobi would converge to a good solution. We use the running time of QA as our estimate because at the time of writing we do not have a good way of measuring the running time of QAOA due to the architecture of the IBM Q Experience. We expect QAOA to have similar performance. This assumes that quantum methods would scale well to larger problems, which is a strong assumption. However, the goal here is to motivate the exploration of quantum optimization heuristics by showing the limitations of classical state-of-the-art and not to demonstrate quantum advantage. Using a better solution within the local search enables 25% (4 iterations) improvement in time to convergence (convergence is defined as three iterations with no improvement). This demonstrates that the subproblems become computationally hard even for sizes that are small enough to potentially fit on near-term devices. It is important to note that even though in our experiments Gurobi performed better than other integer programming solvers, it is quite possible that other solvers can perform better on this problem, especially after tuning. Indeed, in the past the improvements in classical heuristics have forced researchers to downgrade claims of quantum advantage.<sup>[83,84]</sup> However, demonstrating quantum advantage is outside of scope of this paper. Instead, we use these results to motivate our hybrid approach by showing the computational complexity of the subproblems offloaded to quantum solvers. As quantum solvers improve and become capable of providing speedups at *subproblem* level, our QLS will enable us to leverage these speedups at the global problem level.

## 7. Conclusion

In the next few years, a number of QC hardware implementations are expected to become mature enough to be applied to practically important problems. QC using trapped ions<sup>[85]</sup> and Rydberg atom arrays<sup>[86]</sup> are just two examples of quantum hardware now moving out of the laboratory, with the potential to realize quantum advantage. However, none of them promise to deliver more than a few hundred qubits in the near future. Therefore, we believe the future of QC is hybrid, with algorithms combining both classical and quantum computation. QLS presents a path to integrate NISQ-era devices into computational workflows in a flexible way, both in terms of adding different hardware backends and extending to different problems. Classical local search heuristics have proven useful for a variety of problems in many fields.<sup>[40]</sup> We believe that QLS can be similarly extended to problems beyond network community detection. We also believe that the decomposition approaches like QLS can improve dramatically the speed and performance of QAOA algorithms on universal quantum computers, which might be the key to achieve quantum advantage on NISQ devices.

## Acknowledgements

This research used quantum computing system resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. This research used the resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy (DOE) Office of Science User Facility supported under Contract DE-AC02-06CH11357. Y.A. and R.S. were supported by the DOE Office of Science. The authors gratefully acknowledge the computing resources provided and operated by the Joint Laboratory for System Evaluation (JLSE) at Argonne National Laboratory. The authors would also like to acknowledge the NNSA's Advanced Simulation and Computing (ASC) program at Los Alamos National Laboratory (LANL) for use of their Ising D-Wave 2000Q quantum computing resource. LANL is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract

No. 89233218NCA000001). Clemson University is acknowledged for generous allotment of compute time on Palmetto cluster.

## Conflict of Interest

The authors declared no conflict of interest.

## Keywords

graph clustering, quantum annealing, quantum approximate optimization algorithm

Received: April 11, 2019

Revised: May 12, 2019

Published online:

- [1] J. Preskill, arXiv:1801.00862, **2018**.
- [2] V. Dunjko, Y. Ge, J. I. Cirac, arXiv:1807.08970, **2018**.
- [3] O. D. Parekh, C. Ryan-Anderson, S. Gharibian, *Quantum Optimization and Approximation Algorithms*, Technical Report, Sandia National Laboratories, Albuquerque, NM and Livermore, CA **2019**, <https://doi.org/10.2172/1492737>
- [4] Z. Wang, S. Hadfield, Z. Jiang, E. G. Rieffel, *Phys. Rev. A* **2018**, 97, 022304.
- [5] G. Guerreschi, A. Matsuura, arXiv:1812.07589, **2018**.
- [6] E. Tang, arXiv:1807.04271, **2018**.
- [7] E. Grumblin, M. Horowitz (Eds.), *Quantum Computing: Progress and Prospects*, The National Academies Press, Washington D.C. **2018**, <https://doi.org/10.17226/25196>
- [8] Z. Bian, F. Chudak, R. B. Israel, B. Lackey, W. G. Macready, A. Roy, *Front. ICT* **2016**, 3, 14.
- [9] *Local Search in Combinatorial Optimization*, 1st ed. (Eds: E. Aarts, J. K. Lenstra), John Wiley & Sons, New York **1997**.
- [10] B. Selman, H. A. Kautz, B. Cohen, in *Cliques, Coloring, and Satisfiability* (Eds: D. S. Johnson, M. A. Trick), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 26, American Mathematical Society **1993**, p. 521.
- [11] Y. Wu, T. Weise, R. Chiong, in *2015 IEEE 14th Int. Conf. on Cognitive Informatics & Cognitive Computing (ICCI\*CC)*. IEEE, Piscataway, NJ **2015**, p. 213.
- [12] D. Johnson, in *Proc. of 17th Int. Colloquium on Automata Languages and Programming* (Lecture Notes in Computer Science), Springer-Verlag, Berlin **1990**, p. 443.
- [13] N. W. Hanson, K. M. Konwar, A. K. Hawley, T. Altman, P. D. Karp, S. J. Hallam, *BMC genomics* **2014**, 15, 619.
- [14] J. Chen, B. Yuan, *Bioinformatics* **2006**, 22, 2283.
- [15] M. E. J. Newman, *Proc. Natl. Acad. Sci. USA* **2006**, 103, 8577.
- [16] E. Aarts, E. H. Aarts, J. K. Lenstra, *Local Search in Combinatorial Optimization*, Princeton University Press, Princeton, NJ **2003**.
- [17] A. D. King, J. Carrasquilla, J. Raymond, I. Ozfidan, E. Andriyash, A. Berkley, M. Reis, T. Lanting, R. Harris, F. Altomare, K. Boothby, P. I. Bunyk, C. Enderud, A. Fréchet, E. Hoskinson, N. Ladizinsky, T. Oh, G. Poulin-Lamarre, C. Rich, Y. Sato, A. Yu. Smirnov, L. J. Swenson, M. H. Volkman, J. Whittaker, J. Yao, E. Ladizinsky, M. W. Johnson, J. Hilton, M. H. Amin, *Nature* **2018**, 560, 456.
- [18] J. Romero, R. Babbush, J. McClean, C. Hempel, P. Love, A. Aspuru-Guzik, *Quantum Sci. Technol.* **2018**, 4, 014008.
- [19] A. Ambainis, K. Balodis, J. Iraids, M. Kokainis, K. Prūsīs, J. Vihrovs, arXiv:1807.05209 **2018**.
- [20] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, M. D. Lukin, arXiv:1812.01041, **2018**.
- [21] U. Brandes, D. Delling, M. Gaertler, R. Görke, M. Hofer, Z. Nikoloski, D. Wagner, arXiv: physics/0608255, **2006**.
- [22] M. Newman, *Networks: An Introduction*, Oxford University Press, Oxford **2010**.
- [23] M. C. Nascimento, A. C. De Carvalho, *Eur. J. Oper. Res.* **2011**, 211, 221.
- [24] H. Ushijima-Mwesigwa, C. F. A. Negre, S. M. Mniszewski, in *Proc. of the Second Int. Workshop on Post Moores Era Supercomputing*. ACM, Denver, CO **2017**, pp. 22–29.
- [25] S. M. Mniszewski, C. F. A. Negre, H. Ushijima-Mwesigwa, *Hybrid Quantum Computing for Graph Partitioning and Community Detection*, ISTI Rapid Response 2018 Hands-on Quantum Computing, **2018**.
- [26] C. F. Negre, H. Ushijima-Mwesigwa, S. M. Mniszewski, arXiv:1901.09756, **2019**.
- [27] I. Zintchenko, M. B. Hastings, M. Troyer, *Phys. Rev. B* **2015**, 91, 024201.
- [28] G. Rosenberg, M. Vazifeh, B. Woods, E. Haber, *Comput. Opt. Appl.* **2016**, 65, 845.
- [29] E. Zahedinejad, D. Crawford, C. Adolphs, J. S. Oberoi, arXiv:1901.04873, **2019**.
- [30] IBM Q: Quantum devices and simulators, <https://www.research.ibm.com/ibm-q/technology/devices/> (accessed: September 2018).
- [31] V. Giovannetti, L. Maccone, T. Morimae, T. G. Rudolph, *Phys. Rev. Lett.* **2013**, 111, 230501.
- [32] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, P. Walther, *Science* **2012**, 335, 303.
- [33] Y.-B. Sheng, L. Zhou, *Sci. Rep.* **2015**, 5, 7815.
- [34] Y.-B. Sheng, L. Zhou, *Phys. Rev. A* **2018**, 98, 052343.
- [35] Y.-B. Sheng, L. Zhou, *Sci. Bull.* **2017**, 62, 1025.
- [36] G. L. Long, L. Xiao, *Phys. Rev. A* **2004**, 69, 052303.
- [37] L. K. Grover, in *Proc. of The Twenty-Eighth Annual ACM Symp. on Theory of Computing*. ACM, Philadelphia, PA **1996**, pp. 212–219.
- [38] L. Xiao, G. Long, *Phys. Rev. A* **2002**, 66, 052320.
- [39] G. L. Long, L. Xiao, *J. Chem. Phys.* **2003**, 119, 8473.
- [40] R. Rotta, A. Noack, *JEA* **2011**, 16, 2.
- [41] R. K. Ahuja, Ö. Ergun, J. B. Orlin, A. P. Punnen, *Discrete Appl. Math.* **2002**, 123, 75.
- [42] G. Tavares, Ph.D. thesis, Rutgers University – Graduate School, New Brunswick, NJ **2008**.
- [43] Y. Wang, Z. Lü, F. Glover, J.-K. Hao, in *European Conf. on Evolutionary Computation in Combinatorial Optimization*. Springer, Torino, Italy **2011**, pp. 72–83.
- [44] H. Karimi, G. Rosenberg, *Quantum Inf. Process.* **2017**, 16, 166.
- [45] H. Karimi, G. Rosenberg, H. G. Katzgraber, *Phys. Rev. E* **2017**, 96, 043312.
- [46] W. Liu, D. Wilkins, B. Alidaee, *A Hybrid Multi-Exchange Local Search for Unconstrained Binary Quadratic Program*, Hearin Center for Enterprise Science, Working paper series HCES-09-05, The University of Mississippi **2005**.
- [47] Z. Bian, F. Chudak, R. Israel, B. Lackey, W. G. Macready, A. Roy, *Front. Phys.* **2014**, 2, 56.
- [48] G. Rosenberg, M. Vazifeh, B. Woods, E. Haber, *Comput. Optim. Appl.* **2016**, 65, 845.
- [49] I. Zintchenko, M. B. Hastings, M. Troyer, *Phys. Rev. B* **2015**, 91, 024201.
- [50] T. T. Tran, M. Do, E. G. Rieffel, J. Frank, Z. Wang, B. O’Gorman, D. Venturelli, J. C. Beck, in *Ninth Annual Symposium on Combinatorial Search*, AAAI Press, New York **2016**.
- [51] D-Wave system documentation: Problem decomposition, **2019**, [https://docs.dwavesys.com/docs/latest/c\\_handbook\\_7.html#toolsdecomposition](https://docs.dwavesys.com/docs/latest/c_handbook_7.html#toolsdecomposition) (accessed: March 2019).
- [52] S. Leyffer, I. Safro, *J. Complex Netw.* **2013**, 1, 183.
- [53] W. W. Hager, J. T. Hungerford, I. Safro, *Comput. Opt. Appl.* **2018**, 69, 189.

- [54] D-Wave Systems Inc., 2018, [www.dwavesys.com/tutorials/background-reading-series/introduction-d-wave-quantum-hardware](http://www.dwavesys.com/tutorials/background-reading-series/introduction-d-wave-quantum-hardware)
- [55] Gurobi optimizer reference manual, 2014, <http://www.gurobi.com>
- [56] C. C. McGeoch, *Synth. Lect. Quantum Comput.* **2014**, 5, 1.
- [57] E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, arXiv:quant-ph/0001106, 2000.
- [58] Quantum Enhanced Optimization (QEO), <https://www.iarpa.gov/index.php/research-programs/qeo> (accessed: September 2018).
- [59] S. Novikov, R. Hinkey, S. Disseler, J. I. Basham, T. Albash, A. Risinger, D. Ferguson, D. A. Lidar, K. M. Zick, arXiv:1809.04485, 2018.
- [60] D. O'Malley, V. V. Vesselinov, B. S. Alexandrov, L. B. Alexandrov, arXiv:1704.01605 2017.
- [61] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, B. Parney, *Front. ICT* **2017**, 4, 1.
- [62] W. Peng, B. Wang, F. Hu, Y. Wang, X. Fang, X. Chen, C. Wang, *Sci. China: Phys., Mech. Astron.* **2019**, 62, 60311.
- [63] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, K. Boothby, P. Bunyk, C. Deng, C. Enderud, S. Huang, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, T. Lanting, R. Li, T. Medina, R. Molav, R. Neufeld, T. Oh, I. Pavlov, I. Perminov, G. Poulin-Lamarre, C. Rich, A. Smirnov, L. Swenson, N. Tsai, M. Volkmann, J. Whittaker, *Science* **2017**, 361, 162.
- [64] D. Deutsch, *Proc. R. Soc. Lond. A* **1985**, 400, 97.
- [65] P. W. Shor, in *1994 Proc. of 35th Annual Symp. on Foundations of Computer Science*, IEEE **1994**, pp. 124–134.
- [66] C. Ballance, T. Harty, N. Linke, M. Sepiol, D. Lucas, *Phys. Rev. Lett.* **2016**, 117, 060504.
- [67] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, J. M. Martinis, *Nature* **2014**, 508, 500.
- [68] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, A. C. Gossard, *Science* **2005**, 309, 2180.
- [69] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, F. K. Wilhelm, *New J. Phys.* **2018**, 20, 080201.
- [70] M. Saffman, T. G. Walker, K. Mølmer, *Rev. Mod. Phys.* **2010**, 82, 2313.
- [71] E. Farhi, J. Goldstone, S. Gutmann, arXiv:1411.4028, 2014.
- [72] E. Farhi, J. Goldstone, S. Gutmann, arXiv:1412.6062, 2014.
- [73] E. Farhi, A. W. Harrow, arXiv:1602.07674, 2016.
- [74] A. A. Hagberg, D. A. Schult, P. J. Swart, in *Proc. of the 7th Python in Science Conf.* (Eds: G. Varoquaux, T. Vaught, J. Millman), Pasadena, CA **2008**, pp. 11–15.
- [75] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, S. De La Puente González, E. De La Torre, D. Ding, E. Dumitrescu, I. Duran, P. Eendebak, M. Everitt, I. F. Sertage, A. Frisch, A. Fuhrer, J. Gambetta, B. G. Gago, J. Gomez-Mosquera, D. Greenberg et al., *Qiskit: An Open-source Framework for Quantum Computing*, Zenodo **2019**.
- [76] T. Boothby, A. D. King, A. Roy, *Quantum Inf. Process.* **2016**, 15, 495.
- [77] IBM QISKit Aqua: Variational forms, [https://github.com/Qiskit/aqua/blob/master/qiskit\\_aqua/algorithms/components/variational\\_forms/ryrz.py](https://github.com/Qiskit/aqua/blob/master/qiskit_aqua/algorithms/components/variational_forms/ryrz.py) (accessed: September 2018).
- [78] E. Jones, T. Oliphant, P. Peterson, SciPy: Open source scientific tools for Python, <http://www.scipy.org/> (accessed: September 2018).
- [79] M. J. Powell, in *Advances in Optimization and Numerical Analysis* (Eds: S. Gomez, J. P. Hennart), Springer, The Netherlands **1994**, pp. 51–67.
- [80] C. J. Wood, J. A. Smolin, A C++ quantum circuit simulator with realistic noise, <https://github.com/Qiskit/qiskit-terra/tree/master/src/qasm-simulator-cpp> (accessed: September 2018).
- [81] J. Kunegis, in *Proc. of the 22nd Int. Conf. on World Wide Web*. ACM, Rio de Janeiro, Brazil **2013**, pp. 1343–1350.
- [82] P. Sah, L. O. Singh, A. Clauset, S. Bansal, *BMC Bioinf.* **2014**, 15, 220.
- [83] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, H. G. Katzgraber, *Phys. Rev. A* **2016**, 94, 022337.
- [84] S. Mandrà, H. G. Katzgraber, C. Thomas, *Quantum Sci. Technol.* **2017**, 2, 038501.
- [85] M. Chmielewski, J. Amini, K. Hudek, J. Kim, J. Mizrahi, C. Monroe, K. Wright, D. Moehring, *Bull. Am. Phys. Soc.* **2018**.
- [86] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, M. D. Lukin, arXiv:1808.10816, 2018.