

Network Community Detection On Small Quantum Computers

Ruslan Shaydulin

E-mail: rshaydu@g.clemson.edu

School of Computing, Clemson University, Clemson, SC 29634

Hayato Ushijima-Mwesigwa

School of Computing, Clemson University, Clemson, SC 29634

Ilya Safro

School of Computing, Clemson University, Clemson, SC 29634

Susan Mniszewski

Computer, Computational, & Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545

Yuri Alexeev

Computational Science Division, Argonne National Laboratory, Argonne, IL 60439

Abstract. In recent years a number of quantum computing devices with small numbers of qubits became available. We present a hybrid quantum local search (QLS) approach that combines a classical machine and a small quantum device to solve problems of practical size. The proposed approach is applied to the network community detection problem. QLS is hardware-agnostic and easily extendable to new quantum devices as they become available. We demonstrate it to solve the 2-community detection problem on networks from the KONECT dataset collection using the 16-qubit IBM quantum device and D-Wave 2000Q. We also show the potential for quantum advantage as larger quantum devices become available in the future.

Reproducibility: our code and data are available at <http://bit.ly/QLSCommunity>.

1. Introduction

The recent years saw rapid progress in 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, near-term intermediate scale quantum (NISQ) devices are widely expected to be limited up to a few hundred, or perhaps a few thousand qubits [1], severely restricting the size of the problems that can be tackled directly. As the potential of these NISQ-era quantum devices is becoming evident [2], there is an increasing interest in developing algorithms that leverage the small quantum devices that are becoming available. This requires the use of hybrid classical-quantum approaches where a problem is solved across a CPU 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, it has been shown that quantum speedup using the hybrid quantum approximate optimization algorithm (QAOA) for a similar problem (max-cut) requires at least several hundred qubits [3]. Research and development into quantum algorithms is necessary as the number and quality of qubits is improving. These quantum algorithms can also be used to improve classical algorithms [4]. The need for development of new quantum algorithms was highlighted in the recent National Academy of Science report [5]. One of the important directions to make quantum computing feasible in the near future is to use various divide-and-conquer 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.

In this work, we introduce the quantum local search (QLS) algorithm based on the local search (LS) method [6]. LS has been applied to numerous hard computational problems like the Metropolis-Hastings algorithm, WalkSAT, and the 2-opt algorithm for the Traveling Salesman Problem. LS is used for problems where an overall solution can be found by iterative solution in the space of candidate solutions (sub-problems), until optimal.

The important feature of QLS is that it is a hybrid 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. We apply our QLS algorithm to the problem of community detection.

Community detection in complex networks (also known as graph clustering) is an unsupervised machine learning method used to identify sub-structure as communities in computer and information networks, social networks, and biological networks. 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 [7],

whereas in a protein interaction network, communities correspond to proteins with similar functionality inside a biological cell [8]. In this work we focus on using Newman’s modularity-based community detection [9].

QLS was applied to state-of-the-art on real networks of up to 410 nodes solving the 2-community detection problem, 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 to outperform classical local-search analogs.

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.

The local search approach 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 [10]. 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 for producing quantum speedups for a number of important problems [11, 12, 13, 2]

The rest of the paper is organized as follows. We begin by introducing the community detection problem and our hybrid local-search scheme. 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 [14] with a variety of applications in complex systems [15, 16, 17, 18]. 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 [9]:

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

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 [9, 19].

Community detection using a hybrid classical-quantum approach targeted for specific quantum architectures has been demonstrated previously. The 2-community problem was solved using *qbsolv* and the D-Wave quantum annealer [20] and extended for k -communities [21]. Solving for 2-communities using QAOA and the IBM Q Experience was shown in [21].

3. Quantum Local Search

To address the challenges outlined in Section 1, we introduce the QLS algorithm. QLS is a hybrid classical-quantum local-search approach, inspired by numerous existing local-search heuristics (see [22] for a survey). 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 provide an efficient way to perform the local search, provided the neighborhood is small enough to fit on the QC device. 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 [9] 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 QA
  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 [23, 24]. Denoting fixed assignments by \tilde{s}_j , the subproblem can be formulated as:

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

Clearly, maximizing (2) can only increase global modularity (1). The objective defined in Eq. (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 [25] and D-Wave 2000Q [26]. Additionally, we implement a subset optimization routine that uses the classical Gurobi solver [27] for quality comparison.

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 [25] 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.

4. Quantum Computing Paradigms

Quantum annealing (QA) is a form of adiabatic quantum computation (AQC) [28]. 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 (± 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 [28].

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 [28].

Since AQC was introduced in 2000 by Farhi et al. [29], D-Wave Systems Inc [26], IARPA’s QEO effort [30] and other researchers [31] have achieved a lot of progress in developing a system implementing QA [28] and applying it to a variety of problems, including optimization problems on graphs [20], machine learning [32], traffic flow optimization [33], and simulation problems [34]. 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 [35] and has seen great theoretical advances since. Shor’s [36] and Grover’s [37] algorithms are two most famous examples of quantum algorithms with theoretically proven speed-ups over classical state-of-the-art. Universal quantum computing has been implemented in hardware by a number of companies, national laboratories and universities [38, 39, 40, 41, 42].

To optimize (2) on a universal quantum computer, we use a hybrid classical-quantum approach, Quantum Approximate Optimization Algorithm (QAOA) [43, 44]. 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 θ . 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 θ_0 . At each step, a multi-qubit state $|\psi(\theta)\rangle$ parameterized by the variational parameters θ 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

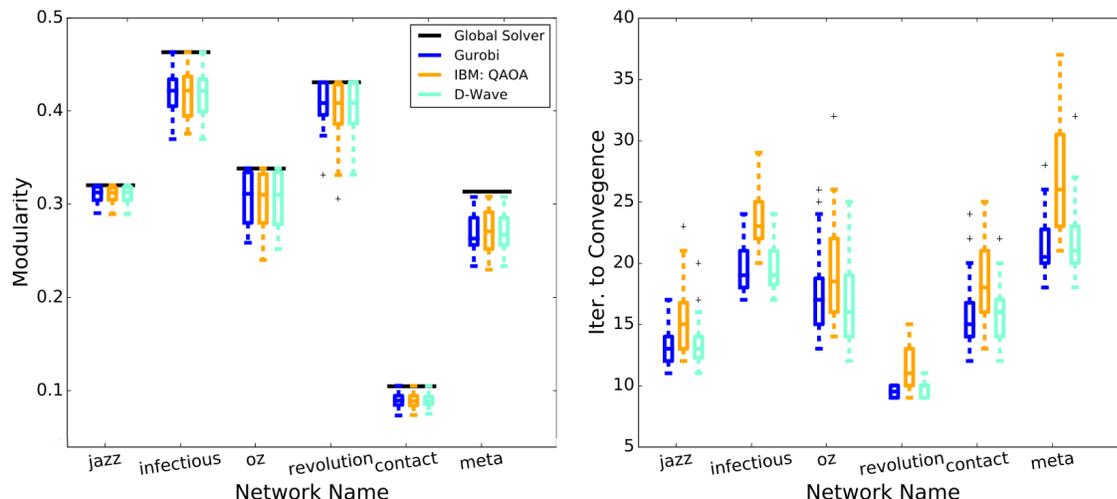


Figure 1. Box-plots comparing modularity scores for 2-community detection (left, greater is better) and number of solver calls (right, less is better) respectively for the three different subproblem solvers. The results show that the proposed approach is capable of achieving results close to the state-of-the-art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

optimizer to choose new parameters θ 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 [45], making it a good algorithm to explore on near-term quantum computers.

5. Results and Discussion

We implement the classical part of QLS in Python 3.6, using NetworkX [46] for network operations. The subproblem solvers are implemented using QA (D-Wave SAPI), QAOA (IBM QISKit [47]) and classical Gurobi solver [27]. 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>.

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 [48] 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 [49] on the IBM 16 Q Rueschlikon [25] with 16 qubits.

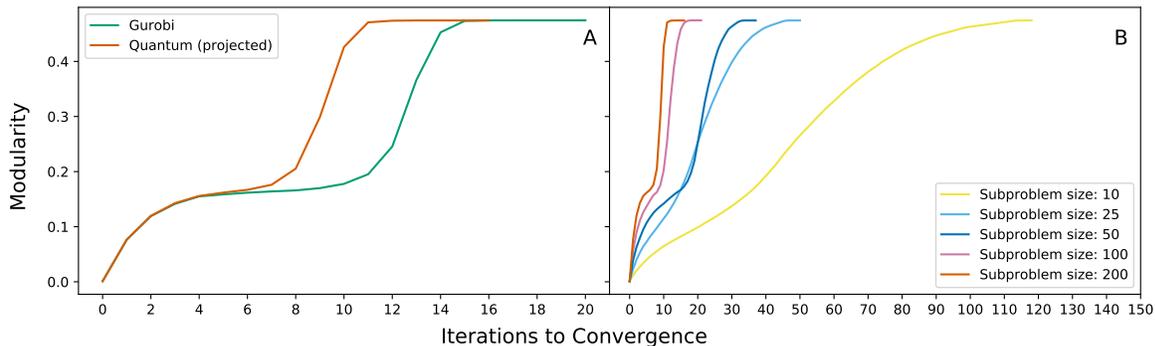


Figure 2. On the left (A), a projection of QLS performance as the quality of the solution from the local search solver is improved. The projection is performed by comparing the performance of classical solver Gurobi with time limit fixed at 0.25s (D-Wave time to solution) and Gurobi with time limit 1000s (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 1000s while taking approximately the same time to solve the problem as they do today. On the right (B), a projection of QLS performance improvement as larger devices become available (projection performed by using Gurobi as a subproblem solver).

For optimization of the variational parameters we used a SciPy [50] implementation of Constrained Optimization BY Linear Approximation (COBYLA) method [51]. For each subproblem, we performed optimization of the variational parameters θ using a high-performance simulator [52] and ran QAOA with optimized parameters on a quantum device using the IBM Q Experience [25] cloud service. We allowed COBYLA 100 function evaluations (i.e. 100 QAOA runs on the simulator) to find optimal parameters θ . 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 current setup we achieved very promising results, indicating great potential for applying variational classical-quantum methods to combinatorial optimization problems.

Our results are presented in Figure 1. We ran our algorithm on six real-world networks from the KONECT dataset [53] 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 hours 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 [27] to simulate the performance 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 [54]. 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 [27] 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 0.25s (the running time of QA on D-Wave) with Gurobi with time limit 1000s, 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 as the quantum devices and quantum optimization algorithms improve enough to handle problems of size 200 or more (i.e. produce solutions comparable to Gurobi after 1000s), there is potential for them to outperform the state-of-the-art classical optimization solvers. It's important to note that since Gurobi has trouble finding a good solution even for

a problem of size 200 *within* the local search, in general it cannot tackle the global problem directly. These results make a clear case for why we need a *quantum* local search framework: current state-of-the-art *classical* optimization solvers cannot provide the solution of desired quality quickly enough even for subproblems small enough to potentially fit on NISQ-era quantum devices.

6. 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 [55] and Rydberg atom arrays [56] 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 [22]. We believe that QLS can be similarly extended to problems beyond network community detection.

7. 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. We 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 NNSAs Advanced Simulation and Computing (ASC) program at Los Alamos National Laboratory (LANL) for use of their Ising D-Wave 2X quantum computing resource and D-Wave Systems Inc. for use of their 2000Q resource. The LANL research contribution has been funded by LANL Laboratory Directed Research and Development (LDRD) and the Advanced Simulation and Computing (ASC) program. 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.

- [1] John Preskill. Quantum computing in the nisq era and beyond. *arXiv preprint arXiv:1801.00862*, 2018.
- [2] Vedran Dunjko, Yimin Ge, and J Ignacio Cirac. Computational speedups using small quantum devices. *arXiv preprint arXiv:1807.08970*, 2018.
- [3] GG Guerreschi and AY Matsuura. Qaoa for max-cut requires hundreds of qubits for quantum speed-up. *arXiv preprint arXiv:1812.07589*, 2018.

- [4] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. *arXiv preprint arXiv:1807.04271*, 2018.
- [5] Emily Grumbling and Mark Horowitz. Quantum computing: Progress and prospects. *National Academies of Sciences, Engineering, and Medicine*, 2018.
- [6] Hisao Ishibuchi and Tadahiko Murata. A multi-objective genetic local search algorithm and its application to flowshop scheduling. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 28(3):392–403, 1998.
- [7] Niels W Hanson, Kishori M Konwar, Alyse K Hawley, Tomer Altman, Peter D Karp, and Steven J Hallam. Metabolic pathways for the whole community. *BMC genomics*, 15(1):619, 2014.
- [8] Jingchun Chen and Bo Yuan. Detecting functional modules in the yeast protein–protein interaction network. *Bioinformatics*, 22(18):2283–2290, 2006.
- [9] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Science*, 103:8577–8582, June 2006.
- [10] Emile Aarts, Emile HL Aarts, and Jan Karel Lenstra. *Local search in combinatorial optimization*. Princeton University Press, 2003.
- [11] Andrew D. King, Juan Carrasquilla, Jack Raymond, Isil Ozfidan, Evgeny Andriyash, Andrew Berkley, Mauricio Reis, Trevor Lanting, Richard Harris, Fabio Altomare, et al. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature*, 560(7719):456, 2018.
- [12] Jonathan Romero, Ryan Babbush, Jarrod McClean, Cornelius Hempel, Peter Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 2018.
- [13] Andris Ambainis, Kaspars Balodis, Jānis Iraids, Martins Kokainis, Krišjānis Prūsis, and Jevgēnijs Vihrovs. Quantum speedups for exponential-time dynamic programming algorithms. *arXiv preprint arXiv:1807.05209*, 2018.
- [14] Ulrik Brandes, Daniel Dellinger, Marco Gaertler, Robert Görke, Martin Hofer, Zoran Nikoloski, and Dorothea Wagner. Maximizing modularity is hard. *arXiv preprint physics/0608255*, 2006.
- [15] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814, 2005.
- [16] Gang Su, Allan Kuchinsky, John H Morris, David J States, and Fan Meng. Glay: community structure analysis of biological networks. *Bioinformatics*, 26(24):3135–3137, 2010.
- [17] Giampiero Bardella, Angelo Bifone, Andrea Gabrielli, Alessandro Gozzi, and Tiziano Squartini. Hierarchical organization of functional connectivity in the mouse brain: a complex network approach. *Scientific reports*, 6:32060, 2016.
- [18] Carlo Nicolini, Cécile Bordier, and Angelo Bifone. Community detection in weighted brain connectivity networks beyond the resolution limit. *Neuroimage*, 146:28–39, 2017.
- [19] Maria CV Nascimento and Andre CPLF De Carvalho. Spectral methods for graph clustering—a survey. *European Journal of Operational Research*, 211(2):221–231, 2011.
- [20] Hayato Ushijima-Mwesigwa, Christian F. A. Negre, and Susan M. Mniszewski. Graph partitioning using quantum annealing on the d-wave system. In *Proceedings of the Second International Workshop on Post Moores Era Supercomputing*, pages 22–29. ACM, 2017.
- [21] Susan M. Mniszewski, Christian F. A. Negre, and Hayato Ushijima-Mwesigwa. Hybrid Quantum Computing for Graph Partitioning and Community Detection. ISTI Rapid Response 2018 Hands-on Quantum Computing, 2018.
- [22] Randolph Rotta and Andreas Noack. Multilevel local search algorithms for modularity clustering. *Journal of Experimental Algorithmics (JEA)*, 16:2–3, 2011.
- [23] Sven Leyffer and Ilya Safro. Fast response to infection spread and cyber attacks on large-scale networks. *Journal of Complex Networks*, 1(2):183–199, 2013.
- [24] William W Hager, James T Hungerford, and Ilya Safro. A multilevel bilinear programming algorithm for the vertex separator problem. *Computational Optimization and Applications*, 69(1):189–223, 2018.
- [25] IBM Q: Quantum devices and simulators. <https://www.research.ibm.com/ibm-q/technology/>

- devices/. [Online; accessed September 21, 2018].
- [26] D-Wave Systems Inc. Introduction to the D-Wave quantum hardware. 2018.
- [27] Gurobi Optimization. Inc.,gurobi optimizer reference manual, 2015. 2014.
- [28] Catherine C McGeoch. Adiabatic quantum computation and quantum annealing: Theory and practice. *Synthesis Lectures on Quantum Computing*, 5(2):1–93, 2014.
- [29] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000.
- [30] Quantum Enhanced Optimization (QEO). <https://www.iarpa.gov/index.php/research-programs/qeo>. [Online; accessed September 25, 2018].
- [31] Sergey Novikov, Robert Hinkey, Steven Disseler, James I Basham, Tameem Albash, Andrew Risinger, David Ferguson, Daniel A Lidar, and Kenneth M Zick. Exploring more-coherent quantum annealing. *arXiv preprint arXiv:1809.04485*, 2018.
- [32] D. O’Malley, V. V. Vesselinov, B. S. Alexandrov, and L. B. Alexandrov. Nonnegative/binary matrix factorization with a D-Wave quantum annealer. *arXiv preprint arXiv:1704.01605*, 2017.
- [33] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, and B. Parney. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:1–6, 2017.
- [34] 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, and J. Yao. Phase transitions in a programmable quantum spin glass simulator. *Science*, 361:162–165, 2017.
- [35] David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400(1818):97–117, 1985.
- [36] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 124–134. Ieee, 1994.
- [37] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996.
- [38] CJ Ballance, TP Harty, NM Linke, MA Sepiol, and DM Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Physical review letters*, 117(6):060504, 2016.
- [39] Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C White, Josh Mutus, Austin G Fowler, Brooks Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500, 2014.
- [40] Jason R Petta, Alexander Comstock Johnson, Jacob M Taylor, Edward A Laird, Amir Yacoby, Mikhail D Lukin, Charles M Marcus, Micah P Hanson, and Arthur C Gossard. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309(5744):2180–2184, 2005.
- [41] Antonio Acín, Immanuel Bloch, Harry Buhrman, Tommaso Calarco, Christopher Eichler, Jens Eisert, Daniel Esteve, Nicolas Gisin, Steffen J Glaser, Fedor Jelezko, et al. The quantum technologies roadmap: a european community view. *New Journal of Physics*, 20(8):080201, 2018.
- [42] Mark Saffman, Thad G Walker, and Klaus Mølmer. Quantum information with rydberg atoms. *Reviews of Modern Physics*, 82(3):2313, 2010.
- [43] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- [44] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem. *arXiv preprint arXiv:1412.6062*, 2014.
- [45] Edward Farhi and Aram W Harrow. Quantum supremacy through the quantum approximate optimization algorithm. *arXiv preprint arXiv:1602.07674*, 2016.
- [46] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics,

- and function using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, pages 11–15, Pasadena, CA USA, 2008.
- [47] Andrew Cross. The ibm q experience and qiskit open-source quantum computing software. *Bulletin of the American Physical Society*, 2018.
- [48] Tomas Boothby, Andrew D King, and Aidan Roy. Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing*, 15(1):495–508, 2016.
- [49] IBM QISKit Aqua: Variational forms. https://github.com/Qiskit/aqua/blob/master/qiskit_aqua/algorithms/components/variational_forms/ryrz.py. [Online; accessed September 21, 2018].
- [50] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed September 21, 2018].
- [51] Michael JD Powell. A direct search optimization method that models the objective and constraint functions by linear interpolation. In *Advances in optimization and numerical analysis*, pages 51–67. Springer, 1994.
- [52] Christopher J. Wood and John A. Smolin. A C++ quantum circuit simulator with realistic noise. <https://github.com/Qiskit/qiskit-terra/tree/master/src/qasm-simulator-cpp>. [Online; accessed September 21, 2018].
- [53] Jérôme Kunegis. Konect: the koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1343–1350. ACM, 2013.
- [54] Pratha Sah, Lisa O Singh, Aaron Clauset, and Shweta Bansal. Exploring community structure in biological networks with random graphs. *BMC bioinformatics*, 15(1):220, 2014.
- [55] Mika Chmielewski, Jason Amini, Kai Hudek, Jungsang Kim, Jonathan Mizrahi, Christopher Monroe, Kenneth Wright, and David Moehring. Cloud-based trapped-ion quantum computing. *Bulletin of the American Physical Society*, 2018.
- [56] Hannes Pichler, Sheng-Tao Wang, Leo Zhou, Soonwon Choi, and Mikhail D Lukin. Quantum optimization for maximum independent set using rydberg atom arrays. *arXiv preprint arXiv:1808.10816*, 2018.