

Noisy intermediate-scale quantum computing algorithm for solving an n -vertex MaxCut problem with $\log(n)$ qubits

Marko J. Rančić 

TotalEnergies, Nano-INNOV Bt. 861, 8, Boulevard Thomas Gobert, 91120 Palaiseau, France



(Received 2 March 2022; accepted 20 December 2022; published 15 February 2023)

Quantum computers are devices, which allow more efficient solutions of problems as compared to their classical counterparts. As the timeline to developing a quantum-error corrected computer is unclear, the quantum computing community has dedicated much attention to developing algorithms for currently available noisy intermediate-scale quantum computers (NISQ). Thus far, within NISQ, optimization problems are one of the most commonly studied and are quite often tackled with the quantum approximate optimization algorithm (QAOA). This algorithm is best known for computing graph partitions with a maximal separation of edges (*MaxCut*), but can easily calculate other problems related to graphs. Here, I present a novel quantum optimization algorithm, which uses exponentially less qubits as compared to the QAOA while requiring a significantly reduced number of quantum operations to solve the *MaxCut* problem. Such an improved performance allowed me to partition graphs with 32 nodes on publicly available 5 qubit gate-based quantum computers without any preprocessing such as division of the graph into smaller subgraphs. These results represent a 40% increase in graph size as compared to state-of-art experiments on gate-based quantum computers such as Google Sycamore. The obtained lower bound is 54.9% on the solution for actual hardware benchmarks and 77.6% on ideal simulators of quantum computers. Furthermore, large-scale optimization problems represented by graphs of a 128 nodes are tackled with simulators of quantum computers, again without any predivision into smaller subproblems and a lower solution bound of 67.9% is achieved. The study presented here paves way to using powerful genetic optimizer in synergy with quantum computers.

DOI: [10.1103/PhysRevResearch.5.L012021](https://doi.org/10.1103/PhysRevResearch.5.L012021)

Introduction. A universal quantum computer has been the holy grail of quantum technology [1]. Such a device would allow more efficient searching through databases [2], prime number factorization [3], and more efficient solutions of systems of linear equations [4], just to name a few. However, universal quantum computers require millions of qubits with quantum error correction implemented and with an implementation timeline, which is difficult to predict [5]. On the other hand, devices with up to 127 noisy qubits are readily available. This steered the scientific community towards exploring potential computational advantages, which such devices could bring. In the rapidly expanding field of noisy intermediate-scale quantum computing (NISQ) [6,7] two algorithms stand out in prospect: the variational quantum eigensolver (VQE) [8,9] and the aforementioned quantum approximate optimization algorithm (QAOA) [10–12]. The VQE is mainly applied to problems in chemistry and material science while the QAOA is best known for computing graph partitions with a maximal separation of edges (*MaxCut*), but can easily calculate other properties of graphs, such as MaxIn-

dependent set and the partition problem, just to name a few. Given a graph $G = (V, E)$ comprising of $|V|$ vertices and $|E|$ nodes the QAOA requires $n = |V|$ qubits and $p(|E| + |V|)$ quantum operations to implement the ansatz for the *MaxCut* problem [11,12]. Here, p is the phenomenological depth parameter.

In this paper a variational *MaxCut* algorithm requiring $n = \lceil \log_2 |V| \rceil$ qubits is introduced, where $\lceil \cdot \rceil$ stands for the ceiling function. For example if $x = 2.1$, $\lceil x \rceil = 3$, if $x = 2.7$, $\lceil x \rceil = 3$. In similarity with all other NISQ algorithms the algorithm presented here iteratively improves a trial solution (ansatz) in a hybrid quantum-classical optimization loop. The ansatz is implemented with at most $2^n - 2n + 5$ single qubit gates and at most $2^n - 2$ two qubit CNOT gates (in total up to $2^{n+1} - 2n + 3$ gates). Exploiting the fact that large graphs can be treated with the algorithm at a low resource overhead I demonstrate the calculation of a *MaxCut* of a graph of 32 nodes on a publicly available device of only 5 qubits. This is a 40% increase in graph size compared to state-of-the-art experiments with QAOA on gate-based quantum computers such as the Google Sycamore [13,14]. The algorithm presented here opens perspective for immediate quantum speedup with contemporary quantum processors, given that the quantum hardware community is still some years away from producing processors with hundreds of qubits required for quantum speedup with QAOA [15]. Furthermore, a graph of 128 nodes is partitioned on contemporary simulators of quantum computers. With this methodology, simulators of

*marko.rancic@totalenergies.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

quantum computers become a powerful tool for graph partitioning, being able to tackle graphs of hundreds of nodes without dividing the problem into smaller subgraphs, such as the work done in [16] or focusing on correlation between pairs of classical independent variables such as work done in Ref. [17].

The study presented here maps a *MaxCut* cost function to multimodal multidimensional cost function, which can be evaluated on quantum processing units (QPUs). In the subsequent step it relies on the powerful class of global optimizers called “genetic optimizers,” which are proven to excel in finding minima of multidimensional multimodal black-box cost functions [18–20].

Methodology. a. Variable reduction. In this subsection a variable reduction technique compatible with the *MaxCut* problem is going to be presented. This is done in order to make the problem amendable to classical optimizers, which cannot easily treat multidimensional data such as for example the surrogate model EGO optimizer [21]. Also, as shown in Fig. S2 within the Supplemental Material (SM) [22], variable reduction yields better results for lower graph densities and lower result variance for all graph densities. Due to the fact that the binary optimization problem of finding a *MaxCut* is NP-hard, a first approach to approximately solve the problem would be to linearly relax the problem. Meaning that instead of assuming that binary optimization variables in the *MaxCut* problem are integers 0 or 1, one assumes that they are continuous variables [0,1]. In the field of semidefinite programming [23] a different, more efficient approach is taken, binary variables are substituted with vectors. Such classical method of approximately solving the *MaxCut* problem is state-of-the-art and has a maximum possible performance guarantee of $\alpha = 0.87856$ as proved by Goemans and Williamson and can be performed in polynomial time [23].

Here an alternative approach compatible with quantum computing is presented. Let me now introduce a continuous, differentiable function of the following form

$$R_f(x, q, m) = \exp(-\exp(2^{m-q} \sin(2^q x + x_0(q, m))))), \quad (1)$$

where $x_0(q, m) = \arcsin[\ln(-\ln(0.5))/2^{m-q}]$ and the integer $m \geq |V|$ and $0 \leq q \leq |V| - 2$. R_f is defined in such a way to be n differentiable and to have a minimum (maximum) at 0(1) and rapidly changing multi-oscillatory behavior in between extremas. It should be noted that this definition is not by any means unique. Assume a graph where $|V| \gg 1$, consequently m is a large number. $R_f(x, 0, m)$ is such a function, which is mostly 0 in the region of $0 \leq x \leq \pi$ and rapidly changes to 1 in the region $\pi < x \leq 2\pi$ (see red line in Fig. 1). I can therefore substitute any binary variable θ_1 taking values 0 or 1 with $\theta_1 \rightarrow R_f(x, 0, m)$. The second function $R_f(x, 1, m) = 0$ for $0 \leq x \leq \pi/2$ and $\pi \leq x \leq 3\pi/2$ and $R_f(x, 1, m) = 1$ for $\pi/2 < x < \pi$ and $3\pi/2 < x < 2\pi$. Therefore, I substitute the second binary variable $\theta_2 \rightarrow R_f(x, 1, m)$ (see green-dashed line in Fig. 1). By substituting all $\theta_1, \dots, \theta_{|V|-1}$ variables with $R_f(x, 0, m), \dots, R_f(x, |V| - 2, m)$ I have mapped the $|V| - 1$ dimensional binary optimization problem to one-dimensional multimodal continual variable optimization problem. The role of $x_0(q, m)$ is to center all R_f 's to 0.5 at $x = 0$, consequently leading to a more regular optimization landscape (avoiding the




FIG. 1. Equation (1) for $q = 0$ red, $q = 1$ dashed green and $q = 2$ dotted black for $m = 4$.

potential splitting of minima around 0.0 and 1.0). It is derived by setting $R_f(0, q, m) = 0.5$ for an arbitrary $x_0(q, m)$.

b. The algorithm. A Laplacian of a graph $L(G)$, where $G = (V, E)$ is a $|V| \times |V|$ matrix with $|V|$ positive terms on the diagonal and $2|E|$ off-diagonal terms. The i th diagonal term of the graph Laplacian corresponds to the number of connections the node i has with remaining nodes in the graph and the i th off-diagonal term of the matrix is the negative weight between the i th and the j th node.

Similarly to a real-valued Hamiltonian in quantum mechanics, the graph Laplacian is symmetric, furthermore it has a spectrum (eigenvalue range) between 0 and its largest eigenvalue. Now, I introduce a partition vector $\mathcal{V}$ of length $|V|$ with i th term equaling $+1$ if i th vertex of G belongs to the first subgraph in the graph partition and i th term -1 if the i th node belongs to the second sub-graph in the graph partition. Then, the number of cuts in the graph bipartition is $N_{\text{cuts}} = \mathcal{V}^T L \mathcal{V} / 4$ [24], and this formula is a central piece of the algorithm presented here. By finding the vector $\mathcal{V}$, which maximizes N_{cuts} a *MaxCut* of the graph is found. Vector $\mathcal{V}$ has $2^{|V|}$ possible values and there is no known algorithm which can exactly find $\mathcal{V}$, which maximizes N_{cuts} with a computational complexity, which is a polynomial of $|V|$.

Now I present the structure of the algorithm for a preselection of r optimization variables $x_1 \dots x_r$.

1. Trivial unconnected vertices are added to the graph so it has a dimension, which is a power of two.

2. The graph Laplacian $L(G)$ is represented as a sum of tensor products of unitary matrices, and denoted as $\mathcal{L}(G)$ in such form. The decomposition is a prerequisite for measuring the expectation value of the graph Laplacian on a QPU (step 5) and also follows the logic of the implementation in IBMQ's Qiskit.

3. If a graph has $|V|$ nodes a Hadamard gate is applied to $\lceil \log_2 |V| \rceil$ qubits. This operation is denoted with H_G .

4. A variational ansatz in a form of diagonal gate is applied

$$U = \text{diag}(\exp(i\pi R_f(x_1, 0, m_r)), \dots, \exp(i\pi R_f(x_1, 2^n/r, m_r)), \dots, \exp(i\pi R_f(x_r, 0, m_r)), \dots, \exp(i\pi R_f(x_r, 2^n/r, m_r)), 1, \dots, 1), \quad (2)$$

where the number of variables r is between $1 \leq r \leq 2^n$ and $r \bmod 2 = 0$, and $m_r \geq 2^n/r + 2$.




FIG. 2. (a) A simple graph where w_{ij} are the off-diagonal elements of a graph Laplacian. (b) The mapping of the graph to a set of coupled energy levels.

5. The number of cuts is calculated as

$$N_{\text{cuts}} = 2^{n-2} \langle 0 | H_G U^\dagger \mathcal{L}(G) U H_G | 0 \rangle. \quad (3)$$

6. Variational parameter x is adjusted with a classical optimizer and steps 3-5 are repeated until a maximum is reached.

The algorithm presented here maps the *MaxCut* problem of a graph $G = (V, E)$ comprising of $|V|$ vertices and $|E|$ edges to a problem of $|V|$ energy levels coupled with $|E|$ coupling terms described by a Hamiltonian $\mathcal{L}(G)$. The weight between the nodes w_{ij} becomes a coupling strength between energy levels i and j . Energy levels i and j are residing at an energy equal to the connectivity of the node i and j respectively.

In Fig. 2, I represent a simple example of a graph with four nodes. For instance the node 1 is connected with two other nodes [Fig. 2(a)]. Therefore, the energy level 1 lies at an energy $E = 2$ in Fig. 2(b). Node 1 is connected with nodes 2 and 3 so the level 1 is coupled with levels 2 and 3 in the energy scheme. Such logic applies for all nodes of any graph. The algorithm searches for a unitary transformation of the Hamiltonian, which maximizes the number of cuts.

c. Circuit depth, computational complexity, and quantum advantage. Given a graph with $|V|$ nodes $n = \lceil \log_2(|V|) \rceil$ qubits are required to implement the algorithm. The multi-control multitarget qubit gate on n qubits required to realize the diagonal gate U in Eq. (3) can be straightforwardly realized with Grey codes [25] or in the context of follow up work [26] at a cost of $(23/48) \times 4^n - (3/2) \times 2^n + 4/3$ CNOT gates. However, exploiting the fact that U is a diagonal gate and following on the works of [26] and [27] it can be realized with $2^n - 2$ CNOT gates. This means that the ansatz is implemented with less than $2|V|$ two qubit CNOT gates, which is in stark contrast with the QAOA ansatz, which requires $p|E|$ two qubit gates, where p is the depth parameter. Given that $|E| \gg |V|$ the algorithm presented here is much more efficient in the number of two qubit gates as compared to even the lowest depth $p = 1$ QAOA.

The algorithm presented here is a heuristic, meaning that its depth is case dependent. However, the quantum implementation of the heuristic can be compared with its classical counterpart for every step of the evaluation. A classical computing variant of Eq. (3) is a vector-matrix-vector multiplication. For a $|V| \times |V|$ matrix the computational complexity of such an evaluation is $O(|V|^2)$.

In the course of this study the graph Laplacian is decomposed into sum of Pauli matrices with PennyLane's *qml.Hermitian* function [28], afterwards a conversion to *Qiskit*

TABLE I. A table summarizing the difference between the approach presented here and QAOA. Complexities are given for one evaluation step and $n = \lceil \log_2 |V| \rceil$.

	QAOA	New alg.	New alg. d -regular gr.
Ansatz	$O(p(E + V))$	$O(V)$	$O(V)$
Algorithm	$O(p(E + V))$	$O(V ^3)$	$O(V d^2(d + \log^* n))$
Qubit nr.	$ V $	n	n
Connectivity	graph inspired	all-in-all	all-in-all

was performed. Mathematically, the Pauli basis decomposition of a general Hermitian matrix H is conducted in the following fashion:

$$H = \frac{1}{4} \sum_{i=1}^{4^n - 2|V|^2} \text{Tr}(P_i H) P_i, \quad (4)$$

where P_i are all possible combinations of tensor products of the Pauli group composed of $\{I, X, Y, Z\}$ acting on n qubits. It should be noted that the process of decomposing a Hermitian matrix into sum of Paulis [see Eq. (4)] is completed in $O(|V|^2)$: up to $|V|^2$ summands, $|V|^3$ for the matrix-matrix product of P_i and H and $|V|$ for computing the trace. For example, for $n = 2$ qubits (up to $|V| = 4$ vertex graph) $4^n = |V|^2 = 16$ possible products exist

$$P_i \in \{II, IX, IY, IZ, XI, XX, XY, XZ, YI, YX, YY, YZ, YI, ZX, ZY, ZZ\}. \quad (5)$$

It should be noted that all tensor products, which are complex in nature (involving odd numbers of Y Pauli matrices per tensor product) yield a zero trace due to the real nature of the graph Laplacian in our study. Furthermore, in our case the graph Laplacian is a real Hermitian matrix and is thus symmetric, so only the decomposition for lower or upper triangular parts could be performed. Finally, such decompositions are easily parallelized on HPC architectures as all computations are independent one of another.

On a quantum computer the computational complexity of evaluating Eq. (3) is $O(|V|^3)$ [one power of $|V|$ coming from the ansatz and up to $|V|^2$ summands yielding $\mathcal{L}(G)$]. This number, although polynomial in $|V|$ can still be quite high for large graphs and thus would require an estimation of the expectation value of the large number of summands on a quantum computer. However, simulating a d -sparse Hamiltonian (d -regular graph) is done in maximally $O(d^2(d + \log^* n))$ queries with the so-called star decomposition of the Hamiltonian [29,30], and efficiently simulating a Hamiltonian is equivalent to calculating an expectation value of a Hamiltonian [31]. So for d -regular graphs every step in the heuristic algorithm executes in $O(|V|d^2(d + \log^* \lceil \log |V| \rceil))$, which is smaller than the classical $O(|V|^2)$ for small d . Further reductions of the number of measurements could be conducted by the command `group_paulis=True` in Pauli expectation class of *qiskit*. This allows grouping of Pauli strings, which commute into same measurements. Also, with larger processors the tasks could be parallelized certain qubits just estimating certain sets of Pauli strings.

In Table I, I summarize the main differences between the algorithm presented here and the QAOA. The diagonal gates




FIG. 3. The number of cuts of a 4-node graph obtained with 2 qubits on a simulator of quantum computers (left) and an actual quantum device (right).

required to preform the algorithm presented here requires an all-in-all connectivity between qubits for optimal performance. On the other hand, QAOA performs best when the qubits are connected in the same way as the nodes of the graph [13].

As of 2021 the second-most powerful supercomputer in the world is Fugaku [32] with 158 967 nodes each having 32 Gb of RAM. In total, this supercomputer can store $158\,967 \times 32$ Gigabytes of data, equaling to it being able to handle 5×10^{15} bytes of data. Commonly one requires 8 bytes to store a real number [33,34], indicating that 0.64×10^{15} double precision numbers can be stored in the RAM memory in such a device. Given that a graph Laplacian is a square matrix, the dimensions of a weighted graph Laplacian, which such a supercomputer could handle is $\sqrt{0.64 \times 10^{15}} \times \sqrt{0.64 \times 10^{15}} = 25.2 \times 10^6 \times 25.2 \times 10^6$. On the other hand side, a 25 qubit device (such as those already

available at IBM and Google) could handle a $2^{25} \times 2^{25} = 33.5 \times 10^6 \times 33.5 \times 10^6$ graph Laplacian. Of course, noise would be a limiting factor in handling such sizes of optimization problems in the pre-error correction era. However, next year is going to be the year in which large scale error mitigation is going to be implemented on IBMQ systems [5]. This will increase the depth of circuits, which could be executed on IBM QPUs to a larger level, which is difficult to estimate at this point. It should also be noted that the way of calculating the number of cuts of a graph on a quantum computer as given by Eq. (3) and can also be applicable to a plethora of algorithms handling different aspects of graph cuts, not only *MaxCut* with the genetic optimizer as done in this study.

Results. In Fig. 3, I compare the output of a simulator with the output of publicly available IBMQ Santiago. Although pure dephasing, shot noise and relaxation may distort the optimization landscape, maxima are clearly noticed although equal local maxima become unequal. The N_{cuts} is estimated for 100 equidistant values of x .

I further present results obtained by benchmarking randomly generated 3-regular graphs of 32 nodes on actual quantum computers and simulators and randomly generated 3-regular graphs of 128 nodes on a simulator of quantum computers (Qiskit). Intensive testing showed that the algorithm performs best when the number of variables is kept at 8 or 16 for graphs of the size 32–128 nodes. Intensive numerical testing also showed that a genetic optimizer is best suited for finding the maximum of the function—not too surprising as genetic optimizer is indeed best used for multimodal cost functions. On top of the genetic algorithm, a number of classical optimizers were tried (COBYLA, Nelder-Mead,




FIG. 4. Randomly generated 3-regular graphs of 32 nodes. Nodes belonging to different partitions are marked in green and orange respectively and the *MaxCut* value is written on top of the graph. Graphs are randomly initialized by Python's networkx package where seed 20 is used for graphs (a)–(c) and seed 30 for graphs (e)–(g). GW stands for Goemans-Williamson.

TABLE II. A table summarizing the results of benchmarking graphs of 128 nodes on a quantum simulator. A random seed in Python is used to generate graph instances. GW stands for the Goemans-Williamson algorithm.

seed	Maxcut (GW)	Maxcut (New alg.)
7	172	133
8	162	137
9	166	134

Basin-Hopping, Particle Swarm, EGO). Further details about the setting of the genetic optimizer can be found in the SM [22].

For the case of 3-regular graphs of 32 nodes variable reduction is preformed so that the optimization landscape has 8 variables. The *MaxCut* is calculated with 5 qubits. The Goemans-Williamson algorithm (GW) is a classical approximate algorithm for the *MaxCut* problem has a performance guarantee of 0.87856 [23]. I define the approximation ratio of the algorithm presented here with respect to the exact solution as

$$0.87856 \text{MaxCut} / \text{MaxCut}_{\text{GW}} \leq r \leq \text{MaxCut} / \text{MaxCut}_{\text{GW}}, \quad (6)$$

where *MaxCut* is the value obtained with the algorithm presented here and *MaxCut*_{GW} is the value obtained with Goemans-Williamson. For the first set of graphs in Figs. 4(a)–4(c) the algorithm performed on a simulator of quantum computers yields $0.776 \leq r \leq 0.884$ and the algorithm executed on IBMQ Quito $0.552 \leq r \leq 0.628$. For the second set of graphs in Figs. 4(d)–4(f) the algorithm preformed on a simulator of quantum computers yields $0.857 \leq r \leq 0.975$ and the algorithm executed on IBMQ Quito $0.549 \leq r \leq 0.625$. For both realizations there is a clear difference between actual hardware benchmarks and ideal simulation. I assume that the main reason for this is the distortion of the optimization landscape due to pure dephasing and relaxation. I expect that shot noise contributed less as the expectation value was estimated for the maximally allowed 8192 shots.

For the case of 3-regular graphs of 128 nodes variable reduction is performed so that the optimization landscape has 16 variables. The *MaxCut* is calculated with seven simulated qubits under the assumption of no noise processes. Given that devices larger than five qubits are unavailable to the author,

for these three graphs I stayed in the domain of quantum simulators. For the first graph Table II seed 7: $0.679 \leq r \leq 0.773$, second graph Table II seed 8: $0.743 \leq r \leq 0.846$, third graph Table II seed 9: $0.709 \leq r \leq 0.807$. Values do not converge as nicely as for smaller graphs likely because the genetic algorithm gets trapped in a local minimum with increasing system size. These results are visually represented in Fig. S3 in the SM [22].

A d -regular graph with $|V|$ nodes has $d \times |V|/2$ edges [35]. An average random bipartition of such a graph is $d \times |V|/4$ [36], or in the case of 3-regular graphs with 32 nodes the average random bipartition is 24. So both the quantum hardware results and the simulator results stay above the average random bipartition value. An average random bipartition of a 3-regular graph of 128 nodes is $3 \times 128/4 = 96$.

Conclusions. In conclusion I have presented a novel algorithm for noisy intermediate-scale quantum computers requiring logarithmically less qubits and significantly less quantum gates as compared to the contemporary state-of-the-art algorithm—QAOA. I went through to calculate the *MaxCut* of a randomly generated 3-regular graph of 32 nodes, a 40% increase compared to experiments of state-of-the-art gate-based quantum computers (Google Sycamore). I did so with publicly available IBM hardware, and obtained a lower bound of 54.9% on the solution for actual hardware benchmarks and 77.6% on ideal simulators of quantum computers. Furthermore, I calculated the *MaxCut* of a 3-regular graph of 128 nodes with quantum simulator obtaining a lower bound of 67.9% on the solution and with no pre-processing of the graph what-so-ever.

Output from quantum computers and codes to support the claim are available at zenodo [38].

Acknowledgments. I acknowledge discussions with Adrien Suau, Vedran Dunjko, Thomas Bäck, and Zaid Allybokus, with the latter especially in the domain of use cases to which the algorithm may be applied. I am thankful to the EU for funding within the H2020 project (NE|AS|QC). M.J.R. envisioned, developed, and conducted the whole study himself. The work presented here is a part of a broader provisional patent claim “Method for optimizing a functioning relative to a set of elements and associated computer program product” submission number EP21306155.9 submitted on 26.8.2021. with Marko J. Rančić and Zaid Allybokus being listed as inventors.

The author declares no further competing interests.

- [1] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, 1996), pp. 212–219.
- [3] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, NW Washington, DC, 1994), pp. 124–134.
- [4] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [5] IBM quantum roadmap, *Research IBM com*, <https://www.ibm.com/quantum/roadmap>.
- [6] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [7] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).

- [8] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [9] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth *et al.*, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [10] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028*.
- [11] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [12] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, *Phys. Rev. X* **10**, 021067 (2020).
- [13] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [14] The reader should be made aware of the state-of-the-art QAOA experiments with 40 trapped ion qubits [37]. However, here the focus was on finding the ground state of the linear Ising model rather than optimizing a partition of a realistic graph.
- [15] G. G. Guerreschi and A. Y. Matsuura, Qaoa for max-cut requires hundreds of qubits for quantum speed-up, *Sci. Rep.* **9**, 6903 (2019).
- [16] J. Li, M. Alam, and S. Ghosh, Large-scale quantum approximate optimization via divide-and-conquer, *arXiv:2102.13288*.
- [17] B. Tan, M.-A. Lemonde, S. Thanasilp, J. Tangpanitanon, and D. G. Angelakis, Qubit-efficient encoding schemes for binary optimisation problems, *Quantum* **5**, 454 (2021).
- [18] F. Hoffmeister and T. Bäck, Genetic algorithms and evolution strategies: Similarities and differences, In *International Conference on Parallel Problem Solving from Nature* (Springer, New York, 1990), pp. 455–469.
- [19] N. N. Glibovets and N. M. Gulayeva, A review of niching genetic algorithms for multimodal function optimization, *Cybern. Syst. Anal.* **49**, 815 (2013).
- [20] N. Casas, Genetic algorithms for multimodal optimization: A review, *arXiv:1508.05342*.
- [21] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. R. A. Martins, A python surrogate modeling framework with derivatives, *Adv. Eng. Softw.* **135**, 102662 (2019).
- [22] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.5.L012021> for the optimizer settings and larger graph results.
- [23] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [24] A. Pothén, H. D. Simon, and K.-P. Liou, Partitioning sparse matrices with eigenvectors of graphs, *SIAM J. Matrix Anal. Appl.* **11**, 430 (1990).
- [25] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Quantum Circuits for General Multiqubit Gates, *Phys. Rev. Lett.* **93**, 130502 (2004).
- [26] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum-logic circuits, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **25**, 1000 (2006).
- [27] E. Malvetti, R. Iten, and R. Colbeck, Quantum circuits for sparse isometries, *Quantum* **5**, 412 (2021).
- [28] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, *arXiv:1811.04968*.
- [29] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient quantum algorithms for simulating sparse Hamiltonians, *Commun. Math. Phys.* **270**, 359 (2007).
- [30] A. M. Childs and R. Kothari, Simulating sparse Hamiltonians with star decompositions, in *Conference on Quantum Computation, Communication, and Cryptography* (Springer, New York, 2010), pp. 94–103.
- [31] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [32] The top 500 supercomputers list, <https://www.top500.org/lists/top500/2022/11/>, 2022.
- [33] D. Goldberg, What every computer scientist should know about floating-point arithmetic, *ACM Comput. Surv.* **23**, 5 (1991).
- [34] M. L. Overton, *Numerical Computing with IEEE Floating Point Arithmetic* (SIAM, Philadelphia, 2001).
- [35] J. M. Aldous and R. J. Wilson, *Graphs and Applications: An Introductory Approach* (Springer Science & Business Media, New York, 2003).
- [36] R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge University Press, Cambridge, 1995).
- [37] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin *et al.*, Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 25396 (2020).
- [38] M. J. Rancic, An exponentially more efficient optimization algorithm for noisy quantum computers, zenodo (2021), doi: 10.5281/zenodo.5592834.