Light Cone Cancellation for Variational Quantum Eigensolver Ansatz

Xinjian Yan

Graduate School of Science and Technology University of Tsukuba Tsukuba, Ibaraki, Japan yanxinjian@cavelab.cs.tsukuba.ac.jp

Ningyi Xie Graduate School of Science and Technology University of Tsukuba Tsukuba, Ibaraki, Japan nyxie@cavelab.cs.tsukuba.ac.jp

Leo Kurosawa Graduate School of Computer Science and Engineering University of Aizu Aizu-Wakamatsu, Fukushima, Japan m5271018@u-aizu.ac.jp

Dongsheng Cai Faculty of Engineering, Information and Systems University of Tsukuba Tsukuba, Ibaraki, Japan cai@cs.tsukuba.ac.jp Xinwei Lee

School of Computing and Information Systems Singapore Management University 80 Stamford Rd, Singapore xwlee@smu.edu.sg

Yoshiyuki Saito Graduate School of Computer Science and Engineering University of Aizu Aizu-Wakamatsu, Fukushima, Japan d8241104@u-aizu.ac.jp

Nobuyoshi Asai School of Computer Science and Engineering University of Aizu Aizu-Wakamatsu, Fukushima, Japan nasai@u-aizu.ac.jp

HoongChuin Lau

School of Computing and Information Systems Singapore Management University 80 Stamford Rd, Singapore hclau@smu.edu.sg

Abstract—Variational Quantum Algorithms (VQAs) represent a class of algorithms that utilize a hybrid approach, combining classical and quantum computing techniques. In this approach, classical computers serve as optimizers that update circuit parameters to find approximate solutions to complex problems. In this study, we apply a method known as Light Cone Cancellation (LCC) to optimize variational circuits, effectively reducing the required number of qubits and gates for circuit simulation. We then evaluate the performance of LCC one of the VOAsthe Variational Ouantum Eigensolver (VOE)-to address the Max-Cut problem. Compared with the Quantum Approximate Optimization Algorithm (QAOA), VQE offers greater degrees of freedom at lower circuit depths. By applying LCC to VQE, we can shift the complexity of circuit simulation from the number of qubits to the number of edges in the graph, i.e., from exponential time to polynomial time. This enables us to solve large problems up to 50 vertices, without actually simulating the entire circuit. From our simulation in a 7-qubit and a 27-qubit noisy devices, we show that LCC yields higher approximation ratios than those cases without LCC, implying that the effect of noise is reduced when LCC is applied.

Index Terms—Variational Quantum Algorithms, Variational Quantum Eigensolver, Light Cone Cancellation, NISQ era

I. INTRODUCTION

Many combinatorial optimization problems (COP) are considered to be difficult to address using traditional computational approaches. COPs aim to find the optimal combination of variables that minimizes (or maximizes) a given objective function, while simultaneously satisfying a set of constraints. Recent years, people focus on using quantumclassical hybrid methods, known as the variational quantum algorithms (VQA) [1] to heuristically solve COPs. The quantum approximate optimization algorithm (QAOA) [2] is one of the VQAs that is intensively explored due to its predictable patterns in the variational parameters [3]-[5], and also its relation with quantum annealing [6], [7]. Another VOA, the variational quantum eigensolver (VQE) [8], is also capable of solving COPs, while it is more known for its application to quantum chemistry. Unlike QAOA which has a problemdependent ansatz, the structure of the VQE ansatz is static and does not depend on the problem solved. Moreover, the VOE ansatz offers a greater degree of freedom in the sense that it has greater expressibility [9] and more number of variational parameters compared to the QAOA ansatz. In our previous work, we have shown that VQE generally has better performance than QAOA and the Multi-angle QAOA [10] in solving the Max-Cut problem using noiseless simulations.

The VQAs have shown the potential quantum advantage on Noisy Intermediate Scale Quantum (NISQ) [11] devices. However, an increase in the number of qubits often leads to higher error rates when building actual quantum hardware. Although recent advancements have prominently featured quantum error correction algorithms, it involve intricate designs that must effectively address the inherent noise and decoherence in quantum environments. It seems that reducing the number of qubits and the number of gates is a more feasible and efficient approach while maintaining the accuracy of the algorithm.

In this paper, we apply a method known as Light Cone Cancellation (LCC) [12], [13] to solve the Max-Cut problem using VQE. When computing the expectation function of variational circuits, there are many redundant operators that need not be included in the computations. LCC exploits the preliminary knowledge of which operators are redundant, so that we do not include them in the calculation at the first place. LCC is widely applied for circuits like QAOA, and also inspired the tensor network for QAOA [14]-[16]. The contribution of our work is that we exploit LCC on the two-local ansatz of VQE. Particularly, we design the subcircuits for LCC-VQE and analyze its performance, supported by experimental evidences with noisy devices. To verify the impact of LCC for VQE in a noisy environment, we conduct experiments using simulated noisy backends with 7 qubits and 27 qubits. Experiments show that compared with original VQE, the implementation of LCC has better performance.

The rest of the paper is structured as follows. Section II provides background information and the construction architecture of LCC. The detailed results of experiments are articulated in Section III. Section IV contains the concluding remarks of this study.

II. BACKGROUND

A. Max-Cut

Max-Cut is a fundamental and widely studied NP-complete problem in the field of graph theory. The primary objective of Max-Cut is to partition the nodes of an undirected graph into two disjoint subsets such that the number of edges connecting the two subsets is maximized. This problem finds relevance in diverse fields, including network design, VLSI layout, community detection, and social network analysis [17].

Consider an *n*-node undirected graph G = (V, E), where |V| = n. V represents the nodes and E represents the edges of graph G. A cut is defined as a partition of the original set V into two subsets. The cost function $C(\mathbf{x})$ to be maximized is the sum of the edges connecting points in the two different subsets, which can be expressed as:

$$C(\mathbf{x}) = \sum_{(i,j)\in E} (x_i \oplus x_j), \tag{1}$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and x_i represent the binary variable of node *i*. The symbol \oplus denotes the XOR (exclusive-



Fig. 1: Two-local ansatz used in the experiments in our work. The architecture of single qubit R_y gates with circular CZ entanglement and one repetition is used. The figure shows a 4-qubit example. The parameters $\theta_{k,l}$ are based on the notation specified in Eq. (4).

OR) operation. We want to find the combinations of x such that the cost function is maximized, which also means that make the number of edges cut is maximum. A brute-force approach in a classical computer will require an $\mathcal{O}(2^n)$ time to solve this problem.

In the quantum realm, the cost function in Eq. (1) can be formulated as the cost Hamiltonian H_C , in which its expectation is to be maximized:

$$H_C = \frac{1}{2} \sum_{(i,j)\in E} (I - Z_i Z_j),$$
(2)

where Z_i represents Pauli-Z operator acting on qubit *i*.

B. Variational Quantum Eigensolver

The variational quantum eigensolver (VQE) is initially developed for calculating the minimum energy states of molecules. When re-formulated to address the Max-Cut problem, the expectation value of the cost Hamiltonian H_C over a trial state $|\psi(\theta)\rangle$ is defined as

$$E(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | H_C | \psi(\boldsymbol{\theta}) \rangle.$$
(3)

The objective is to maximize $E(\theta)$, which is equivalent to minimizing $-E(\theta)$, using a classical optimizer. θ is the collection of variational parameters for the VQE ansatz circuit. In this paper, we employ a two-local circuit with $R_y(\theta)$ single-qubit rotation gates and CZ (controlled-Z) circular entanglement, with only one repetition. Fig. 1 shows an example of a 4qubit ansatz circuit. The circular entanglement has CZ gates between adjacent qubits, and also between the first and the last qubit. Instead of using a flattened array θ , we use a matrix Θ to represent the variational parameters to suit the geometrical position of those parameters in the circuit:

$$\Theta = \begin{bmatrix} \theta_{1,1} & \theta_{1,2} \\ \theta_{2,1} & \theta_{2,2} \\ \vdots & \vdots \\ \theta_{n,1} & \theta_{n,2} \end{bmatrix}.$$
 (4)

It is important to note that Θ is simply the re-shaped version of θ , and the notation $\theta_{k,l}$ denotes the *l*-th parameter for the



Fig. 2: (a) The light cone cancellation (LCC) in a two-local ansatz with repetition of 1. The expectation function $\langle \psi(\theta) | Z_i Z_j | \psi(\theta) \rangle$ is visualized as a quantum circuit on the left figure. R_y 's are the single-qubit gates and C_z 's are the two-qubit gates. The gates on the left of the red dashed line show $|\psi(\theta)\rangle$ and the gates on the right show $\langle \psi(\theta) |$. The blue shaded regions show the redundant gates that can be cancelled during the calculation of the expectation. The figure on the right shows the resulting circuit after cancellation. (b), (c) and (d) shows the possible resulting circuit for the Max-cut Hamiltonian, depending on the indices *i* and *j* (positions of the Pauli-Z operators). (b) When Z_i and Z_j are adjacent; (c) when Z_i and Z_j are one qubit apart; and (d) when Z_i and Z_j are two or more qubits apart.

k-th qubit in the original circuit. Also, due to the periodicity of the R_y rotation gate, the range of the values for $\theta_{k,l}$ are restricted to $[0, 2\pi)$.

Numerous metrics are available for assessing the performance of VQE. Within the scope of unconstrained COP, our focus is on studying the approximation ratio (AR). This metric compares the expected solutions obtained through VQE with that of the optimal solutions, essentially measuring the efficiency of VQE in approximating the best possible outcome. It is defined as:

$$AR = \frac{E(\boldsymbol{\theta}^*)}{MaxCut(G)},$$
(5)

where θ^* is the quasi-optimal parameters returned by the optimizer, $E(\theta^*)$ is its corresponding expectation, and MaxCut(G) is the real solution of graph G. The closer the value of AR is to 1, the closer it is to the true solution of Max-Cut.

The performance of noiseless VQE and QAOA in addressing Max-Cut problems is evaluated in a previous research. When both algorithms are initialized with the same number of parameters, our findings indicate that VQE outperforms QAOA using random initialization. Furthermore, a comparison with Multi-angle quantum approximate optimization algorithm (ma-QAOA) [10] reveals that VQE also achieves superior performance in different undirected graphs [18]. The reason for this is that at the lower circuit depth, the degree of freedom in VQE is greater than that in ma-QAOA, which in turn is greater than that in QAOA.

C. Light Cone Cancellation

Light cone cancellation (LCC) is a method that utilizes the intrinsic property of the expectation function, so that the redundant operators in the expectation function are not included in the computation of the expectation function in the first place [19]. The LCC property was originally used in QAOA in reducing the problem graphs into their constituent subgraphs, hence simplifying the problem to be solved [20]– [22]. It also inspired applications like the tensor network of QAOA [14]–[16]. Referring to Eq. (2), for a Hamiltonian H_C that can be written as the sum of local operators, the trial wavefunction (or ansatz) of both sides, $\langle \psi(\theta) |$ and $|\psi(\theta) \rangle$, can be cancelled out partially. This is because some of the unitary gates used to prepare $|\psi(\theta)\rangle$ commutes through the central local operators. Fig. 2(a) shows an example of the LCC of



Fig. 3: Comparison of the approximation ratio (AR) for the VQE solved with a 7-qubit device FakeCascablanca (with LCC) and a 27-qubit device FakeParis (without LCC). Each point shows the best AR (highest) chosen out of 24 trials. The lines are linear fits of their respective data.

a two-local circuit used to prepare the trial wavefunction for VQE. The figure visualizes the expectation function Eq. (3) as a quantum circuit. The circuit on the left of the red dashed line shows the term $|\psi(\theta)\rangle$, and the circuit on the right of the dashed line shows $\langle \psi(\theta) |$, with the central operator $Z_i Z_j$ (between the red dashed lines) acting on qubit *i* and *j*. Since $|\psi(\theta)\rangle$ is just the conjugate transpose of $\langle \psi(\theta) |$, they are the counterpart of each other in the circuit. The blue shades show the gates that are not related to qubits *i* and *j* can commute through the center are cancelled. The result of the cancellation is shown on the right of Fig. 2(a), which has reduced the number of qubits and number of gates.

The Max-Cut Hamiltonian in Eq. (2) has a two-local Z operator on every term. After LCC on the two-local ansatz that we considered (single layer R_u and circular CZ entanglement), we can get 3 different types of sub-circuits as shown in Fig. 2(b), (c) and (d), which requires a 4-, 5-, and 6-qubit quantum circuit, respectively. Circular entanglement means the adjacent qubits, as well as the first and the last qubits, are entangled. Note that one sub-circuit corresponds to the expectation of a local term in the Hamiltonian, so the simulation of the sub-circuits can be done separately. Also, the sub-circuit in Fig. 2(d) can be further divided into two separate circuits as the first 3 qubits and the last 3 qubits are not entangled. Thus, we only require a maximum of 5 qubits to simulate the expectation of the entire Max-Cut Hamiltonian, regardless of the problem size. In fact, only at most 2k + 1 qubits are required to simulate the expectation of a k-local Hamiltonian for this kind of ansatz (one layer of single-qubit gates and circular entanglement). The architecture of the entanglement is crucial in deciding how many qubits we can reduce. For linear entanglement (adjacent qubits entangled, first and last not entangled), there would be another case of the sub-circuits where the operator is located on the first qubit or on the last qubit. For full entanglement (all qubits entangled), LCC will not be possible.

Since now the number of qubits required stays constant, the complexity of circuit simulation now shifts to the number of sub-circuits, i.e. the number of terms that is in the Hamiltonian. For Max-Cut, the number of terms scales with the number of edges in the graphs, which is at most n(n-1)/2 edges (a complete graph). Without LCC, the simulation complexity would be dominated by the number of qubits n, which takes $\mathcal{O}(2^n)$ space and time when simulated classically. With LCC, the time complexity now depends on the number of edges of the graph, which is $\mathcal{O}(n^2)$. The shift in complexity of quantum circuit simulation, from exponential to polynomial, is a significant advantage for LCC.

Another advantage for LCC is that the number of gates is reduced, which in turn reduces the effect of gate noises in quantum devices. On the other hand, it is also worth noting that even though the number of qubits and the number of gates are reduced, the number of parameters stays unchanged after LCC. This is because the sub-circuits after LCC will have different parameters corresponding to the indices *i* and *j*, depending on where the operators $Z_i Z_j$ are. Therefore, the difficulty in optimization remains the same before and after LCC.

LCC is also applicable for other circuits like QAOA and ma-QAOA. However, QAOA or ma-QAOA usually needs more repetitions (larger circuit depths) to achieve higher ARs. Meanwhile, VQE with even one repetition is enough to reach most of the states and hence easier to reach higher ARs than QAOA. This is due to the difference in expressibility between the two-local ansatz and the QAOA ansatz [9]. When the repetition increases, the number of gates (and the number of qubits) that can be cancelled decreases, resulting in a larger sub-circuit after LCC. Since VQE yields higher ARs with fewer repetitions than QAOA, LCC-VQE can be done with fewer qubits, compared to LCC-QAOA or LCC-ma-QAOA.

III. RESULTS

We solve the Max-Cut problem using VQE with a twolocal ansatz (single layer R_u and CZ entanglement), and we use the COBYLA optimizer [23] for all the experiments. The experiments are to compare the performances of the VQE with LCC and that without LCC. We measure the performance using the approximation ratio (AR), which defines how near the result given by VQE is to the actual solution. To make sure that the optimizer does not converge to a good minimum by chance, we perform 24 trials of random initial parameters for every instance, i.e., random initialization. The experiments are performed under noisy conditions so that we can observe the effect of the reduction in the number of qubits and the number of gates on the amount of noise in the circuit. We use two different fake noisy devices provided by Qiskit [24]: FakeCascablanca (7 qubits) and FakeParis (27 qubits). The fake devices simulate the same noise settings in their respective real quantum devices.

Figure 3 shows the comparison between Max-Cut instances solved with LCC on the 7-qubit device, and those solved without LCC on the 27-qubit device. As only 5 qubits are required to simulate the sub-circuits for LCC, the VQE simulation can be run on a device with 7 qubits. On the other hand, simulation with full number of qubits is required for those without LCC. In both settings, we solve the Max-Cut for 36 non-isomorphic instances, ranging from number of vertices n = 10 to n = 15. Additonally, 20 non-isomorphic instances are solved on the 7-qubit device (with LCC) for n = 20, 30, 40 and 50. The dataset for the experiments is shown in the Appendix. Each point in the plot represents the best AR out of 24 trials for a single problem instance. The red points plot the ARs for the instances solved with LCC on the 7-qubit device; the blue points plot the ARs for the instances solved without LCC on the 27-qubit device. The red line is a linear fit through the red points (with LCC) for n = 10 to n = 50. The blue line linearly fits through the blue points (without LCC) for n = 10to n = 15. The green line is a fit for the red points (with LCC) from n = 10 to n = 15.

There are a few observations worth noting. LCC enabled the simulation of large problems up until n = 50, only with quantum circuits with at most 5 qubits. From the red and blue fitted lines, we can see that the ARs for problems with LCC are generally higher than those without LCC. Although with different environments (7-qubit device and 27-qubit device), the error rate is generally lower on a smaller device, so the AR is not so much deteriorated on a 7-qubit device. Moreover, with less number of gates, the effect of noise on the AR is also reduced. It can also be observed that the AR decreases as the problem size (number of vertices of the graph) increases. Another interesting point to note is that the green line has similar slope as the red line, which means the decreasing trend





Fig. 4: Comparison of the AR for the VQE with LCC and without LCC, using the same FakeParis backend (27 qubits). The diagonal dashed line shows where the AR of both methods are equal. The points in the lower triangle have higher AR with LCC. The points in the upper triangle have higher AR without LCC.

is similar for microscopic (n = 10 to n = 15) and macroscopic (n = 10 to n = 50) number of nodes. Also, we can observe from the blue line that if without LCC, the AR decreases faster due to more noises in the circuits. This implies the possibility that in the case without LCC, the AR would decrease faster than the case with LCC, if the blue line is extrapolated to larger problem size.

Figure 4 shows the comparison for n = 10 to n = 15problems solved with and without LCC, using the same 27qubit backend. The 36 instances used are the same as those in Fig. 3. The diagonal dashed line shows where both methods (with and without LCC) have the same AR. The points in the lower triangle are those which LCC gives higher ARs than without having LCC; while the points in the upper triangle are those which LCC gives lower ARs. Under the same noise conditions, 27 out of the 36 instances gives higher ARs with LCC applied, while the others have higher ARs for without LCC. This result shows the effect of the reduction in the number of gates in a more evident way than the experiment shown in Fig. 3, as the number of qubits and the error rates are the same for both with and without LCC. It is also observed that problems with larger size n benefit more from LCC as their ARs stay in the lower triangle, and those with smaller sizes stay in the upper triangle. This is because larger circuits generally have more noise, causing their ARs to deteriorate more.

IV. CONCLUSION

We studied the effect of LCC on VQE in solving the Max-Cut problem. Our work opens up the possibility of using VQE to solve combinatorial optimization problems, as VQE requires lower repetition than QAOA to achieve the same performance. This lower repetition of VQE allows us to cancel a larger number of qubits when applying LCC, thereby shifting the complexity of circuit simulation from exponential scaling (number of qubits) to polynomial scaling (number of edges in a graph). For the Max-Cut problem with a two-local cost Hamiltonian, only at most five qubits are required to solve the problem of any size. We compare the performance of circuits with and without LCC on a noisy simulator provided by Qiskit. The results show that the circuits with LCC generally yield higher approximation ratios than the circuits without LCC, hence implying that the noise is being mitigated.

APPENDIX

DATASET OF THE EXPERIMENTS

Table I shows the experiment datasets we conducted in this study, we use two categories of unweighted, undirected graphs: the G(n,p) Erdös-Rényi graphs and the regular graphs. The graphs are generated using the NetworkX Python package. The G(n,p) graphs are generated using fast_gnp_random_graph(); the regular graphs are generated using random_regular_graph(), with the seeds specified. Each seed represents one graph instance.

TABLE I: The datasets used in the experiments—the regular graphs with degree d and G(n,p) graphs with edge probability p.

No. of nodes	Graph type	d (reg.) or p (ER)	Seed
10	ER	0.5	0, 1, 2, 3
10	Reg.	3	0, 1, 2, 3
11	ER	0.5	0, 1, 2, 3
11	Reg.	4	0, 1, 2, 3
12	ER	0.5	0, 1, 2, 3
12	Reg.	3	0, 1, 2, 3
13	ER	0.5	0, 1
13	Reg.	2	1, 3
14	ER	0.4	0, 1
14	Reg.	2	2, 5
15	ER	0.3	1, 2
15	Reg.	2	3, 7
20	ER	0.25	0, 3, 5, 17
20	Reg.	3	0, 1, 2, 3
30	ER	0.12	0, 5
30	Reg.	2	3, 8
40	ER	0.06	106, 125
40	Reg.	2	0, 1
50	ER	0.06	126,167,424,561

REFERENCES

- [1] 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," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625– 644, 2021.
- [2] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
 [Online]. Available: https://arxiv.org/abs/1411.4028

- [3] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, "Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices," *Phys. Rev. X*, vol. 10, p. 021067, Jun 2020. [Online]. Available: https://link.aps.org/doi/10.1103/ PhysRevX.10.021067
- [4] J. Cook, S. Eidenbenz, and A. Bärtschi, "The quantum alternating operator ansatz on maximum k-vertex cover," 2020 IEEE International Conference on Quantum Computing and Engineering (QCE), pp. 83–92, 2020.
- [5] X. Lee, N. Xie, D. Cai, Y. Saito, and N. Asai, "A depth-progressive initialization strategy for quantum approximate optimization algorithm," *Mathematics*, vol. 11, no. 9, p. 2176, May 2023. [Online]. Available: http://dx.doi.org/10.3390/math11092176
- [6] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, "Quantum computation by adiabatic evolution," arXiv preprint arXiv:quantph/0001106, 2000. [Online]. Available: https://arxiv.org/abs/quant-ph/ 0001106
- [7] S. H. Sack and M. Serbyn, "Quantum annealing initialization of the quantum approximate optimization algorithm," *Quantum*, vol. 5, p. 491, Jul. 2021. [Online]. Available: http://dx.doi.org/10.22331/ q-2021-07-01-491
- [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," *Nature communications*, vol. 5, no. 1, p. 4213, 2014.
- [9] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies*, vol. 2, no. 12, Oct. 2019. [Online]. Available: http://dx.doi.org/10.1002/qute. 201900070
- [10] R. Herrman, P. C. Lotshaw, J. Ostrowski, T. S. Humble, and G. Siopsis, "Multi-angle quantum approximate optimization algorithm," 2021.
- [11] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018. [Online]. Available: http: //dx.doi.org/10.22331/q-2018-08-06-79
- [12] A. Lowe, M. H. Gordon, P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, "Unified approach to data-driven quantum error mitigation," *Physical Review Research*, vol. 3, no. 3, Jul. 2021. [Online]. Available: http://dx.doi.org/10.1103/PhysRevResearch.3.033098
- [13] L. Leone, S. F. E. Oliviero, L. Cincio, and M. Cerezo, "On the practical usefulness of the hardware efficient ansatz," 2022. [Online]. Available: https://arxiv.org/abs/2211.01477
- [14] C. Huang, F. Zhang, M. Newman, X. Ni, D. Ding, J. Cai, X. Gao, T. Wang, F. Wu, G. Zhang, H.-S. Ku, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, M. Szegedy, Y. Shi, H.-H. Zhao, C. Deng, and J. Chen, "Efficient parallelization of tensor network contraction for simulating quantum computation," *Nature Computational Science*, vol. 1, no. 9, pp. 578–587, Sep 2021. [Online]. Available: https://doi.org/10.1038/s43588-021-00119-7
- [15] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, "Tensor network quantum simulator with step-dependent parallelization," in 2022 IEEE International Conference on Quantum Computing and Engineering (QCE), 2022, pp. 582–593.
- [16] D. Lykov and Y. Alexeev, "Importance of diagonal gates in tensor network simulations," in 2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI), 2021, pp. 447–452.
- [17] A. Jagannath, J. Ko, and S. Sen, "Max κ-cut and the inhomogeneous potts spin glass," *The Annals of Applied Probability*, vol. 28, no. 3, pp. 1536–1572, 2018.
- [18] X. Yan, X. Lee, D. Cai, and N. Asai, "Comparison between the performances of general two-local ansatzes and qaoa in max-cut problem," *IPSJ SIG Technical Report*, vol. 2024-QS-11, no. 16, pp. 1–8, Mar. 2024. [Online]. Available: https://ipsj.ixsq.nii.ac.jp/ej/ ?action=pages_view_main&active_action=repository_view_main_item_ detail&item_id=233690&item_no=1&page_id=13&block_id=8
- [19] M. Benedetti, M. Fiorentini, and M. Lubasch, "Hardware-efficient variational quantum algorithms for time evolution," *Physical Review Research*, vol. 3, no. 3, p. 033083, 2021.
- [20] F. G. S. L. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, "For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances," *arXiv preprint arXiv:1812.04170*, 2018. [Online]. Available: https://arxiv.org/abs/1812.04170

- [21] J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, "The quantum approximate optimization algorithm at high depth for maxcut on large-girth regular graphs and the sherrington-kirkpatrick model." Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [Online]. Available: https://drops.dagstuhl.de/opus/volltexte/2022/16514/
- Available: https://drops.dagstuhl.de/opus/volltexte/2022/16514/
 [22] J. Wurtz and D. Lykov, "Fixed-angle conjectures for the quantum approximate optimization algorithm on regular maxcut graphs," *Phys. Rev. A*, vol. 104, p. 052419, Nov 2021. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.104.052419
- [23] M. J. D. Powell, A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation, 1994, pp. 51–67. [Online]. Available: https://app.dimensions.ai/details/ publication/pub.1046127469
- [24] Qiskit contributors, "Qiskit: An open-source framework for quantum computing," 2023.