



Efficient QAOA Optimization using Directed Restarts and Graph Lookup

Meng Wang

mengwang@ece.ubc.ca

The University of British Columbia
Pacific Northwest National Laboratory
Canada, USA

Ang Li

ang.li@pnnl.gov

Pacific Northwest National Laboratory
USA

Bo Fang

bo.fang@pnnl.gov

Pacific Northwest National Laboratory
USA

Prashant Nair

prashantnair@ece.ubc.ca

The University of British Columbia
Canada

ABSTRACT

Variational Quantum Algorithms (VQA) aim to enhance the capabilities of Noisy Intermediate-Scale Quantum (NISQ) devices. These algorithms utilize parameterized circuits and classical optimizers to iteratively execute circuits with varying parameters. However, VQA faces computational overheads due to repeated iterations and random restarts. Prior work suggests using basic sub-graphs to transfer parameters for the input graph, reducing optimizer overheads but limiting applicability to structured regular graphs. In real-world applications, random irregular graphs are common, and existing methods are not scalable or practical for such graphs. This paper presents a framework that aims to improve random irregular graphs in VQA. The framework uses graph similarity and important features like total edge counts, average edge counts, and variance. It follows an iterative process to choose basis sub-graphs from a small database and adjust parameters accordingly. Classical optimizers then utilize these parameters to determine when to restart and perform gradient descent. This approach increases the chances of reaching global maximum points.

CCS CONCEPTS

• **Computer systems organization** → **Quantum computing**.

KEYWORDS

QAOA parameter optimization, QAOA parameter transfer, parameterized quantum circuits, quantum variational algorithms

ACM Reference Format:

Meng Wang, Bo Fang, Ang Li, and Prashant Nair. 2023. Efficient QAOA Optimization using Directed Restarts and Graph Lookup. In *Proceedings of the 2023 International Workshop on Quantum Classical Cooperative Computing (QCCC '23)*, June 20, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3588983.3596680>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
QCCC '23, June 20, 2023, Orlando, FL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0162-7/23/06.
<https://doi.org/10.1145/3588983.3596680>

1 MOTIVATION

Quantum computing has shown great potential in solving complex problems that are beyond the capabilities of classical computers [1, 6, 15, 17]. In the current era of Noisy Intermediate-Scale Quantum (NISQ) computers, Variational Quantum Algorithms (VQAs) [3] has emerged as a promising approach to tackle optimization problems using quantum computing [10]. Among the VQAs, Quantum Approximate Optimization Algorithm (QAOA) is one of the most widely studied and promising algorithms for solving combinatorial optimization problems [4, 7–9, 14, 18].

However, one of the major challenges in implementing QAOA is finding the optimal parameters that can provide the best possible solution for a given optimization problem. The challenge comes from the impact of noise and execution overheads in the NISQ era. Due to these limitations, it becomes difficult to execute quantum circuits with high precision and efficiency.

2 LIMITATION OF STATE-OF-THE-ART

Table 1: Proportion of Irregular Graphs in Common Datasets

Dataset	Description	Number of Graphs	Percentage Irregular
AIDS [11]	Chemical Compounds	700	98.9%
LINUX [13]	Program-Dependence Graphs	1000	100%
IMDB [16]	Ego-Networks	1500	47.5%

The approach of utilizing parameter concentration [2] has been effective in optimizing regular graphs by transferring optimized parameters from smaller to larger instances. However, it faces limitations when applied to real-world irregular graphs that have more complex local features. Table 1 shows the prevalence of irregular graphs in commonly used datasets such as AIDS, LINUX, and IMDB. These graphs tend to lack the symmetry and structure that make it easy to transfer optimized parameters using the parameter concentration approach.

To address this limitation, prior works [5, 12] have explored the transferability of optimized parameters between graphs with

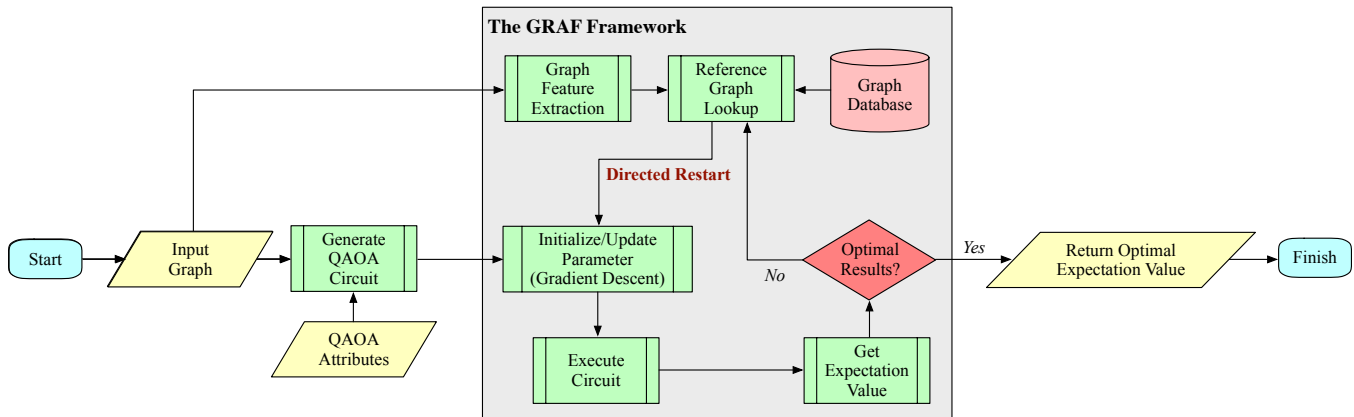


Figure 1: An overview of the proposed optimization framework that avoids random restarts by using graph feature extraction and reference graph lookup. The framework employs a lightweight graph database and performs directed restarts to output the optimal QAOA energy for a given graph and QAOA attributes.

similar structures, defined based on the parity of individual node degrees. If all nodes in a graph have an odd or even degree, then its optimized parameters can be transferred to other graphs with the same overall node parity. Figure 2 shows an example of parameter transfer between similar graphs.

However, the effectiveness of this approach for more general irregular graphs remains an open question. While prior research has shown promise in transferring optimized parameters between similar graphs, the limitations of applying this approach to irregular graphs with different structures and features remain a challenge for optimizing real-world irregular graphs. Ongoing research aims to refine and extend this approach to address these challenges.

3 KEY INSIGHTS

In the seminal QAOA paper [4], it is shown that the QAOA energy can be expressed as a sum of contributions from sub-graphs, and the overall QAOA energy is the sum of local QAOA energies. The local energy calculation involves a subgraph that consists of individual edges in the graph along with the nodes that are at most p distance away, where p is the number of QAOA layers. This way, we can represent each input graph as a weight vector, which is a linear combination of a fixed set of basis sub-graphs. When the normalized weight vectors of two graphs are similar, the QAOA circuit constructed from those two graphs will provide similar results for fixed parameters.




Figure 2: Example parameter transfer between similar graphs using prior work methods. Note that, A is a regular graph and B is an irregular graph.

For regular graphs, the number of basis sub-graphs is limited, but for irregular graphs, the number of basis sub-graphs increases significantly as we scale up the number of graph nodes, making it infeasible to compare graph similarities using basis sub-graphs.

One observation that can be used to find representative sub-graphs for irregular graphs efficiently is that the evaluation of the local QAOA energy involves the target nodes, the edges that connect to these nodes, and their neighboring nodes. Therefore, edge count significantly impacts the graph structure for a fixed graph diameter. We use this observation to perform a path-finding experiment in which we compare the *maxcut* values with the edge count difference of random irregular graphs. This experiment does not decompose these irregular graphs into sub-graphs, but it helps us understand the relationships between different irregular graphs.

Figure 3 displays the correlation between the number of edges and the *maxcut* value in irregular graphs. By identifying all possible sub-graphs, we found that a smaller difference in the number of edges between two graphs leads to a closer *maxcut* value, which




Figure 3: The *maxcut* difference versus edge count difference for pairs of irregular graphs. The pairs of basis graphs are from a set of 457 graphs with 2 to 9 nodes and 2 to 21 edges. There is a clear correlation between the graph’s *maxcut* value and the number of edges in the graph. Essentially, the smaller difference in the number of edges between two graphs is, the closer their *maxcut* values are.

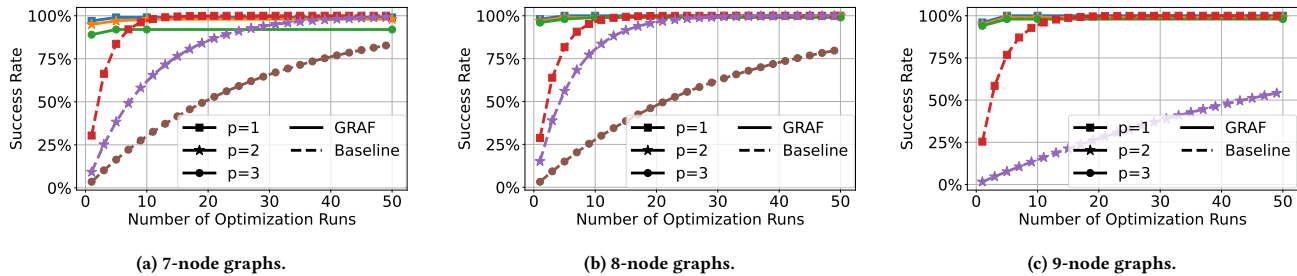


Figure 4: Success rate versus the number of optimization runs while using initial parameters generated from our framework and baseline. As the baseline optimizer performs repeated random initialization (restarts), it typically takes 10+ optimization runs to have a success rate of >90%. On the contrary, our approach achieves this with only three optimization runs.

also results in similar local QAOA energy. However, relying solely on edge count has high variability. Basic sub-graph decomposition helps, but the varied edge counts can limit its usefulness. Therefore, considering the mean and variance of the sub-graphs can serve as additional filters to identify similar graphs with transferable parameters and correlated maxcut values.

4 MAIN ARTIFACT

The main artifact proposed in this work is a framework that enhances the baseline optimizer with three components to find optimal parameters. The first component is a graph feature extraction block that considers various graph features such as total node count, average node count, and variance in node counts. The feature extraction block enables the framework to identify relevant sub-graphs and helps in efficient parameter transfer.

The second component is a lightweight graph database that stores pre-computed optimal parameters of graphs. The database acts as a reference or basis for transferring parameters efficiently, and it grows linearly with an increase in the number of nodes. Unlike the baseline optimizer, which tries to find optimal values through repeated ‘random’ trials, our approach makes use of the database of pre-computed optimal parameters to provide close-to-optimal starting points for optimization.

The third component is a low-cost graph lookup block that iteratively searches the database using the extracted features and provides starting points with close-to-optimal initial parameters. This approach increases the likelihood of finding the global maxima for these parameters, making the optimization process more efficient. The directed approach of using the database for parameter transfer and starting points, along with the feature extraction block, enables the framework to perform significantly fewer restarts than the baseline optimizer while converging on optimal parameters.

Figure 1 provides a high-level overview of the proposed framework. Overall, the framework is more efficient than the baseline optimizer and provides an innovative way to find optimal parameters by leveraging pre-computed optimal parameters, graph feature extraction, and a low-cost graph lookup block.

5 KEY RESULTS AND CONTRIBUTIONS

Figure 4 provides a comparison between the success rate of our framework and the baseline optimizer. The baseline optimizer uses

random initial points across 50 restarts, while our framework uses graph feature extraction and reference graph lookup to avoid random restarts and perform directed restarts. The graphs used in the experiments are randomly selected non-isomorphic 7, 8, and 9-node graphs and the graph dataset contains graphs with 5 and 6 nodes.

As shown in Figure 4, our framework significantly outperforms the baseline optimizer in terms of the achievable success rate across different numbers of restarts. For example, for the 7 node graphs, the baseline optimizer achieves a maximum success rate of around 75% with 50 restarts, while our framework achieves a success rate of over 97% with only 10 restarts. This demonstrates the effectiveness of our approach in finding globally optimal results efficiently.

Overall, the results show that our framework, which leverages graph feature extraction and reference graph lookup, outperforms the baseline optimizer, which uses random initial points across a large number of restarts. This suggests that our approach is more efficient and effective in finding globally optimal results for QAOA optimization problems.

6 CONCLUSION

In conclusion, we propose a framework that employs a directed restart strategy to efficiently find optimal parameters for QAOA optimization problems. By leveraging graph feature extraction and reference graph lookup, our approach avoids the overhead of random restarts and achieves globally optimal results with fewer restarts. Our results demonstrate the effectiveness of our approach, outperforming the baseline optimizer in terms of the achievable success rate across different numbers of restarts, even for larger graphs. These findings suggest that our approach can significantly improve the efficiency and effectiveness of QAOA optimization.

7 ACKNOWLEDGEMENT

This material is based upon work supported by the National Research Council Canada grant AQC 003. This work is also supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center. The Pacific Northwest National Laboratory is operated by Battelle for the U.S. Department of Energy under Contract DE-AC05-76RL01830.

REFERENCES

- [1] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph Bardin, Rami Barends, et al. Quantum supremacy using a programmable superconducting processor, 2019.
- [2] Fernando G. S. L. Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For Fixed Control Parameters the Quantum Approximate Optimization Algorithm's Objective Function Value Concentrates for Typical Instances, December 2018. arXiv:1812.04170 [quant-ph].
- [3] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon Benjamin, Suguro Endo, Keisuke Fujii, Jarrod Ryan McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick Coles. Variational quantum algorithms. *Nature Reviews Physics*, 2021.
- [4] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- [5] Alexey Galda, Xiaoyuan Liu, Danylo Lykov, Yuri Alexeev, and Ilya Safro. Transferability of optimal QAOA parameters between random graphs, June 2021. arXiv:2106.07531 [quant-ph].
- [6] Jay Gambetta. Expanding the IBM Quantum roadmap to anticipate the future of quantum-centric supercomputing. <https://research.ibm.com/blog/ibm-quantum-roadmap-2025>, 2022. [Online; accessed 23-August-2022].
- [7] Sergei V. Isakov, Dvir Kafri, Orion Martin, Catherine Vollgraf Heidweiller, Wojciech Mroczkiewicz, Matthew P. Harrigan, Nicholas C. Rubin, Ross Thomson, Michael Broughton, Kevin Kissell, Evan Peters, Erik Gustafson, Andy C. Y. Li, Henry Lamm, Gabriel Perdue, Alan K. Ho, Doug Strain, and Sergio Boixo. Simulations of quantum circuits with approximate noise using qsim and cirq, 2021.
- [8] Ang Li, Samuel Stein, Sriram Krishnamoorthy, and James Ang. Qasmbench: A low-level qasm benchmark suite for nisq evaluation and simulation, 2021.
- [9] Phillip C. Lotshaw, Travis S. Humble, Rebekah Herrman, James Ostrowski, and George Siopsis. Empirical performance bounds for quantum approximate optimization. *Quantum Information Processing*, 20(12):403, December 2021. arXiv:2102.06813 [physics, physics:quant-ph].
- [10] Andrea Morello and David Reilly. What would you do with 1000 qubits? *Quantum Science and Technology*, 3(3):030201, 2018.
- [11] Kaspar Riesen and Horst Bunke. IAM Graph Database Repository for Graph Based Pattern Recognition and Machine Learning. In Niels da Vitoria Lobo, Takis Kasparis, Fabio Roli, James T. Kwok, Michael Georgiopoulos, Georgios C. Anagnostopoulos, and Marco Loog, editors, *Structural, Syntactic, and Statistical Pattern Recognition*, volume 5342, pages 287–297. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. Series Title: Lecture Notes in Computer Science.
- [12] Ruslan Shaydulin, Phillip C. Lotshaw, Jeffrey Larson, James Ostrowski, and Travis S. Humble. Parameter Transfer for Quantum Approximate Optimization of Weighted MaxCut, January 2022. arXiv:2201.11785 [quant-ph].
- [13] Xiaoli Wang, Xiaofeng Ding, Anthony K.H. Tung, Shanshan Ying, and Hai Jin. An Efficient Graph Indexing Method. In *2012 IEEE 28th International Conference on Data Engineering*, pages 210–221, Arlington, VA, USA, April 2012. IEEE.
- [14] Zhihui Wang, Stuart Hadfield, Zhang Jiang, and Eleanor G. Rieffel. Quantum approximate optimization algorithm for MaxCut: A fermionic view. *Physical Review A*, 97(2):022304, February 2018.
- [15] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Strong quantum computational advantage using a superconducting quantum processor, 2021.
- [16] Pinar Yanardag and S.V.N. Vishwanathan. Deep Graph Kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1365–1374, Sydney NSW Australia, August 2015. ACM.
- [17] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-cheng Chen, Li-Chao Peng, and Yi-Han Luo. Quantum computational advantage using photons, 12 2020.
- [18] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. *Physical Review X*, 10(2):021067, June 2020.