



OPEN Performance analysis of multi-angle QAOA for $p > 1$

Igor Gaidai[✉] & Rebekah Herrman

In this paper we consider the scalability of multi-angle QAOA with respect to the number of QAOA layers. We found that MA-QAOA is able to significantly reduce the depth of QAOA circuits, by a factor of up to 4 for the considered data sets. Moreover, MA-QAOA is less sensitive to system size, therefore we predict that this factor will be even larger for big graphs. However, MA-QAOA was found to be not optimal for minimization of the total QPU time. Different optimization initialization strategies are considered and compared for both QAOA and MA-QAOA. Among them, a new initialization strategy is suggested for MA-QAOA that is able to consistently and significantly outperform random initialization used in the previous studies.

The Quantum Approximate Optimization Algorithm (QAOA)^{1–4} is a promising variational quantum algorithm designed to find approximate solutions to combinatorial optimization problems. With the appropriate transformations^{5,6}, the optimized cost function can be translated to the cost Hamiltonian C , and the original problem can be reformulated as finding the maximum eigenpair of C .

The basic version of QAOA finds large eigenpairs of C as follows. First, a second Hamiltonian B , called the mixing Hamiltonian, is introduced. It is usually defined as a sum of Pauli X operators applied to each qubit,

$$B = \sum_i X_i, \quad (1)$$

although other mixers can be found in the literature^{7–10}. Then an equal superposition state $|+\rangle^{\otimes n}$ (an eigenstate of B) is subjected to alternating evolution under the cost (C) and mixing (B) Hamiltonians for a total of p layers.

$$|\gamma, \beta\rangle = U(B, \beta_p)U(C, \gamma_p)\dots U(B, \beta_1)U(C, \gamma_1)|+\rangle^{\otimes n}, \quad (2)$$

$$U(B, \beta) = e^{-i\beta B}, \quad (3)$$

$$U(C, \gamma) = e^{-i\gamma C}. \quad (4)$$

The total number of layers (p) is a convergence parameter of the algorithm and the individual evolution times of each Hamiltonian ($\gamma_1\dots\gamma_p, \beta_1\dots\beta_p$, also known as QAOA angles) are selected to maximize the expectation of the cost Hamiltonian in the final QAOA state $\langle\gamma, \beta|C|\gamma, \beta\rangle$, usually via a classical optimizer. Finally, a series of repeated measurements of $|\gamma, \beta\rangle$ in the computational basis yields a bit string with a cost of at least $\langle\gamma, \beta|C|\gamma, \beta\rangle$.

QAOA has been a subject of numerous recent studies (e.g.^{11–23}), but many challenges still exist. One of such challenges is an efficient implementation of QAOA on the existing Noisy Intermediate Scale Quantum (NISQ) devices. NISQ devices suffer from high levels of noise^{24,25}, which means that only circuits of limited depths²⁶ can be executed before the qubits lose coherence and the result is completely dominated by noise. The number of QAOA layers necessary for efficient solution of practically relevant problems significantly exceeds the capacity of the modern quantum hardware. Recent work has studied methods for decomposing practical-sized problems to fit on hardware with fewer qubits, however some of these methods have drawbacks such as trying to piece together disjoint subproblem solutions^{27–29}. Therefore, to make QAOA more useful on NISQ devices it is important to design a version of QAOA that minimizes the necessary number of layers.

One way to achieve this is to allow independent evolution times for each term of the cost and mixer Hamiltonians, which changes the definitions of Eqs. (3) and (4) to

$$U(B, \vec{\beta}) = e^{-i\sum_i \beta_i X_i} = \prod_i e^{-i\beta_i X_i}, \quad (5)$$

Department of Industrial and Systems Engineering, University of Tennessee at Knoxville, 37996 Knoxville, TN, USA.
✉email: igaidai@utk.edu

$$U(C, \vec{\gamma}) = e^{-i \sum_i \gamma_i C_i} = \prod_i e^{-i \gamma_i C_i}. \quad (6)$$

Obviously, the performance of this modification cannot be worse than that of the original QAOA, since any QAOA result can be reproduced simply by setting all $\gamma_i = \gamma$ and all $\beta_i = \beta$, but how much better can it be? Does it justify the cost of having to optimize a significantly larger number of parameters?

This idea was originally suggested by Farhi et al.³⁰ who demonstrated optimistic, but very limited results. Later on, the performance of this modification, dubbed Multi-Angle QAOA (MA-QAOA), was further investigated on a larger data set and analyzed in Refs.^{31,32}, but the analysis was limited to $p = 1$. A similar recent development³³ suggested to combine the ideas of MA-QAOA with an XY-mixer⁸ to achieve even better performance, but they also considered the case of $p = 1$ only. Other related works include additional parameters on well-defined subproblems³⁴ and adding problem-independent ansatz layers with multiple parameters³⁵.

In practice we would like to take advantage of as many layers as our hardware can support to achieve the desired performance level, which can extend beyond $p = 1$ even on NISQ devices. How many layers will we need for that and how well does MA-QAOA scale beyond $p = 1$? Can we realistically find good angles even as p gets large? All of these questions need to be addressed to establish if MA-QAOA is useful for practical implementations on NISQ devices and the goal of this paper is to provide the answers to them.

Results

The performance of the methods considered in this paper is compared based on the approximation ratio that they can achieve for a given instance of the unweighted MaxCut problem, which is to partition the vertices of a graph into two subsets such that the number of edges between the sets is maximized. MaxCut is an important problem with a multitude of real-world applications (e.g., circuit design, statistical physics, image processing)^{36,37}

The approximation ratio is defined in the usual way as the ratio of the largest found expectation of the cost Hamiltonian to the maximum achievable cut on a given graph (found by brute-forcing through all possible partitions):

$$AR = \frac{\langle \gamma, \beta | C | \gamma, \beta \rangle}{C_{max}}. \quad (7)$$

The approximation ratio is a number in the range of $[0, 1]$.

The graphs for the MaxCut problem were generated randomly with fixed number of nodes and edge probability (Erdős–Rényi model). In order to explore how the performance of the methods scales with respect to the graph characteristics we generated a total of 7 data sets with 1000 connected, non-isomorphic graphs of specific QAOA covering depth in each (also referred to as *c-depth*, for brevity), which is defined as a minimum number of QAOA layers necessary for it to see the whole graph (starting from any edge)¹. This terminology is derived from Ref.³⁸, where the authors discuss the minimum number of iterations p that QAOA needs to “cover” the entire graph. As such, it is expected to correlate with the performance of QAOA. More specifically, *c-depth* is defined as the maximum depth of edge-BFS (i.e., Breadth First Search through the edges instead of nodes), taken over all edges.

This metric is closely related to the more conventional graph diameter as

$$\text{diameter} - 1 \leq \text{c-depth} \leq \text{diameter} + 1, \quad (8)$$

where the graph diameter is the maximum shortest distance between all pairs of vertices in the graph. Therefore, graph diameter and *c-depth* can be regarded as approximately equal. The exact diameter distribution and other characteristics of our data sets are summarized in Table 1.

The results of the quantum algorithms compared in this work are calculated on a classical state vector simulator implemented in Python. The source code of the simulator and the graph files are available at Ref.³⁹.

Selection of angles for QAOA

In general, it is difficult to find optimal QAOA parameters for the MaxCut problem analytically because the expected value function is a complex trigonometric function that depends on the degree of each vertex in the graph⁴⁰. Therefore, the optimal parameters are typically found by numerical optimization procedures.

Set #	Nodes	Edge probability	c-depth	#(diameter = c-depth - 1)	#(diameter = c-depth)	#(diameter = c-depth + 1)
1	9	0.6	3	663	332	5
2	10	0.6	3	762	237	1
3	11	0.6	3	799	201	0
4	12	0.6	3	823	176	1
5	12	0.2	4	32	481	487
6	12	0.1	5	7	223	770
7	12	0.1	6	12	211	777

Table 1. Characteristics of the seven data sets considered in this paper. The last three columns show the number of graphs with specific diameter within each data set.

One of the biggest challenges with this approach is the problem of selecting good initial angles for such optimization procedures. Selecting random initial angles may require an exponential (in p) number of restarts to find the best angles due to a large number of local minima and barren plateaus in the optimization landscape of QAOA, especially for large problem sizes and at high depth^{41,42}. As such, a large number of publications in the field has been devoted to finding heuristic strategies for selection of good initial angles^{43–50}.

For the sake of fair comparison with MA-QAOA, we wanted to use the best angles that we can find for QAOA within a reasonable budget of optimization attempts. Therefore, in this section we will compare empirical performance of several such heuristics reported in the literature and select the one that performs best. Specifically, the best heuristic is defined as the one that is able to exceed $AR = 16/17 \approx 0.941$ for all graphs in a given data set, using the smallest number of QAOA layers. The number 16/17 was chosen because it is NP-hard to approximate MaxCut with $AR > 16/17$ ^{33,51,52}, therefore this level, in general, cannot be achieved with polynomial amount of resources (time) classically and it is interesting to see how much resources (number of layers) QAOA needs for the same task.

The heuristics are compared on the data set with 9 nodes (the first row of Table 1). Each method was used to calculate all values of p consecutively, until one of them exceeds $AR = 16/17$ on every single graph in every data set. If a method failed to find better angles at a given level p , then the best angles found on level $p - 1$ are appended with zeros, and are declared as the best at level p . Therefore, the performance of each method can only increase monotonically with p . The average and worst case performances of the considered methods are shown in Fig. 1. The left frame shows the case when only 1 optimization attempt was allowed for each value of p . On the right frame, the methods were allowed to perform p optimizations at a given value of p .

The following heuristics were considered: Constant¹⁷, TQA⁴⁶, Interp⁴³, Fourier⁴³, Greedy⁴⁸, Random.

The first considered heuristic, labeled as “Constant” here, was taken from Ref.¹⁷. There, the authors suggest to initialize all values of γ with -0.01 and all values of β with 0.01 , i.e. a fixed constant that does not depend on anything and has opposite signs on γ and β . We examined convergence with constants in the set $\{0.01, 0.05, 0.1, 0.2, 0.4, 1\}$ (see Supplemental Information for details) and found that QAOA converged fastest when the value was set to 0.2 and appeared to decrease with the distance away from 0.2. This constant was used to generate the corresponding line in Fig. 1. This heuristic is only plotted in the left frame of Fig. 1, since it cannot make use of more than 1 optimization attempt.

The next heuristic, labeled as “TQA” (Trotterized Quantum Annealing), Ref.⁴⁶ assumes, as the name suggests, that good starting angles can be obtained by using linear parameter schedules typical for quantum annealing. Specifically, the values of γ are given by $\gamma_i = (i - 0.5)\Delta t/p$, and the values of β are given by $\beta = \Delta t - \gamma$, where Δt is the only unknown parameter, found by optimization in the one-dimensional Δt -space, after which the angles corresponding to the optimal Δt are used as a starting point for the full $2p$ -dimensional QAOA optimization. In the case of p optimization attempts (the right frame of Fig. 1), the first one-dimensional optimization was not counted in the optimization budget as a separate optimization. Additionally, since the number of local minima in the one-dimensional space can be rather limited, if the optimization procedure returned the value of Δt that has already been tried, then a random new value of Δt was selected. In our experience, we observed that starting the full optimization from a random non-optimal value of Δt can often perform better than starting from the optimal one.

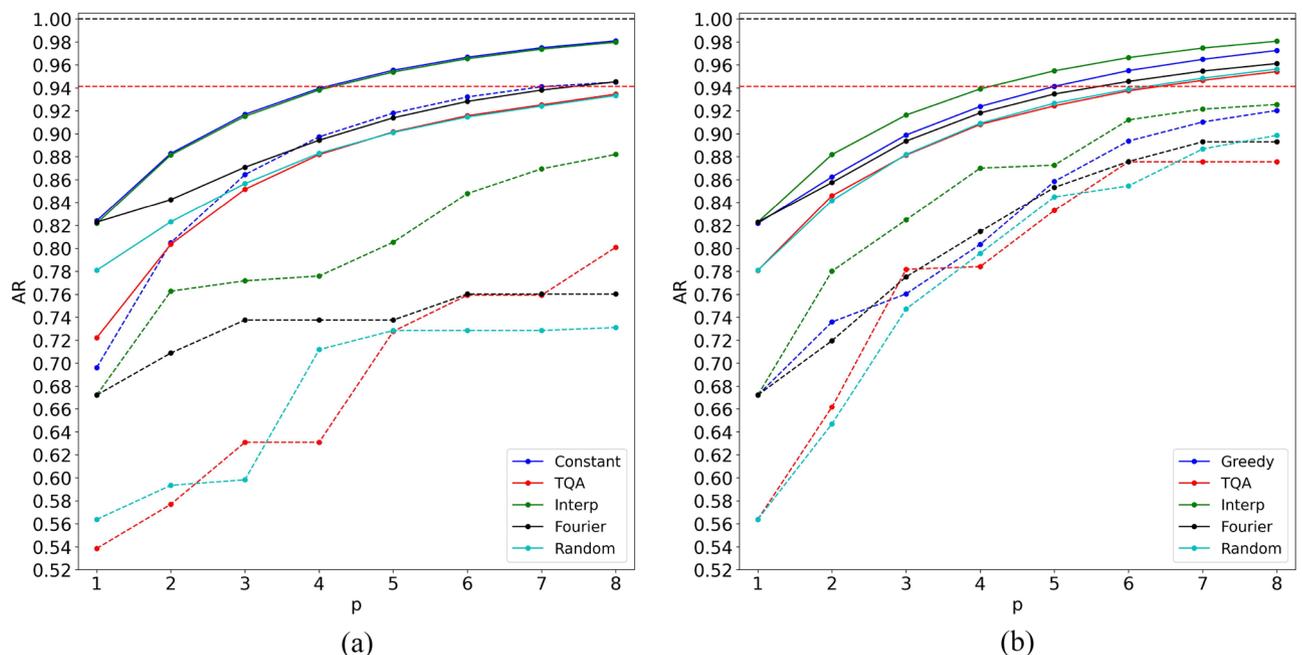


Figure 1. Comparison between different initialization heuristics for QAOA. Solid (dashed) lines show average (worst case) approximation ratio. The horizontal red dashed line shows the desired $AR = 16/17$.

The next two heuristics are taken from Ref.⁴³ and labeled as “Interp” and “Fourier”. These heuristics are similar in that they assume that the optimal angles change smoothly and re-evaluate the best angles found at level $p - 1$ on a denser grid or use Fourier transform to generate a good starting guess for level p . The extra optimization attempts on right frame of Fig. 1 are used to apply random perturbations to the best angles found at level $p - 1$ and start the optimization from there, in accordance with the prescription of Ref.⁴³.

Another heuristic, labeled “Greedy” can be found on the right frame of Fig. 1 only, and replaces “Constant” from the left frame. This heuristic (introduced in Ref.⁴⁸) is based on the idea that one can always insert an “empty” QAOA layer (i.e. a layer with $\gamma_i = \beta_i = 0$) before or after any other QAOA layer, which will have no effect on the expectation value of the cost Hamiltonian. As such, these zero insertions to the angle vectors can be used to generate good initial guesses for level p , using the best angles found at level $p - 1$. Indeed, starting the optimization from such angles guarantees that the expectation value achieved at level p will be no worse than the one achieved at level $p - 1$. There are p possible choices for the insertion position to angles found at level $p - 1$, therefore p optimizations are necessary to explore all initial guesses obtained in this way, which, in fact, determined the choice of the optimization budget for the right frame of Fig. 1.

Finally, the last considered method is to select the angles randomly, which defines the lowest reasonable performance threshold for any method. This method was used as an experimental control against which to compare the other initialization heuristics.

For all heuristics, except Greedy, L-BFGS-B optimization method from `scipy.optimize.minimize` python package was used. For Greedy, we used Nelder-Mead, since BFGS is a first-order method and is not suitable for optimization from a transition state, where the first-order derivatives are zero. For heuristics that require the best angles from the previous layer (Interp, Fourier and Greedy), at $p = 1$, we used a random initialization with 10 attempts (in both frames of Fig. 1) to start them with globally optimal angles.

As one can see from Fig. 1, Constant demonstrates superior average and worst-case performance for all p , even compared to methods that were allowed p optimization attempts. Interp has nearly the same performance on average, but is less consistent and does not do as well in the worst case even with added perturbations (on the right frame). Constant was the first to reach the NP-hard region of $AR > 16/17$ in the worst case at $p = 8$ and thus, Constant was selected to represent the performance of QAOA in the comparison against MA-QAOA (“Performance comparison between QAOA and MA-QAOA” section).

Selection of angles for MA-QAOA

In the previous MA-QAOA studies^{30–32} the initial angles were selected randomly. However, as seen in the “Selection of angles for QAOA” section, other techniques for finding initial angles tend to significantly outperform random initialization. Is it possible to adapt some of the QAOA heuristics to MA-QAOA? Will they still work well? In this section we introduce several new angle strategies for MA-QAOA and answer these questions.

The primary candidates for adaptation are the ones that perform best for QAOA. Looking at Fig. 1, we select Constant and Interp as such. Constant applies to MA-QAOA straightforwardly (set all $\gamma_i = 0.2$, and all $\beta_i = -0.2$). For Interp, we interpolate each angle γ_i and β_i separately, in the same way as it was for QAOA.

Another thing that we could try for MA-QAOA is to start the optimization from the optimal angles found by QAOA (with the Constant strategy, as discussed in the previous section), which we will call “QAOA Relax”.

To determine the role of quality of angles in QAOA Relax strategy, we also consider “Random QAOA”, where the initial values of γ_i and β_i are random, but equal within the same layer, as in QAOA. Finally, the last strategy is a completely random independent guess, which was the approach used in the previous studies^{30–32}.

The results of these strategies are shown in Fig. 2 on the nine-vertex data set. Random QAOA and Random strategies show nearly the same performance on average, both worse than QAOA Relax. This indicates that simply keeping the angles the same within each layer does not improve performance and that the quality of QAOA angles for the QAOA Relax strategy is important to achieve quick convergence.

In contrast to QAOA, Interp does not perform well as an MA-QAOA initialization approach. In fact, the AR with the Interp initialization strategy gets even worse than the Random initialization strategy by $p = 3$, which indicates that optimal MA-QAOA angles tend to not change smoothly.

Both Constant and QAOA Relax are nearly identical on average and both reach the target convergence at $p = 3$, but QAOA Relax performs better in the worst case and is able to solve all MaxCut problems in the data set exactly. As such, we select it as the best strategy for comparison with QAOA below.

Performance comparison between QAOA and MA-QAOA

Using the best angle strategies picked in the two previous sections (Constant for QAOA and QAOA Relax for MA-QAOA), we calculated the AR on all data sets from Table 1. The results of this are shown in Fig. 3. As one can see, the average AR of both QAOA and MA-QAOA converges to 1 smoothly. The worst case AR is generally less smooth, but still follows a similar trend. At sufficiently large values of p , the performance of both QAOA and MA-QAOA expectedly drops as the system size is increasing, but MA-QAOA is less affected by it. Note, that at low values of p , the order of the lines in Fig. 3a is reversed, which indicates that the performance of QAOA-like methods in general cannot be adequately judged based on the results for $p = 1$ only, and convergence with respect to p has to be analyzed.

Interestingly, from Fig. 3b, one can see that QAOA converges faster for graphs of larger c -depth, while for MA-QAOA the behavior is the opposite. This might prompt the conclusion that QAOA could become better than MA-QAOA for graphs of large c -depth. Note, however, that the maximum c -depth is limited by the number of nodes in a graph and increasing the number of nodes has the opposite and stronger effect on the performance of QAOA. Additionally, the performance of MA-QAOA cannot be worse than performance of QAOA since we start the optimization from the optimal QAOA angles.

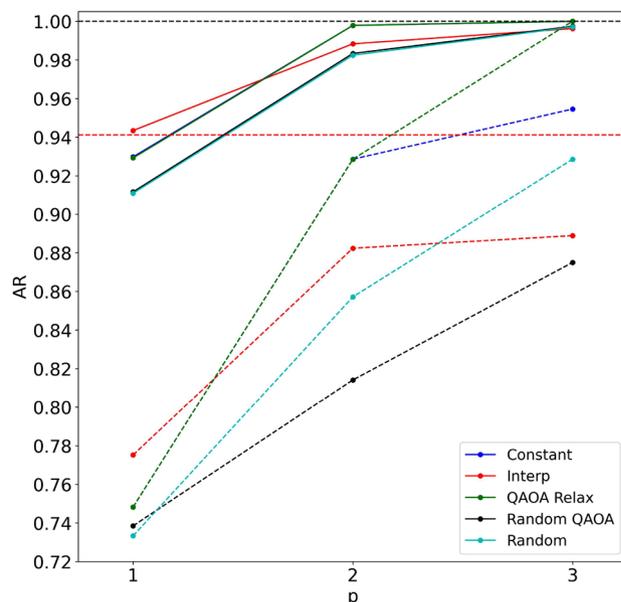


Figure 2. Comparison between different initialization heuristics for MA-QAOA on the data set with 9 nodes. Solid (dashed) lines show average (worst case) approximation ratio. The horizontal red dashed line shows the desired AR = 16/17.

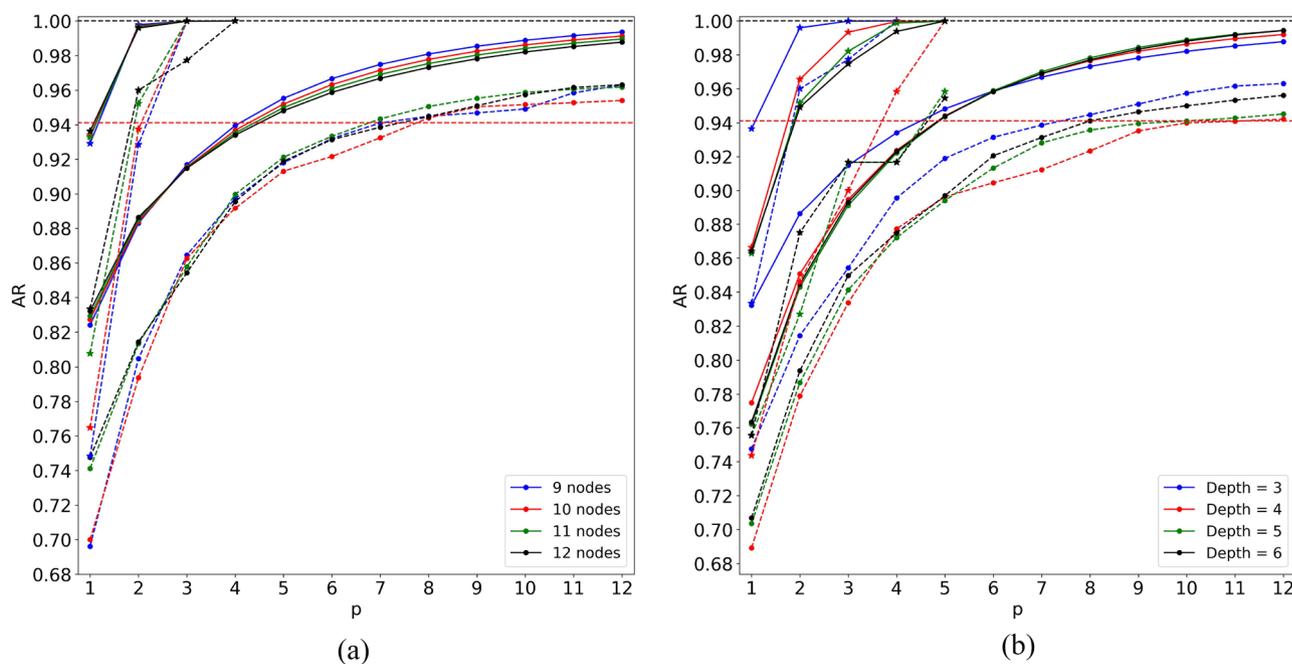


Figure 3. Approximation ratio as a function of p obtained with the best angle selection strategies for QAOA and MA-QAOA. Solid (dashed) lines show average (worst case) approximation ratio. Circles (stars) show the results of QAOA (MA-QAOA). The horizontal red dashed line shows the desired AR = 16/17.

In both frames of Fig. 3, we see that MA-QAOA achieves substantially larger ARs and thus converges much faster than QAOA. In particular, the number of layers necessary to converge the average AR required by QAOA was 5 for all data sets, while MA-QAOA was able to achieve the same performance level in 2 layers, reducing the depth of the circuit by a factor of 2.5. The number of layers necessary for the worst case AR is more variable and is given in Table 2. As one can see, in this case MA-QAOA allows to reduce the depth of the circuit by a factor of up to 4. As a consequence of the trends described above, the depth reduction factors tend to increase with the graph size (sets 1–4), but decrease with c -depth (sets 4–7).

Set #	# layers QAOA	# layers MA-QAOA	Depth reduction factor
1	8	3	2.7
2	8	3	2.7
3	7	2	3.5
4	8	2	4.0
5	12	4	3.0
6	11	5	2.2
7	9	5	1.8

Table 2. The number of layers necessary to converge the worst-case approximation ratios for QAOA and MA-QAOA. The set numbers are the same as in Table 1.

Prospects of MA-QAOA in fault-tolerant era

In the previous section we demonstrated that the depth of QAOA circuit can be made several times smaller by using MA-QAOA. But what if we have a hypothetical fault-tolerant computer and the depth does not matter as much anymore? In that case the new efficiency metric that one can define instead of circuit depth, is the total time it takes to solve the problem, which can be approximated as the number of calls to QPU times depth of the corresponding circuits. Both QAOA and MA-QAOA have the same depth per layer, therefore for the sake of comparison between them, one can measure circuit depth simply in the number of QAOA layers.

$$\text{Cost} = n_c \times p, \tag{9}$$

where n_c symbolizes the number of calls to a QPU.

Higher dimensionality of the optimization space means more calls to QPU. Will MA-QAOA still be better than QAOA with this metric, i.e. will it have better AR for a given cost? In Fig. 4, we re-plotted the data of Fig. 3, using the metric of Eq. (9) as the argument. As one can see, overall, QAOA and MA-QAOA are closely matched in the region of smaller values of p . On average, QAOA achieves the desired accuracy level at a substantially smaller cost compared to MA-QAOA. The same is true for most (5 out of 7) data sets in the worst case. However, starting at the sufficiently high values of p ($p > 1$, data set dependent), MA-QAOA is able to achieve consistently larger ARs at any given cost. Considering the scaling with respect to the system size, the minimum value of p at which this happens is expected to grow.

Conclusions

To summarize, the main novel contributions of this paper are the following

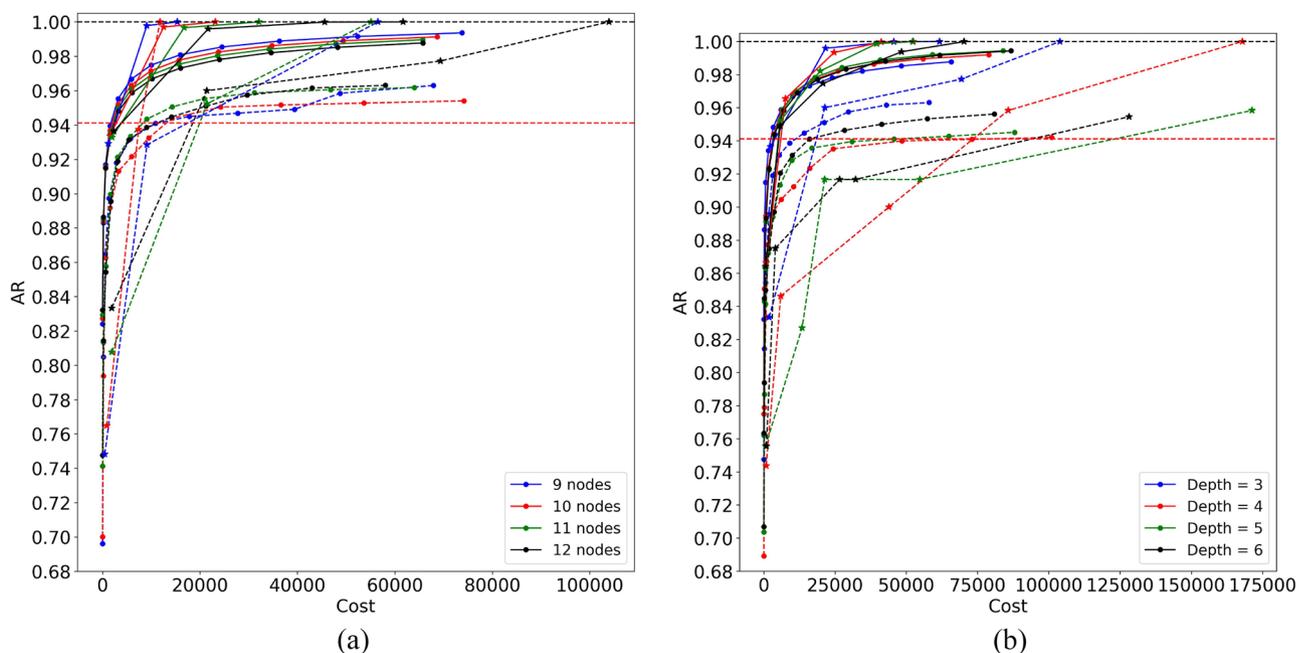


Figure 4. The data of Fig. 3, re-plotted using the cost defined in Eq. (9).

1. We investigated the numerical scaling of MA-QAOA for $p > 1$ and found that the advantage of MA-QAOA over QAOA is preserved even when multiple layers are considered, which allows it to converge to the desired accuracy in a much smaller number of layers, compared to QAOA. In addition to that, MA-QAOA was found to be less susceptible to performance decay associated with the problem size.
2. We found a new initialization strategy for MA-QAOA that allows to consistently outperform random initialization used in the previous studies. Additionally, this strategy guarantees that the performance of MA-QAOA will be no worse than that of QAOA.
3. We analyzed the prospects of MA-QAOA on a hypothetical fault-tolerant computer, where the depth of the circuit is no longer the primary concern, and found that in this scenario MA-QAOA no longer has an advantage over QAOA.

Additionally, we compared the performance of multiple parameter setting strategies for QAOA (Constant, TQA, Interp, Fourier, Greedy, Random) and MA-QAOA (Constant, Interp, QAOA Relax, Random QAOA, Random). Most of these initialization strategies have already been reported in the literature, but not all of them have been directly compared against each other on the same data sets. We consider this a minor, but interesting contribution.

These contributions are important because they establish the advantage of MA-QAOA over QAOA in a practically relevant setting of $p > 1$ and demonstrate that good angles can be found efficiently.

One limitation of the present study is that we considered random graphs only. It is possible that different classes of graphs might display different characteristics, therefore affecting the performance of MA-QAOA. Exploring this possibility could be a subject of future research.

As noted in this work, some QAOA angle initialization strategies do not work well for MA-QAOA. It is reasonable to assume that there may be MA-QAOA parameter initialization strategies that are not well-defined for QAOA due to the drastic difference in the number of parameters the algorithms require. As such, another direction for future work might include developing new angle initialization techniques for MA-QAOA that may not necessarily make sense for QAOA.

Another research direction could be to reduce the optimization cost of MA-QAOA by selecting only most relevant degrees of freedom.

Data availability

The full dataset of input graphs and raw results of calculations can be found at <https://zenodo.org/records/12627513>. Additionally, the source code can be found at <https://github.com/Gaidailgor/MA-QAOA>.

Received: 17 May 2024; Accepted: 7 August 2024

Published online: 14 August 2024

References

1. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. <http://arxiv.org/abs/1411.4028> (2014).
2. Choi, J. & Kim, J. A tutorial on quantum approximate optimization algorithm (qaoa): Fundamentals and applications. In *2019 International Conference on Information and Communication Technology Convergence (ICTC)* 138–142. <https://doi.org/10.1109/ICTC46691.2019.8939749> (2019).
3. Blekos, K. *et al.* A review on quantum approximate optimization algorithm and its variants. *Phys. Rep.* **1068**, 1. <https://doi.org/10.1016/j.physrep.2024.03.002> (2024).
4. Hogg, T. Quantum search heuristics. *Phys. Rev. A* **61**, 052311. <https://doi.org/10.1103/PhysRevA.61.052311> (2000).
5. Hadfield, S. *et al.* From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **12**, 34. <https://doi.org/10.3390/a12020034> (2019).
6. Hadfield, S. On the representation of Boolean and real functions as Hamiltonians for quantum computing. *ACM Trans. Quantum Comput.* **2**, 1. <https://doi.org/10.1145/3478519> (2021).
7. Bärttschi, A. & Eidenbenz, S. Grover mixers for qaoa: Shifting complexity from mixer design to state preparation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* 72–82. <https://doi.org/10.1109/QCE49297.2020.00020> (IEEE, 2020).
8. Wang, Z., Rubin, N. C., Dominy, J. M. & Rieffel, E. G. *xy* mixers: Analytical and numerical results for the quantum alternating operator ansatz. *Phys. Rev. A* **101**, 012320. <https://doi.org/10.1103/PhysRevA.101.012320> (2020).
9. Fuchs, F. G., Lye, K. O., Møll Nilsen, H., Stasik, A. J. & Sartor, G. Constraint preserving mixers for the quantum approximate optimization algorithm. *Algorithms* **15**, 202. <https://doi.org/10.3390/a15060202> (2022).
10. Zhu, L. *et al.* Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer. *Phys. Rev. Res.* **4**, 033029. <https://doi.org/10.1103/PhysRevResearch.4.033029> (2022).
11. Crooks, G. E. Performance of the quantum approximate optimization algorithm on the maximum cut problem. <http://arxiv.org/abs/1811.08419> (2018).
12. Wang, Z., Hadfield, S., Jiang, Z. & Rieffel, E. G. Quantum approximate optimization algorithm for maxcut: A fermionic view. *Phys. Rev. A* **97**, 022304. <https://doi.org/10.1103/PhysRevA.97.022304> (2018).
13. Wurtz, J. & Love, P. Maxcut quantum approximate optimization algorithm performance guarantees for $p > 1$. *Phys. Rev. A* **103**, 042612. <https://doi.org/10.1103/PhysRevA.103.042612> (2021).
14. Basso, J., Farhi, E., Marwaha, K., Villalonga, B. & Zhou, L. The quantum approximate optimization algorithm at high depth for maxcut on large-girth regular graphs and the Sherrington–Kirkpatrick model. In *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, Vol. 232, 1–7. <https://doi.org/10.4230/LIPIcs.TQC.2022.7> (2022).
15. Marwaha, K. Local classical max-cut algorithm outperforms $p = 2$ qaoa on high-girth regular graphs. *Quantum* **5**, 437 (2021).
16. Basso, J., Gamarnik, D., Mei, S. & Zhou, L. Performance and limitations of the qaoa at constant levels on large sparse hypergraphs and spin glass models. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* 335–343. <https://doi.org/10.1109/FOCS54457.2022.00039> (2022).
17. Boulebnane, S. & Montanaro, A. Solving Boolean satisfiability problems with the quantum approximate optimization algorithm. <http://arxiv.org/abs/2208.06909> (2022).
18. Hadfield, S., Hogg, T. & Rieffel, E. G. Analytical framework for quantum alternating operator ansätze. *Quant. Sci. Technol.* **8**, 015017. <https://doi.org/10.1088/2058-9565/ACA3CE> (2022).

19. Herrman, R. Relating the multi-angle quantum approximate optimization algorithm and continuous-time quantum walks on dynamic graphs. <http://arxiv.org/abs/2209.00415> (2022).
20. Ozaeta, A., van Dam, W. & McMahon, P. L. Expectation values from the single-layer quantum approximate optimization algorithm on ising problems. *Quant. Sci. Technol.* **7**, 045036. <https://doi.org/10.1088/2058-9565/ac9013> (2022).
21. Lykov, D. *et al.* Sampling frequency thresholds for the quantum advantage of the quantum approximate optimization algorithm. *NPJ Quant. Inf.* **9**, 73. <https://doi.org/10.1038/s41534-023-00718-4> (2023).
22. Shaydulin, R. *et al.* Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem. *Sci. Adv.* **10**, 6761. <https://doi.org/10.1126/sciadv.adm6761> (2024).
23. Stechly, M., Gao, L., Yogendran, B., Fontana, E. & Rudolph, M. Connecting the Hamiltonian structure to the qaoa energy and fourier landscape structure. <http://arxiv.org/abs/2305.13594>
24. Sun, J. *et al.* Mitigating realistic noise in practical noisy intermediate-scale quantum devices. *Phys. Rev. Appl.* **15**, 034026. <https://doi.org/10.1103/PhysRevApplied.15.034026> (2021).
25. Dasgupta, S. & Humble, T. S. Stability of noisy quantum computing devices. <http://arxiv.org/abs/2105.09472> (2021).
26. Herrman, R., Ostrowski, J., Humble, T. S. & Siopsis, G. Lower bounds on circuit depth of the quantum approximate optimization algorithm. *Quant. Inf. Process.* **20**, 1. <https://doi.org/10.1007/s11128-021-03001-7> (2021).
27. Zhou, Z., Du, Y., Tian, X. & Tao, D. Qaoa-in-qaoa: Solving large-scale maxcut problems on small quantum machines. *Phys. Rev. Appl.* **19**, 024027. <https://doi.org/10.1103/PhysRevApplied.19.024027> (2023).
28. Ponce, M. *et al.* Graph decomposition techniques for solving combinatorial optimization problems with variational quantum algorithms. <http://arxiv.org/abs/2306.00494> (2023).
29. Li, J., Alam, M. & Ghosh, S. Large-scale quantum approximate optimization via divide-and-conquer. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **42**, 1852. <https://doi.org/10.1109/TCAD.2022.3212196> (2023).
30. Farhi, E., Goldstone, J., Gutmann, S. & Neven, H. Quantum algorithms for fixed qubit architectures. <http://arxiv.org/abs/1703.06199> (2017).
31. Herrman, R., Lotshaw, P. C., Ostrowski, J., Humble, T. S. & Siopsis, G. Multi-angle quantum approximate optimization algorithm. *Sci. Rep.* **12**, 6781. <https://doi.org/10.1038/s41598-022-10555-8> (2022).
32. Shi, K. *et al.* Multiangle qaoa does not always need all its angles. In *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)* 414. <https://doi.org/10.1109/SEC54971.2022.00062> (2022).
33. Vijendran, V., Das, A., Koh, D. E., Assad, S. M. & Lam, P. K. An expressive ansatz for low-depth quantum approximate optimisation. *Quant. Sci. Technol.* **9**, 025010. <https://doi.org/10.1088/2058-9565/ad200a> (2024).
34. Wurtz, J. & Love, P. J. Classically optimal variational quantum algorithms. *IEEE Trans. Quant. Eng.* **2**, 1. <https://doi.org/10.1109/TQE.2021.3122568> (2021).
35. Chalupnik, M., Melo, H., Alexeev, Y. & Galda, A. Augmenting qaoa ansatz with multiparameter problem-independent layer. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* 97–103. <https://doi.org/10.1109/QCE53715.2022.00028> (2022).
36. Gu, S. & Yang, Y. A deep learning algorithm for the max-cut problem based on pointer network structure with supervised learning and reinforcement learning strategies. *Mathematics* **8**, 298. <https://doi.org/10.3390/math8020298> (2020).
37. Barahona, F., Grötschel, M., Jünger, M. & Reinelt, G. An application of combinatorial optimization to statistical physics and circuit layout design. *Oper. Res.* **36**, 493. <https://doi.org/10.1287/opre.36.3.493> (1988).
38. Farhi, E., Gamarnik, D. & Gutmann, S. The quantum approximate optimization algorithm needs to see the whole graph: A typical case. <http://arxiv.org/abs/2004.09002> (2020).
39. <https://github.com/GaidaiIgor/MA-QAOA>.
40. Hadfield, S. A. Quantum algorithms for scientific computing and approximate optimization. <http://arxiv.org/abs/1805.03265> (Columbia University, 2018)
41. Lee, X., Saito, Y., Cai, D. & Asai, N. Parameters fixing strategy for quantum approximate optimization algorithm. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* 10–16. <https://doi.org/10.1109/QCE52317.2021.00016> (2021).
42. Wang, S. *et al.* Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.* **12**, 6961. <https://doi.org/10.1038/s41467-021-27045-6> (2021).
43. Zhou, L., Wang, S.-T., Choi, S., Pichler, H. & Lukin, M. D. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X* **10**, 021067. <https://doi.org/10.1103/PhysRevX.10.021067> (2020).
44. Boulebnane, S. & Montanaro, A. Predicting parameters for the quantum approximate optimization algorithm for max-cut from the infinite-size limit. <http://arxiv.org/abs/2110.10685> (2021).
45. Galda, A., Liu, X., Lykov, D., Alexeev, Y. & Saftro, I. Transferability of optimal qaoa parameters between random graphs. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* 171. <https://doi.org/10.1109/QCE52317.2021.00034> (2021).
46. Sack, S. H. & Serbyn, M. Quantum annealing initialization of the quantum approximate optimization algorithm. *Quantum* **5**, 491. <https://doi.org/10.22331/q-2021-07-01-491> (2021).
47. Farhi, E., Goldstone, J., Gutmann, S. & Zhou, L. The quantum approximate optimization algorithm and the Sherrington–Kirkpatrick model at infinite size. *Quantum* **6**, 759. <https://doi.org/10.22331/q-2022-07-07-759> (2022).
48. Sack, S. H., Medina, R. A., Kueng, R. & Serbyn, M. Recursive greedy initialization of the quantum approximate optimization algorithm with guaranteed improvement. *Phys. Rev. A* **107**, 062404. <https://doi.org/10.1103/PhysRevA.107.062404> (2023).
49. Sud, J., Hadfield, S., Rieffel, E., Tubman, N. & Hogg, T. Parameter-setting heuristic for the quantum alternating operator ansatz. *Phys. Rev. Res.* **6**, 023171. <https://doi.org/10.1103/PhysRevResearch.6.023171> (2024).
50. Sureshbabu, S. H. *et al.* Parameter setting in quantum approximate optimization of weighted problems. *Quantum* **8**, 1231. <https://doi.org/10.22331/q-2024-01-18-1231> (2024).
51. Hästad, J. Some optimal inapproximability results. *J. ACM* **48**, 798–859. <https://doi.org/10.1145/502090.502098> (2001).
52. Trevisan, L., Sorkin, G. B., Sudan, M. & Williamson, D. P. Gadgets, approximation, and linear programming. *SIAM J. Comput.* **29**, 2074. <https://doi.org/10.1137/S0097539797328847> (2000).

Acknowledgements

R. Herrman acknowledges the National Science Foundation award CCF-2301120.

Author contributions

I.G.: Study design, Code, Data acquisition, Figures, Original draft of the text R.H.: Supervision, Text review and editing, Funding acquisition.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-69643-6>.

Correspondence and requests for materials should be addressed to I.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024