

Limitations of quantum approximate optimization in solving generic higher-order constraint-satisfaction problems

Thorge Müller^{1,2,*}, Ajainderpal Singh³, Frank K. Wilhelm^{2,3} and Tim Bode³

¹*German Aerospace Center (DLR), Institute for Software Technology, Department High-Performance Computing, 51147 Cologne, Germany*

²*Theoretical Physics, Saarland University, 66123 Saarbrücken, Germany*

³*Institute for Quantum Computing Analytics (PGI-12), Forschungszentrum Jülich, 52425 Jülich, Germany*



(Received 2 December 2024; accepted 4 February 2025; published 19 May 2025)

The ability of the quantum approximate optimization algorithm (QAOA) to deliver a quantum advantage on combinatorial optimization problems is still unclear. Recently, a scaling advantage over a classical solver was postulated to exist for random 8-SAT at the satisfiability threshold. At the same time, the viability of quantum error mitigation for deep circuits on near-term devices has been put into doubt. Here we analyze the QAOA's performance on random Max- k XOR as a function of k and the clause-to-variable ratio. As a classical benchmark, we use the mean-field approximate optimization algorithm and find that it performs better than or equal to the QAOA on average. Still, for large k and numbers of layers p , there may remain a window of opportunity for the QAOA. However, by extrapolating our numerical results, we find that reaching high levels of satisfaction would require extremely large p , which must be considered rather difficult both in the variational context and on near-term devices.

DOI: [10.1103/PhysRevResearch.7.023165](https://doi.org/10.1103/PhysRevResearch.7.023165)

I. INTRODUCTION

Quantum combinatorial optimization was applied to satisfiability problems early on [1], underlining the naturalness of the question if, in essence, quantum superposition could aid combinatorial searches in exponentially large spaces. The interest in such questions only grew after the introduction of the quantum approximate optimization algorithm (QAOA) by Farhi *et al.* [2]. Whether either adiabatic quantum computation [3] or the QAOA can deliver an advantage over any known best classical algorithm has remained an open question, however, not least because beating classical methods is a moving target.

Even so, Boulebnane and Montanaro [4], by analyzing the performance of the QAOA for the *decision* problem k SAT with even k , recently found a potential speedup of the QAOA over a state-of-the-art classical solver when applied to random 8SAT. A practical caveat to this result is the large number of QAOA layers required to push the (still exponential) scaling below the classical benchmark. A very recent result also hinted at possible speedups for almost symmetric optimization problems from *low-depth* QAOA [5].

This practical difficulty should not be underestimated, as a unique learning-theoretic proof [6] has exponentially

tightened the limitations of quantum error mitigation, thus severely reducing the prospects of variational quantum algorithms on near-term devices, which naturally includes the QAOA. One key message is that even shallow circuits will have a hard time estimating expectation values at current noise levels. In the worst case, this would mean that the QAOA is effectively limited to low numbers of layers until quantum error correction arrives. The QAOA had originally drawn attention [7] as one of the first quantum algorithms that might show an advantage over classical algorithms without fault-tolerant quantum computation.

In the case of optimization problems involving higher-order Pauli-Z strings, as explored below, the need to decompose these many-body interactions into two-body couplings further increases the circuit depth on platforms such as superconducting devices.

In a similar vein, a number of results imply that models with a provable absence of barren plateaus turn out to be efficiently classically simulable [8,9]. Reference [10] goes further by conjecturing that the Lie algebras of most problem graphs grow either exponentially or polynomially in system size, meaning that the QAOA will either suffer from barren plateaus or turn out to be, yet again, classically simulable.

The interest in the QAOA's potential naturally extends beyond k SAT and their respective optimization problems, Max- k SAT, in particular, to the related Max- k XOR problems, with Max-2XOR having received the most thorough scrutiny in the literature [11–13]. We remark that these problems belong to the complexity class NP-hard for $k \geq 2$ [14]. Compared to the nonexclusive OR ($\vee$) of Max- k SAT, the literals per constraint of Max- k XOR problem instances are connected via an *exclusive* OR ($\oplus$). The conversion from a classical

*Contact author: Thorge.Mueller@dlr.de

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

Max- k XOR instance to a quantum operator turns a k -literal constraint into a Pauli-Z string of length k .

Crooks [15] provided numerical evidence that the QAOA could indeed yield an advantage over classical algorithms on these problem types, showing that for an increasing number of parameters $2p$, the QAOA can at least surpass the lower bound established by the classic Max-2XOR algorithm due to Goemans and Williamson [16], which derives from semidefinite programming. Note that Makarychev and Makarychev [17] showed that a general k -constrained satisfaction problem (k -CSP) has a lower bound than a $(k - 1)$ -CSP, i.e., increasing k should be an interesting avenue to pursue in the search for an advantage.

Consequently, here delve into the performance of the QAOA on problems with $k > 2$ literals (constraints) per clause, where Akshay *et al.* [18] conducted a performance investigation of the QAOA as a function of the clause-to-variable ratio r for Max- k SAT. They showed that for $k = 3$, one needs higher p than for $k = 2$ to achieve the same accuracy of the locally optimal solution.

Another interesting result was given by Marwaha and Hadfield [19] for *uniform* Max- k XOR problems, who observed better performance than the best-known classical algorithm, the local threshold algorithm [20]. The graphs associated with the problem instances in their investigation were *regular* and *triangle-free*. Similarly, the works of Chou *et al.* [21], Chen *et al.* [22], and Basso *et al.* [23] focused on random Max- k XOR problems where the related graph is *sparse*. They especially drew attention to the so-called overlap gap property [24,25] for even k ; loosely speaking, this states that if two suboptimal (or optimal) sets of satisfied clauses are close in energy, their bitstrings are either very close or very far from each other in Hamming distance. They showed analytically that, due to this property, the QAOA needs increasingly high p values for Max- k XOR with even k .

In contrast to these previous results, we explore *arbitrary* random Max- k XOR problems. Additionally, we examine the performance specifically for *odd* values of k , which complements the studies conducted in Refs. [21–23].

This paper is structured as follows: We first introduce the QAOA, the mean-field approximate optimization algorithm (MF-AOA), and the Max- k XOR problem. Then we proceed by applying the QAOA and the MF-AOA to ensembles of randomly generated Max- k XOR instances. We begin by examining the dependence of their performance on the clause-to-variable ratio r . Subsequently, we analyze the impact of the parameter k on the algorithms' behavior. Next, we explore the performance of the QAOA at high circuit depths, including an analysis of the distribution of the optimized parameters. Based on these results, we estimate the circuit depth at which QAOA achieves 99% approximation ratio. Finally, we compare the performance of the QAOA and MF-AOA across different values of k .

II. MAX- k XOR

As mentioned, the distinction between k SAT and k XOR problems is that the variables within a clause of the latter are linked by the exclusive OR ($\oplus$) operator. In fact, a random

TABLE I. Truth table for a 3XOR clause consisting of the $N = 3$ variables x_1, x_2 , and x_3 . The rightmost column shows the results under which assignment to the variables the clause is satisfied.

x_1	x_2	x_3	$x_1 \oplus x_2 \oplus x_3$
0	0	0	0
0	0	1	1
0	1	0	1
0	1	1	0
1	0	0	1
1	0	1	0
1	1	0	0
1	1	1	1

k XOR decision problem can be written as

$$C = \bigwedge_{\{i_1, i_2, \dots, i_k\} \in M} x_{i_1} \oplus x_{i_2} \oplus \dots \oplus x_{i_k}, \quad (1)$$

where the $x_1, \dots, x_N \in \{0, 1\}$ are the N variables defining the problem and M denotes the set of clauses. As an example, if we pick a 2XOR problem with three variables and three clauses, we have

$$C = (x_1 \oplus x_2) \wedge (x_2 \oplus x_3) \wedge (x_1 \oplus x_3). \quad (2)$$

To go from the decision problem to the optimization problem, we transform $\wedge \rightarrow \sum$ and $\oplus \rightarrow +$. Hence, C can also be expressed as

$$C = \sum_{\{i_1, i_2, \dots, i_k\} \in M} (x_{i_1} + x_{i_2} + \dots + x_{i_k} + a_{i_1 \dots i_k}) \bmod 2, \quad (3)$$

where the symbol $a_{i_1 \dots i_k} \in \{0, 1\}$ represents the parity sign corresponding to the selected clause, fixing the preferred truth convention for that specific clause. In the context of Max-2XOR, all parity signs are set to 1 and all variables are non-negated, resulting in the well-known standard formulation of Max-Cut. If instead we negate one variable in every clause by convention, each clause would be satisfied if and only if *both* variables are equal, which is the opposite of standard Max-Cut.

The problem can be translated into a quantum operator $C \rightarrow \hat{C}$. To understand the transformation, we investigate the exemplary truth table of Max-3XOR in Table I. The outcome of the truth table (rightmost column of Table I) can be represented by Pauli-Z strings and the identity matrix. Hence, we identify $\{0, 1\} \rightarrow \{1, -1\}$ and obtain $\hat{C} = 1 - \hat{Z}_1 \hat{Z}_2 \hat{Z}_3$ for this single-clause example. We can easily generalize this result to Max- k XOR, resulting in

$$\hat{C} = \sum_{\{i_1, i_2, \dots, i_k\} \in M} \frac{1}{2} (1 \pm \hat{Z}_{i_1} \hat{Z}_{i_2} \dots \hat{Z}_{i_k}), \quad (4)$$

where the signs depend on the respective parities $a_{i_1 \dots i_k}$ from Eq. (3).

We finish this section by discussing how to sample these random Max- k XOR instances. The clause-to-variable ratio is defined as

$$r = \frac{|M|}{N}. \quad (5)$$

Because we are interested in *random* graphs, we need to determine the clause probability corresponding to specific values of r , N , and k . Now, one might assume that there are 2^k total clauses for a given k literal, given the inclusion of negated variables. This would imply that the number of all possible clauses for Max- k XOR at system size N is $2^k \binom{N}{k}$. However, by looking at the exemplary truth table shown in Table I, we see that the truth value of each of the 2^k clauses only depends on the *parity* of negations, i.e., an *odd* number of negated variables results in the corresponding clause being violated, while an even number of negations results in a true clause. Effectively, then, for the k -literal of Table I, we have only *two* distinct outcomes with *opposite* parity signs in the cost function. Clearly, including both of these would render the cost function trivial. For this reason, we randomly choose among these two possibilities for each k literal, which results in $2 \binom{N}{k}$ being the effective total number of clauses. Therefore, to find a specific random instance of Max- k XOR for given r , one can randomly pick from the set of all possible clauses with a probability

$$\mathcal{P} = rN \binom{N}{k}^{-1}. \quad (6)$$

To calculate the performance of QAOA applied to a specific Max- k XOR problem with fixed r , we then average over a set of random instances per value of k . Furthermore, we only take pure k -literal constraints into account [instead of $(k-1)$ -literals, etc.] because we want to investigate the performance of the QAOA in clean dependence on k .

III. METHODS

A. The QAOA

Every constrained satisfaction problem with k literals can be converted into a quantum operator consisting of diagonal quantum Pauli-Z strings, expressed as

$$\hat{C} = \sum_{i_1, \dots, i_k} J_{i_1, \dots, i_k} \hat{Z}_{i_1} \cdots \hat{Z}_{i_k}, \quad (7)$$

where the $J_{i_1, \dots, i_k}$ are real. The corresponding unitary operator is given by

$$\hat{U}_C(\gamma) = \exp(-i\gamma\hat{C}). \quad (8)$$

In the QAOA, the cost function or problem Hamiltonian is complemented by a so-called *driver* Hamiltonian typically inducing a unitary

$$\hat{U}_X(\beta) = \exp\left(-i\beta \sum_{i=1}^N \hat{X}_i\right) \quad (9)$$

comprised of single-qubit X rotation gates. These gates, often called transverse-field operators in the quantum-annealing literature [26], are responsible for driving transitions between computational basis states, i.e., solution states. The QAOA ansatz is then constructed from p pairs of these unitaries,

$$|\gamma, \beta\rangle = \hat{U}_X(\beta_p) \hat{U}_C(\gamma_p) \cdots \hat{U}_X(\beta_1) \hat{U}_C(\gamma_1) |+\rangle^{\otimes n}, \quad (10)$$

where γ, β are classical parameters. From the ansatz state, the expectation value of the cost function is computed as

$$F_p(\gamma, \beta) = \langle \gamma, \beta | \hat{C} | \gamma, \beta \rangle, \quad (11)$$

which serves as the loss function to be optimized, i.e., the aim is to determine the values of γ and β that maximize $F_p(\gamma, \beta)$. To quantify the performance of the QAOA, we introduce the approximation ratio

$$\mathcal{M}_p = \frac{\max_{\gamma, \beta} F_p(\gamma, \beta) - E_{\min}}{E_{\max} - E_{\min}}, \quad (12)$$

where $E_{\min}$ and $E_{\max}$ are the minimum and maximum values of $\hat{C}$, respectively. To determine the optimal angles for Eq. (12), we adopt the linear interpolation strategy proposed by Zhou *et al.* [27]:

$$[\gamma_{(p+1)}]_i = \frac{i-1}{p} [\gamma_{(p)}^*]_{i-1} + \frac{p-i+1}{p} [\gamma_{(p)}^*]_i, \quad (13)$$

where $\gamma_{(p+1)}$ serves as the initial point for the optimizer (we use COBYLA) with $p+1$ parameters for γ . All parameters $[\gamma_{(p+1)}]_i$ are then independently optimized to obtain $\gamma_{(p+1)}^*$. Here, $\gamma_{(p)}^*$ represents the optimized parameters for p . The subscript i denotes the individual parameters of each point. To compute the initial point $[\gamma_{(p+1)}]_i$, we interpolate between the corresponding values $[\gamma_{(p)}^*]_i$. For the initial parameters of $\beta_{(p+1)}$, we replace $\gamma \rightarrow \beta$ in Eq. (13). For $p \leq 3$, we use a random initialization strategy, generating 1000 random starting points and optimizing them with COBYLA to create a shape suitable for interpolation. Unless stated otherwise, this linear interpolation strategy is applied to all instances independently.

B. Mean-field approximate optimization

The MF-AOA [28] can be viewed the classical counterpart to the QAOA, where quantum evolution is replaced by classical spin dynamics via the mean-field approximation. A simplified version of it has indeed been considered in the early days of quantum optimization [29,30]. Here, we generalize the algorithm to arbitrary k -spin interactions as occurring in k -CSPs. Through the analogy of the fixed-angle QAOA with adiabatic quantum computation, we introduce the adiabatic Hamiltonian as

$$\hat{H}(s) = (1-s(t)) \sum_{i=1}^N \hat{X}_i + s(t) \hat{C}, \quad (14)$$

where $s(t) = t/T_f$, $s(t) \in [0, 1]$ is the scaled time. The mean-field form of this Hamiltonian is then given by

$$H(s) = (1-s(t)) \sum_{i=1}^N n_i^x(t) + s(t) \sum_{i_1, \dots, i_k} J_{i_1, \dots, i_k} n_{i_1}^z(t) \cdots n_{i_k}^z(t), \quad (15)$$

with $\mathbf{n}_i(t) = (n_i^x(t), n_i^y(t), n_i^z(t))^T$ being the classical spin vectors living on their respective Bloch spheres. The effective magnetization

$$m_{ij}(t) = \sum_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_{k-1}} J_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_{k-1}} \times n_{i_1}^z(t) \cdots n_{i_{j-1}}^z(t) n_{i_{j+1}}^z(t) \cdots n_{i_{k-1}}^z(t) \quad (16)$$

TABLE II. Standard deviations σ used in the random catalysts of Eq. (18) for different k and r .

k	3	4	5	6	7	8	9	10
$r = 0.5$	0.2	0.5	1.0	1.0	1.5	1.5	2.0	3.0
$r = 1.5$	0.5	1.0	1.0	1.0	1.5	1.5	1.5	3.0

then determines the classical equations of motion as

$$\begin{aligned}\dot{n}_i^x(t) &= -2s(t)m_i(t)n_i^y(t), \\ \dot{n}_i^y(t) &= 2s(t)m_i(t)n_i^x(t) - 2(1-s(t))n_i^z(t), \\ \dot{n}_i^z(t) &= 2(1-s(t))n_i^y(t).\end{aligned}\quad (17)$$

This set of $3N$ nonlinear ordinary differential equations can be solved numerically in an efficient way [31]. These equations can be derived, e.g., from the Heisenberg equations of motion via a product ansatz for the density matrix. We also remark that a similar product ansatz is used to derive near-optimal approximation ratios for the quantum generalization of Max-Cut in Ref. [32].

For $\mathbb{Z}_2$ -symmetric problem Hamiltonians, i.e., those without a *local* tensor J_{i_1} , the effective magnetization evidently remains zero throughout. To circumvent this issue, in the MF-AOA it is necessary to explicitly break the $\mathbb{Z}_2$ symmetry. A useful way of achieving this is to add a so-called catalyst Hamiltonian [33]

$$\hat{H}_{\text{cat}} = \sum_{i=1}^N \Lambda_i(t) \hat{Z}_i \quad (18)$$

to Eq. (14), where the coefficients $\lambda_i(t)$ need to obey the boundary conditions $\lambda_i(0) = \lambda_i(T_f) = 0$. For our purposes, it suffices to take $\Lambda_i(t) = \lambda_i s^2(t)(1-s(t))$, where the λ_i are chosen from a normal distribution $\mathcal{N}(0, \sigma^2)$. For the purpose of this paper, the standard deviations σ are chosen according to Table II. Note that this additional Hamiltonian term does not affect the initial and final spectrum, as it is zero at both the beginning and the end of the adiabatic evolution. The effective magnetization can now be written as

$$\begin{aligned}m_{ij}(t) &= \Lambda_{ij}(t) + \sum_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_{k-1}} J_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_{k-1}} \\ &\quad \times n_{i_1}^z(t) \cdots n_{i_{j-1}}^z(t) n_{i_{j+1}}^z(t) \cdots n_{i_{k-1}}^z(t).\end{aligned}\quad (19)$$

The equations of motion are then evolved adaptively up to a final time $T_f = 2^{15}$ in units of the inverse transverse field, and the solution bit string is obtained by projecting the z components of the final spin vectors $\mathbf{n}_i(T_f)$ to the poles of their Bloch spheres via

$$\sigma = (\text{sign}(n_1^z(T_f)), \dots, \text{sign}(n_N^z(T_f)))^T. \quad (20)$$

IV. RESULTS

A. QAOA approximation ratio as a function of the clause-to-variable ratio

The clause-to-variable ratio r is a critical property of specific instances of CSPs. As the clause-to-variable ratio increases, they transition from an underconstrained regime




FIG. 1. The normalized number of degenerate ground states P_0 defined in Eq. (21) against the clause-to-variable ratio r . We investigate P_0 for $k = 3, 4, 5, 6$ at $N = 15$ and average over 400 instances per k and r .

to an overconstrained regime. Generally speaking, overconstrained instances will be harder than underconstrained ones. This statement should be understood in a statistical sense: at high r , hard random instances will be more prevalent, while easy instances may still occur.

For small r , the instances are underconstrained, and typically not every variable appears in every clause; hence multiple solutions can exist. Such instances are easy to solve for a given algorithm. Additionally, the effective problem size N can be reduced. To construct underconstrained and overconstrained problems, we define

$$P_0 = 2^{-N} \sum_i E_{0,i} \quad (21)$$

as the ratio of the number of degenerate ground states and the total number of possible variable assignments, which is 2^N . Figure 1 illustrates how P_0 scales with r , e.g., for small $r = 0.1$, we have $P_0 \approx 0.15$, which means that 15% of all states correspond to the optimal solution, making these instances highly underconstrained. As r increases, the number of optimal solutions drops quickly. For $r = 1.5$, we already have $P_0 \approx 10^{-5}$ for odd $k = 3, 5$ and $P_0 \approx 10^{-4}$ for even $k = 4, 6$. For $k = 3, 5$, in particular, the instances approach unique solutions for high r . Such instances with only a few ground states are overconstrained because every variable is typically included in *at least* one clause. An exact threshold for under- and overconstrained random instances does not exist, but as a rule of thumb, we can state that $r < 1$ corresponds to less than one clause per variable. We also note that the decrease in the approximation ratio slows down for all p and k at high values of r .

The upper panel of Fig. 2 displays the results for the QAOA applied to Max- k XOR as a function of r . We fix $k = 5$ and investigate the ensemble-averaged approximation ratio for $p = 1, 2, 3, 4$; we observe a clear dependence on r . Crucially, $\mathcal{M}_p$, as defined in Eq. (12), exhibits a significant decrease with growing r , regardless of the specific Max- k XOR problem and the value of p . For high r , the instances become hard to solve. Importantly, higher values of p merely result in a vertical *shift* along the $\mathcal{M}_p$ axis and *not* in a systematically




FIG. 2. Ensemble-averaged approximation ratio of the QAOA at $N = 18$. For every r, k, p , we average over 10 000 random instances. Upper panel: $\mathcal{M}_p$ against the clause-to-variable ratio r for $k = 5$ and $p = 1, 2, 3, 4$. Lower panel: $\mathcal{M}_p$ as a function of the parameter r for $k \in [3, 10]$.

improved scaling. We remark that increasing r until all-to-all connected problem graphs are reached would lead to another underconstrained regime in which the approximation ratio of the QAOA should improve again. The Max-5XOR problem is used as a representative; we did not observe qualitatively different results for other values of k .

B. Mean-field approximation ratio as a function of the clause-to-variable ratio

To investigate how the mean-field approximation ratio depends on the clause-to-variable ratio r , we apply the MF-AOA to random instances of the Max- k XOR problem across the entire range of possible r values for fixed N and k . For this problem, the number of clauses ranges from 1 to $2^{\binom{N}{k}}$. We plot the relative energy deviation $(E^* - E_0)/E_0$ as a function of r , as this metric quantifies how close the system is to the true ground state of the problem Hamiltonian. It also serves as a useful hardness measure for MF-AOA across different problem instances.

In the decision variant of the k -XOR problem, a well-known phase transition occurs at $M = N$, corresponding to $r = 1$. Instances with $r < 1$ or $r \gg 1$ are computationally less challenging to solve compared to those around $r = 1$. For the optimization variant, the results presented in Fig. 3 demonstrate distinct performance regimes. In the underconstrained regime ($r < 1$), where the associated problem graph is sparse, the system remains closer to the ground state. Conversely, in the hard regime ($r > 1$) the ensemble-averaged approximation ratio plateaus, becoming largely independent of the clause-to-variable ratio.




FIG. 3. The relative energy deviation $(E^* - E_0)/E_0$ plotted as a function of r for $k = 3$ and $N = 8$, covering the entire range of possible r . Each point represents the average over 500 randomly generated instances. The dotted line indicates the phase transition from easy to hard for the k XOR decision problem.

Since the approximation ratio in the hard regime does not vary significantly with r , we fix $r = 0.5$ for the easy case and $r = 1.5$ for the hard case in the remainder of the paper.

C. QAOA approximation ratio as a function of k

In this section, we investigate how the ensemble-averaged approximation ratio of the QAOA changes when applied to random Max- k XOR instances for large k and fixed $p = 4$. The results are shown in the lower panel of Fig. 2. For large r , we observe that $\mathcal{M}_4$ decreases with k , providing a first indication that random Max- k XOR becomes harder and harder to solve for the QAOA. This is in contrast to the results of Marwaha and Hadfield [19] for uniform instances, who observed an increase in the QAOA's approximation ratio with k . In Fig. 2, it is evident that the standard deviation is larger for small r than for large r for all values of k . This is due to a higher variability in the difficulty of random instances when r is closer to zero. We also observe the impact of delocalization on the QAOA. To illustrate the delocalization effect, we choose a constant value for k . Upon selecting a random assignment of the variables for a given instance, flipping a single bit induces a change in the cost function. If we now calculate this change of cost for an instance with $k + 1$, we observe an increase as compared to k . The QAOA is of local structure through its single-qubit X -driver, leading to a decrease in the ensemble-averaged approximation ratio as k increases.

D. QAOA approximation ratio at high circuit depth

This section analyzes how the QAOA performs on hard-to-solve random instances at large p . As previously discussed, solving random instances with large r is more difficult than those with small r , regardless of the value of k , unless the problem-related graph is nearly all to all connected, i.e.,

$$r \approx \binom{N}{k} N^{-1}. \quad (22)$$

In Fig. 4, we plot the results for $\mathcal{M}_p$ with $p \in \{1, \dots, 30\}$ and $k \in \{3, \dots, 10\}$. As k increases, solving Max- k XOR problems with the QAOA becomes significantly more challenging. This trend suggests that an exponentially larger number of layers is required to reach the ground state for high values of k . Extending the findings of Basso *et al.* [23], who




FIG. 4. The ensemble-averaged approximation ratio $\mathcal{M}_p$ against p for Max- k XOR with $k \in \{3, \dots, 10\}$. We set $N = 18$ and $r = 1.5$ to investigate hard-to-solve instances for the QAOA. For each p and k , we average over 100 random instances. The inset shows the collapsed logarithmic fits to the data (i.e., subtracting each constant shift). The solid (red) line is drawn for illustration with a coefficient of $c = 0.06$.

demonstrated the inefficiency of the QAOA for solving Max- k XOR problems with *even* values of k , we thus provide numerical evidence that this assertion seems to hold for *all* values of k .

Figure 5 presents the average results for the calculated parameters. The standard deviation remains low for small values of p and increases for higher p . Beyond $p = 30$, the linear-interpolation method shows no further improvement. In contrast, the Fourier-based approach proposed by Zhou *et al.* [27] demonstrates promising performance for $p > 30$, albeit at a significantly higher computational cost.

To assess the limitations of the QAOA, we estimate the number of layers p required to an approximation ratio of 90% and above for challenging problem instances in Fig. 6. Specifically, we fit the data from Fig. 4 to a logarithmic ansatz (see inset of Fig. 4). To improve the accuracy of the fit for higher p values, we exclude the first data point ($p = 1$). By extrapolating the resulting fits to high p , we can estimate the p value at which a given $\mathcal{M}_p$ is achieved. For lower values of k (e.g., $k = 3$), a circuit depth of $p \approx 50$ is required to reach 99% of the ground-state energy. For larger k (e.g., $k = 10$), this increases to approximately $p = 770$ layers. These findings highlight the limitations of the QAOA on current hardware, where gate fidelities of around 99.5% make it impractical to implement circuit depths as large as $p = 50$.




FIG. 5. The ensemble-averaged optimal parameters $\bar{\beta}$, $\bar{\gamma}$ for the QAOA with $p = 10$ (left plot) and $p = 20$ (right plot). The averages are taken over the ensembles of Fig. 4 with 100 instances per k and p .

E. QAOA approximation ratio at larger N

In the previous chapters, we set $N = 18$ to be able to simulate the QAOA with reasonable computational effort. In this section, we now aim to investigate how the QAOA performs for larger system sizes N . To compute the values of $\mathcal{M}_p$, we use the optimal angles from Fig. 5 found at $N = 18$ and apply these parameters to instances with $N > 18$ for each k , i.e., we do not perform the parameter optimization for $N > 18$. Exemplary results for $p = 20$ are shown in Fig. 7.

We observe that, for each k , the ensemble-averaged approximation ratio of the QAOA remains nearly constant across all investigated values of N , suggesting a certain degree of universality of our optimized angles. These findings are consistent with those of Farhi *et al.* [13], who demonstrated that parameters originally derived for the infinite-size Sherrington-Kirkpatrick model were also applicable to finite-size instances.

To investigate the transition in behavior from small ($N \approx 10$) to larger system sizes ($N \approx 20$), we analyze the ensemble-averaged probability distribution of the QAOA across states ranging from the ground state to the eighth




FIG. 6. The required p to reach various ensemble-averaged approximation ratio thresholds $\mathcal{M}_p$ as a function of k . The value of p is determined by fitting a logarithmic function to Fig. 4. The dotted curves represent exponential and polynomial fits, along with the root-mean-squared error (RMSE) of each fit.




FIG. 7. Ensemble-averaged approximation ratio $\mathcal{M}_{20}$ against system size N . Per N and Max- k XOR, we average over 100 instances.

excited state. These distributions, plotted in Fig. 8 for $k = 3$, reveal a trend that closely follows a Poisson distribution for different values of N . For small systems, the distribution exhibits concentration of states near the ground state. As N increases, this concentration diminishes, accompanied by a noticeable shift toward higher excited states.

F. MF-AOA vs QAOA

In this section, we compare the ensemble-averaged approximation ratio of the QAOA with the MF-AOA by analyzing $\mathcal{M}_p$ across random Max- k XOR instances. This provides insight into how both algorithms compare across varying problem complexities and circuit depths.




FIG. 8. Concentration of the ensemble-averaged probability distributions of the final states reached by the QAOA for $k = 3$. The state indices are ordered by energy, with 0 indicating the ground state. Above the eighth excited state, the probabilities become negligible. Note that the Poisson distributions are *not* fitted but simply evaluated at the mean state index.




FIG. 9. The distribution of $\mathcal{M}_p$ against k for Max- k XOR with $k \in \{3, \dots, 10\}$. We set $N = 18$ and $r = 1.5$ to compare hard-to-solve instances for the QAOA against mean field. While the shaded area indicates the outcome distribution, the points represent averages over 100 random instances.

In Fig. 9, we plot the distribution of $\mathcal{M}_p$ as a function of k , with a fixed $r = 1.5$ and circuit depth $p = 30$ for the QAOA. This allows us to compare both algorithms' ensemble-averaged approximation ratio for increasing values of k , which corresponds to the number of literals in each clause of the Max- k XOR problem. We observe that the MF-AOA exhibits a broadened distribution of $\mathcal{M}_p$, indicating greater variance of the MF-AOA across random instances. Despite the wider distribution, the MF-AOA consistently matches the QAOA on average. Notably, the QAOA's ensemble-averaged approximation ratio relative to MF-AOA improves as k increases. It seems likely that both the increased MF-AOA variance and the slightly lower approximation ratio at larger k can be addressed by sampling more random catalysts according to Eq. (18).

In Fig. 10, we plot $\mathcal{M}_p$ against p for four distinct combinations of k and r , representing both easy and hard cases, as well as low and high values of k . We compare the ensemble-averaged approximation ratio of the QAOA (in dependence on the circuit depth) with that of the MF-AOA.




FIG. 10. Ensemble-averaged approximation ratio $\mathcal{M}_p$ against p for the QAOA (red solid-dotted line) and the MF-AOA (blue solid horizontal line) across four different combinations of k and r . Each line represents an average over 100 random instances, with the shaded areas indicating the standard deviation.




FIG. 11. Ensemble-averaged approximation ratio $\mathcal{M}_p$ against system size N for the MF-AOA and the QAOA ranging from $p = 1$ to $p = 25$ for $k = 3$. Each line represents an average over 100 random instances. The clearly visible convergence of the QAOA's approximation ratio with increasing p raises the question whether anything is to be gained by going to deeper circuits.

For hard instances ($r = 1.5$), we observe that the QAOA approaches and eventually surpasses the MF-AOA's approximation ratio as the circuit depth increases. However, the QAOA requires relatively high circuit depths to match the $\mathcal{M}_p$ value achieved by the classical benchmark for both low and high k values. Notably, for the easy case ($r = 0.5$) at $k = 5$, the QAOA matches the classical approximation ratio at relatively low circuit depth $p = 17$ and even surpasses it at large depths. However, the MF-AOA results require hardly any computational effort, while the angle optimization going into the QAOA results consumes considerable resources (besides being problematic on noisy quantum hardware [6,8–10]).

To identify a region where the QAOA outperforms its classical counterpart (MF-AOA), we compare both algorithms in Fig. 11 across varying system sizes and circuit depths. Our findings show that for small systems ($N \approx 8$), the QAOA outperforms the MF-AOA. However, as the system size increases ($N \approx 20$), the MF-AOA surpasses the QAOA. Interestingly, for small values of p , the QAOA's approximation ratio remains nearly independent of the system size.

V. CONCLUSION AND OUTLOOK

Our numerical analysis demonstrates that the *relative* improvement of the QAOA's ensemble-averaged approximation ratio applied to Max- k XOR declines as both the parameter k and the number of layers p increase. Our results provide explicit numerical evidence detailing the approximation ratio of the QAOA across different Max- k XOR instances (for odd and even k) and tentatively predict the circuit depth required to achieve an approximation ratio $\mathcal{M}_p$ exceeding 90%. Additionally, we observe that the parameters optimized for a fixed problem size ($N = 18$) appear to generalize well to larger systems. These results suggest that achieving solutions close to optimal will require significantly higher circuit depths than presently available. These findings align with the analytic observations of Basso *et al.* [23], who noted challenges in using the QAOA to approximate ground state energies for sparse Max- k XOR problems (even k), especially for larger k , with shallow circuits.

Interestingly, the same approximation-ratio decrease in k is experienced by our classical benchmark, seemingly resulting in a slight average advantage for the QAOA at larger k ; see Fig. 9. It would be intriguing for future work to look into the behavior of other classical solvers in the same limit. We also point out that our mean-field results should be considered a *lower bound* on the possible ensemble-averaged approximation ratio, as there is, at the very least, the option to sample more instantiations of our random catalyst from Eq. (18). Since any proper time-to-solution analysis [3] has to take into account optimizing the QAOA angles, there is a lot of leeway for classical benchmarks in this regard. Overall, we provide evidence that the QAOA, as applied to random Max- k XOR problems with odd and even k , might not show an advantage over classical algorithms.

ACKNOWLEDGMENTS

The authors acknowledge funding, support, and computational resources from the German Aerospace Center (DLR) and Forschungszentrum Jülich. Furthermore, we acknowledge funding from QSolid via the Federal Ministry of Education and Research (BMBF). We also acknowledge funding from AQUAS and QUASIM, both via the Federal Ministry of Economic Affairs and Climate Action (BMWK). Finally, we are grateful for useful conversations with S. Hadfield, S. Gharibian, A. Ciani, and D. Bagrets.

- [1] T. Hogg, Quantum search heuristics, *Phys. Rev. A* **61**, 052311 (2000).
- [2] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [3] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [4] S. Boulebnane and A. Montanaro, Solving Boolean satisfiability problems with the quantum approximate optimization algorithm, *PRX Quantum* **5**, 030348 (2024).
- [5] A. Montanaro and L. Zhou, Quantum speedups in solving near-symmetric optimization problems by low-depth QAOA, [arXiv:2411.04979](https://arxiv.org/abs/2411.04979).
- [6] Y. Quek, D. Stilck França, S. Khatri, J. J. Meyer, and J. Eisert, Exponentially tighter bounds on limitations of quantum error mitigation, *Nat. Phys.* **20**, 1648 (2024).
- [7] A. Abbas, A. Ambainis, B. Augustino, A. Bäertschi, H. Buhrman, C. Coffrin, G. Cortiana, V. Dunjko, D. J. Egger, B. G. Elmegreen, N. Franco, F. Fratini, B. Fuller, J. Gacon, C. Gondiulea, S. Gribbling, S. Gupta, S. Hadfield, R. Heese, G. Kircher *et al.*, Challenges and opportunities in quantum optimization, *Nat. Rev. Phys.* **6**, 718 (2024).
- [8] A. A. Mele, A. Angrisani, S. Ghosh, S. Khatri, J. Eisert, D. S. França, and Y. Quek, Noise-induced shallow circuits and absence of barren plateaus, [arXiv:2403.13927](https://arxiv.org/abs/2403.13927).

- [9] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, E. R. Anschuetz, and Z. Holmes, Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing, [arXiv:2312.09121](#).
- [10] S. Kazi, M. Larocca, M. Farinati, P. J. Coles, M. Cerezo, and R. Zeier, Analyzing the quantum approximate optimization algorithm: Ansätze, symmetries, and Lie algebras, [arXiv:2410.05187](#).
- [11] K. Marwaha, Local classical max-cut algorithm outperforms $p = 2$ QAOA on high-girth regular graphs, *Quantum* **5**, 437 (2021).
- [12] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for MaxCut: A fermionic view, *Phys. Rev. A* **97**, 022304 (2018).
- [13] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, *Quantum* **6**, 759 (2022).
- [14] J. Håstad, Some optimal inapproximability results, *J. ACM* **48**, 798 (2001).
- [15] G. E. Crooks, Performance of the quantum approximate optimization algorithm on the maximum cut problem, [arXiv:1811.08419](#).
- [16] M. X. Goemans and D. P. Williamson, .879-approximation algorithms for MAX CUT and MAX 2SAT, in *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '94 (Association for Computing Machinery, New York, 1994), p. 422–431.
- [17] K. Makarychev and Y. Makarychev, *Approximation Algorithm for Non-Boolean Max k-CSP* (Springer, Berlin, Berlin, 2012), pp. 254–265.
- [18] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, Reachability deficits in quantum approximate optimization, *Phys. Rev. Lett.* **124**, 090504 (2020).
- [19] K. Marwaha and S. Hadfield, Bounds on approximating Max k XOR with quantum and classical local algorithms, *Quantum* **6**, 757 (2022).
- [20] J. Hirvonen, J. Rybicki, S. Schmid, and J. Suomela, Large cuts with local algorithms on triangle-free graphs, [arXiv:1402.2543](#).
- [21] C.-N. Chou, P. J. Love, J. S. Sandhu, and J. Shi, Limitations of local quantum algorithms on random Max- k -XOR and beyond, [arXiv:2108.06049](#).
- [22] W.-K. Chen, D. Gamarnik, D. Panchenko, and M. Rahman, Suboptimality of local algorithms for a class of max-cut problems, *Ann. Probab.* **47**, 1587 (2019).
- [23] J. Basso, D. Gamarnik, S. Mei, and L. Zhou, 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)* (IEEE, Denver, CO, USA, 2022).
- [24] D. Gamarnik and A. Jagannath, The overlap gap property and approximate message passing algorithms for p -spin models, [arXiv:1911.06943](#).
- [25] D. Gamarnik, The overlap gap property: A topological barrier to optimizing over random structures, *Proc. Natl. Acad. Sci. USA* **118**, e2108492118 (2021).
- [26] M. H. S. Amin, Consistency of the adiabatic theorem, *Phys. Rev. Lett.* **102**, 220401 (2009).
- [27] 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).
- [28] A. Misra-Spieldenner, T. Bode, P. K. Schuhmacher, T. Stollenwerk, D. Bagrets, and F. K. Wilhelm, Mean-field approximate optimization algorithm, *PRX Quantum* **4**, 030335 (2023).
- [29] J. A. Smolin and G. Smith, Classical signature of quantum annealing, *Front. Phys.* **2** (2014).
- [30] S. W. Shin, G. Smith, J. A. Smolin, and U. Vazirani, How “quantum” is the D-wave machine? [arXiv:1401.7087](#).
- [31] T. Bode and F. K. Wilhelm, Adiabatic bottlenecks in quantum annealing and nonequilibrium dynamics of paramagnons, *Phys. Rev. A* **110**, 012611 (2024).
- [32] S. Gharibian and O. Parekh, *Almost Optimal Classical Approximation Algorithms for a Quantum Generalization of Max-Cut* (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2019).
- [33] R. Ghosh, L. A. Nutricati, N. Feinstein, P. A. Warburton, and S. Bose, Exponential speed-up of quantum annealing via n -local catalysts, [arXiv:2409.13029](#).