

PAPER • OPEN ACCESS

Grover-QAOA for 3-SAT: quadratic speedup, fair-sampling, and parameter clustering

To cite this article: Zewen Zhang *et al* 2025 *Quantum Sci. Technol.* **10** 015022

View the [article online](#) for updates and enhancements.

You may also like

- [Quantum approximate optimization algorithm in non-Markovian quantum systems](#)
Bo Yue, Shibe Xue, Yu Pan et al.
- [Convergence of digitized-counterdiabatic QAOA: circuit depth versus free parameters](#)
Mara Vizzuso, Gianluca Passarelli, Giovanni Cantele et al.
- [To quantum or not to quantum: towards algorithm selection in near-term quantum optimization](#)
Charles Moussa, Henri Calandra and Vedran Dunjko

Quantum Science and Technology



PAPER

OPEN ACCESS

RECEIVED
14 February 2024

REVISED
24 July 2024

ACCEPTED FOR PUBLICATION
21 October 2024

PUBLISHED
4 November 2024

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Grover-QAOA for 3-SAT: quadratic speedup, fair-sampling, and parameter clustering

Zewen Zhang^{1,*} , Roger Paredes², Bhuvanesh Sundar^{1,5} , David Quiroga³, Anastasios Kyrillidis³,
Leonardo Duenas-Osorio², Guido Pagano^{1,4} and Kaden R A Hazzard^{1,4}

¹ Department of Physics and Astronomy, Rice University, Houston, TX 77005, United States of America

² Department of Civil and Environmental Engineering, Rice University, Houston, TX 77005, United States of America

³ Department of Computer Science, Rice University, Houston, TX 77005, United States of America

⁴ Smalley-Curl Institute, Rice University, Houston, TX 77005, United States of America

⁵ Current address: Rigetti Computing, Berkeley, California 94710, United States of America.

* Author to whom any correspondence should be addressed.

E-mail: zewen.zhang@rice.edu

Keywords: QAOA, speedup, SAT-solver, Grover mixer, parameter clustering, 3-SAT, Max-SAT

Abstract

The SAT problem is a prototypical NP-complete problem of fundamental importance in computational complexity theory with many applications in science and engineering; as such, it has long served as an essential benchmark for classical and quantum algorithms. This study shows numerical evidence for a quadratic speedup of the Grover Quantum Approximate Optimization Algorithm (G-QAOA) over random sampling for finding all solutions to 3-SAT (All-SAT) and Max-SAT problems. G-QAOA is less resource-intensive and more adaptable for these problems than Grover's algorithm, and it surpasses conventional QAOA in its ability to sample all solutions. We show these benefits by classical simulations of many-round G-QAOA on thousands of random 3-SAT instances. We also observe G-QAOA advantages on the IonQ Aria quantum computer for small instances, finding that current hardware suffices to determine and sample all solutions. Interestingly, a single-angle-pair constraint that uses the same pair of angles at each G-QAOA round greatly reduces the classical computational overhead of optimizing the G-QAOA angles while preserving its quadratic speedup. We also find parameter clustering of the angles. The single-angle-pair protocol and parameter clustering significantly reduce obstacles to classical optimization of the G-QAOA angles.

1. Introduction

Combinatorial optimization problems are of great interest for quantum computing and have been explored using amplitude amplification [1], quantum annealing [2], and variational techniques, such as the variational quantum eigensolver [3] and the quantum approximate optimization algorithm (QAOA) [4]. QAOA is suited for combinatorial optimization due to its simple protocol [5] and its ability as a hybrid variational algorithm to balance quantum and classical resources to adapt to the constraints of practical quantum computers.

QAOA has been applied to various combinatorial optimization tasks, with some even demonstrating its capabilities on state-of-the-art quantum devices [6–14]. Recent explorations on combinatorial problems such as MaxCut, Ising, spin glass, and k -SAT focused on the classical optimization of QAOA – including employing strategies centered on fixed [15] or smooth angles [16], finding more intricate angle structures [17–22], or using empirical optimization strategies [23–29].

Among these problems, the Boolean satisfiability problem (SAT) is central for its practical importance. It is an NP-complete problem, encapsulating the challenges inherent to all NP problems [30–32]. SAT problems have applications spanning fields such as hardware and software design [33–36], cryptography [37], bioinformatics [38], network reliability [39] and more [40]. k -SAT problems are defined in terms of Boolean expressions $\phi(m, n)$ consisting of m conjoined clauses, each of which is a disjunction of k

Boolean variables, over a set of n Boolean variables. The density, defined as the ratio $d = m/n$, is an important characteristic that influences the hardness of the problem.

Numerous SAT applications require fair sampling of solutions, namely the capability for SAT solvers to identify comprehensive solution sets rather than just one, a problem known as All-SAT [41] and related to model counting [42]. There has been growing interest in utilizing quantum algorithms to address SAT problems, leveraging techniques like quantum annealing [43–45] and amplitude amplification [46–48]. While quantum annealing-based methods do not consistently ensure speedup, amplitude amplification protocols based on the Grover algorithm [49, 50] guarantee quadratic speedup over random guessing. However, methods rooted in the amplitude amplification necessitate implementing an oracle, demanding ancilla qubits proportionate to the total clause count [51] or increasing the circuit depth (see appendix A). Alternative methods that reduce the number of ancillas or depth for the oracle are desirable.

In this work, we adopt a recently proposed QAOA variant (G-QAOA) in which the mixer operation, usually implemented with a global qubit rotation, is replaced by a Grover mixer that guarantees the uniform sampling of the solution space [8, 52, 53]. This has been recently applied to SAT problems [54, 55] to find solutions, but its scaling with problem size has not yet been studied. Although some research indicates a potential speedup, the numerical demonstrations are either unclear [55, 56] on scaling or not applied to NP-hard problems [52].

Here, we present compelling evidence of the quadratic speedup over random guessing of G-QAOA for finding all solutions to 3-SAT problems. We analyze thousands of random 3-SAT instances up to 26 Boolean variables with densities spanning from 2 to 8, encompassing critical thresholds [57, 58]. Unlike previous classical optimization schemes that simultaneously determine angles for all QAOA rounds, we show that maintaining a single angle pair throughout all G-QAOA rounds retains a quadratic speedup over random guessing. This approach dramatically reduces the classical optimization cost. These angle pairs cluster in a small parameter range across all 3-SAT instances. In addition, we notice that the speedup of G-QAOA shows no dependence on the clause density, which probably leads to an advantage over classical algorithms for hard instances. A comparison of scaling with the leading classical algorithm *Elim-Count* is discussed in the appendix F.

This single angle pair approach obtains a quadratic speedup not only for finding solutions of 3-SAT but also for finding a configuration of the problem variables that satisfies as many clauses as possible, referred to as ‘Max-SAT’ [59, 60]. Max-SAT is particularly important at densities higher than the so-called critical density when typically no solution to the 3-SAT problems exists [61]. Max-SAT is a significant combinatorial optimization problem in its own right, with applications in areas such as system design [62, 63], machine learning [64, 65] and bioinformatics [66, 67].

Finally, we find that G-QAOA guarantees a 10-fold fewer qubits and circuit depth at the critical density compared to the standard Grover algorithm, when implemented on an all-to-all connected trapped-ion machine for 3-SAT problems (see appendix A). This advantage increases proportionally to the density. Thus, reducing the required quantum resources enables us to run this algorithm on state-of-the-art IonQ quantum computers. These demonstrations provide additional validation for potential future applications beyond synthetic benchmarks and further corroborate the fair-sampling attribute of G-QAOA.

The paper is structured as follows: section 2 delineates the 3-SAT problems under examination and describes the G-QAOA methodologies implemented. Section 3 numerically demonstrates the quadratic speedup of G-QAOA on 3-SAT problems and presents evidence that the single-angle-pair approach maintains the speedup. Section 4 shows results from IonQ’s Aria machines, and section 5 outlines this work’s conclusions and outlook.

2. Background: 3-SAT and G-QAOA

2.1. 3-SAT

The 3-SAT problem refers to the particular case of Boolean formulas in *conjunctive normal form* (CNF) in which every clause i has three literals $l_{i_1}, l_{i_2}, l_{i_3}$ (see appendix B). This problem has long been a standard combinatorial setting for algorithmic and theoretical studies. The density d is a crucial indicator characterizing the solution space and its computational hardness. For a 3-SAT instance with small d , the problem is under-constrained and likely to be satisfiable, while at large d , it is over-constrained and likely to be unsatisfiable. The computational difficulty for a broad class of state-of-the-art classical algorithms is maximized at a critical density $d_c \approx 4.26$ [61].

Given a 3-SAT CNF instance ϕ , the problem can be translated into a qubit representation through a problem Hamiltonian H_P [32]

$$H_P = \sum_{i=1}^m H_C^{(i)},$$

where each local Hamiltonian $H_C^{(i)}$ is associated to the i th clause $(l_{i_1} \vee l_{i_2} \vee l_{i_3})$ of the instance and defined as

$$H_C^{(i)} \equiv \frac{1}{8} (\mathbb{1}_{i_1} - s_{i_1} Z_{i_1}) (\mathbb{1}_{i_2} - s_{i_2} Z_{i_2}) (\mathbb{1}_{i_3} - s_{i_3} Z_{i_3}), \quad (1)$$

where $\mathbb{1}_{i_j}$ is an identity operator on the i_j th qubit, Z_{i_j} is the Pauli operator $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$ acting on the i_j th qubit, and $s_{i_j} \in \{\pm 1\}$ denotes the sign of the literal l_{i_j} (-1 if negated, and $+1$ otherwise). Therefore, given an assignment of the Boolean variables $\tau \in \{0, 1\}^{\otimes n}$ corresponding to a quantum state $|\tau\rangle$, the local energy E_i of the i th clause is

$$E_i = \langle \tau | H_C^{(i)} | \tau \rangle = \begin{cases} 0, & \text{if } \tau \text{ satisfies } (l_{i_1} \vee l_{i_2} \vee l_{i_3}), \\ 1, & \text{otherwise.} \end{cases}$$

Thus, the expectation value of $H_P(\tau)$ equals the number of unsatisfied clauses of the instance ϕ under assignment τ . If $|\tau\rangle$ is a zero-energy ground state of H_P , then τ is a solution of ϕ .

2.2. G-QAOA

This section first reviews the G-QAOA with arbitrary variational angles. Then, it discusses the variant we introduce in this paper, the single-angle-pair G-QAOA, where the angles are equal across all G-QAOA rounds.

G-QAOA. For the problems we consider, G-QAOA alternately applies two sets of parameterized unitaries generated by H_P and a mixing operator,

$$H_M = \prod_{i=1}^n (\mathbb{1}_i + X_i) / 2 = |+\rangle\langle +|, \quad (2)$$

to the product initial state $|+\rangle = \otimes_i (|0\rangle_i + |1\rangle_i) / \sqrt{2}$ to boost the probability of solutions. Here, X_i is the Pauli operator $X = |1\rangle\langle 0| + |0\rangle\langle 1|$ acting on the i th qubit. This results in the state

$$|\beta, \gamma\rangle \equiv e^{-i\beta_p H_M} e^{-i\gamma_p H_P} \dots e^{-i\beta_1 H_M} e^{-i\gamma_1 H_P} |+\rangle, \quad (3)$$

where $\beta = (\beta_1, \dots, \beta_p)$ and $\gamma = (\gamma_1, \dots, \gamma_p)$ are real parameters, and p is the number of QAOA rounds. The Grover mixing unitary $e^{-i\beta H_M}$ for n qubit problems can be decomposed with the help of $(n-3)$ ancilla qubits, as shown in figure 8 in appendix E.

The β and γ are optimized as described in appendix C to minimize a cost function

$$C \equiv \langle \beta, \gamma | H_P | \beta, \gamma \rangle, \quad (4)$$


which can be evaluated by averaging over measurements on a quantum computer.

G-QAOA alleviates a significant obstacle to finding all solutions that confronts the usual QAOA based on the transverse-field mixer [5] (X-QAOA): $H_X = \sum_{i=1}^n X_i$. In particular, the X-QAOA fails to sample all solutions equally, often exhibiting exponentially small weights on some solutions [8, 68].

Single-angle-pair G-QAOA. Optimizing many parameters in variational quantum algorithms poses difficulties to classical optimizers [69–71]. Hence, this paper also considers a protocol in which the same pair of (β, γ) angle values is used across all rounds of the G-QAOA. Accordingly, the protocol in equation (3) becomes

$$|\beta, \gamma, p\rangle \equiv \prod_{j=1}^p e^{-i\beta H_M} e^{-i\gamma H_P} |+\rangle \quad (5)$$

for p -round G-QAOA, reducing the number of variational parameters from $2p$ to 2. This simplification significantly alleviates the classical optimization demands (see appendix C inherent to the QAOA protocol. Tests (see section 3) across different problem sizes comparing the single-angle-pair G-QAOA and full G-QAOA reveal that, despite this constraint, the single-angle-pair G-QAOA exhibits a similar quadratic speedup as the full G-QAOA.




3. Quadratic speedup and parameter clustering

This section assesses the performance of G-QAOA and its single-angle-pair variant for 3-SAT problems. We produce random instances $\phi(m, n)$ of the 3-SAT problem for two density ranges, roughly below and above the satisfiability threshold: (i) $d \in (2, 4.5)$, in which the large majority of instances are satisfiable, and (ii) $d \in (4.6, 8)$ in which most are unsatisfiable. In the former, we keep only satisfiable instances; in the latter, we keep only unsatisfiable ones to highlight potential distinctions between the two cases. The two separate density ranges chosen here help to avoid cluttering in plotting, and due to the finite-size effect, there are satisfiable instances when the density is slightly above the threshold. In generating random instances, if an instance contains duplicate clauses, only one of the repeating clauses is retained. We consider problem sizes from $n = 10$ to 26, with at least 100 random 3-SAT instances for each $n < 23$, and at least 50 random instances for each $n \geq 23$ (simulation details in appendix C).

3.1. Evidence of quadratic speedup

Figure 1(a) shows the number of G-QAOA rounds required to create a quantum state with at least $S = 50\%$ probability of being a solution, where each point corresponds to a sampled 3-SAT instance. While the time to find a solution with random guessing would scale as $1/P$, where P is the fraction of Boolean variable assignments that satisfy ϕ , figure 1(a) shows that the number of G-QAOA rounds scales only as $1/\sqrt{P}$. Appendix C shows similar results for other success probabilities, $S = 25\%$ and 75% . The same trend and quadratic speedup hold for results obtained using the simpler single-angle-pair G-QAOA, also shown in figure 1(a). Using the single-angle-pair ansatz, figure 1(b) shows the same scaling for the as-hard or harder problem of finding the variable assignments with the maximum number of satisfied clauses (Max-SAT) when instances are not fully satisfiable. In this part, we record D is the fraction of Boolean variable assignments that give the maximal number of satisfied clauses to ϕ .



The quadratic speedup observed in the number of QAOA rounds is maintained after accounting for the implementation cost of the quantum circuits (number of Mølmer-Sørensen gates [72, 73]), as shown in figure 1(c).

The same data points of figures 1(a) and (b) are plotted in figure 1(d) as a function of density d . The results up to 26 Boolean variables indicate that, for the satisfiable 3-SAT problem, the number of G-QAOA rounds required to boost the probability of solution increases linearly as d increases. The number of rounds is roughly independent of d for unsatisfiable instances. Specifically, the region of hard instances around the critical density, as indicated by the dashed line in figure 1(d), shows no difference in scaling than other regions.

3.2. Parameter clustering

Tracking the angle pairs (β, γ) associated with these 3-SAT problems, a clustering pattern is observed, as shown in figure 2. For all 3-SAT instances in figure 1(a), the distribution of the angle pairs is notably concentrated, contingent upon the size of instances and the number of clauses. Similar clustering behaviors are also observed for the Max-SAT problems (figure 1(b)) when applying the single-angle-pair protocol, as shown in appendix D.

As the number of clauses increases, the distribution of angle pairs concentrates, especially the γ angle. This observation implies that for future handling of large, practically challenging instances, the parameter search space will be significantly narrowed. Parameters within this restricted region are likely to yield high-quality approximate results. Consequently, the classical optimization cost will be sharply reduced, providing an additional advantage besides the single-angle-pair protocol.

4. G-QAOA on IonQ hardware

To benchmark the feasibility of G-QAOA on NISQ devices, we executed sample 3-SAT problems⁶ using IonQ’s Aria machine. The demonstrations were executed at optimal G-QAOA parameters that were pre-optimized classically. The 4-qubit (figure 3) and the 5-qubit (figure 4) examples show that the G-QAOA enhances solution probability over random guessing, with each solution having a comparable relative frequency. In contrast, the X-QAOA enhances solution probabilities with relative frequencies varying significantly between solutions.

The experimental results qualitatively agree with the noiseless numerical simulations. For larger problems ($n > 5$) and rounds of G-QAOA ($p \geq 3$), the results depart significantly from the noiseless simulation, especially for G-QAOA which requires considerably more entangling gates (Mølmer-Sørensen) than its X-QAOA counterpart at the tested densities. However, the relative difference in gate cost will decrease at higher densities. Specifically, given n Boolean variables and the number of clauses m , X-QAOA requires $6m$ entangling gates per round. In contrast, G-QAOA requires $6m + 5(2n - 5)$ per round, where the additional $5(2n - 5)$ entangling gates come from the decomposition of the Toffoli gates needed to implement

⁶ 4-qubit example: $\{[-0, +2, +3], [+0, +2, -3], [-1, +2, -3], [-1, -2, -3], [-1, -2, +3], [+1, +2, -3], [+0, +2, +3], [-0, +1, -3]\}$; 5-qubit example: $\{[+1, -2, -3], [-1, -3, +4], [+0, -2, -4], [-0, -2, +3], [+2, +3, +4], [-0, +1, +2], [+0, -2, 4], [-1, +2, -4]\}$. This notation lists the set of conjunctive clauses, where the notation for each clause is that minus signs imply the negation of a Boolean variable, and the number indexes the Boolean variable. For example $[-0, +2, +3] = (\bar{x}_0 \vee x_2 \vee x_3)$.




Figure 3. Results from IonQ Aria machines for an $n = 4$ (4 qubits + 1 ancilla) problem. Here, 2000 shots are taken for each panel. States displaying pink bins represent solutions, while the rest give non-solutions. The dashed line indicates the probability of finding a solution by random guessing. Open bins give noiseless simulated results for reference. The state index is the decimal number corresponding to the binary string for the quantum state.




Figure 4. Results from IonQ Aria machines for an $n = 5$ (5 qubits + 2 ancillas) problem. Here, 4000 shots are taken for each panel. States displaying pink bins represent solutions, while the rest give non-solutions. The dashed line indicates the probability of finding a solution by random guessing. Open bins give noiseless simulated results for reference. The state index is the decimal number corresponding to the binary string for the quantum state.

the Grover mixer in equation (2). Additionally, G-QAOA requires $n - 3$ ancillas to implement the Grover mixer (see appendix E for details).

Simulations and demonstration on Aria highlight two aspects of G-QAOA's efficacy, whose relative importance differs in different applications: the total probability of solutions and how uniformly the solutions are sampled. We utilize two metrics to assess the latter: (i) the number of shots needed to observe all solutions [8]; and (ii) the number of draws from the experimental solution distribution required to reject (by χ^2 -tests) the null hypothesis of uniform solution sampling, taking the median from multiple rejection trials, a metric introduced in [74]. Results from both metrics are shown in table 1.

While noiseless G-QAOA outperforms its X-QAOA counterpart in all metrics, results from Aria show that the X-QAOA solution probability is often larger than the G-QAOA's since the G-QAOA requires greater circuit depth per round. Despite this, the G-QAOA outperforms the X-QAOA in both shots-to-all-solutions and shots-to-reject-fairness metrics. This shows that G-QAOA's fair-sampling advantage is evident even on noisy devices. Especially, the results for shots-to-all-solutions suggest that, in future applications on noisy quantum platforms, G-QAOA requires fewer resources than X-QAOA to find all solutions.

With both X-QAOA and G-QAOA, the experimental data also shows an enhancement in solution probability from $p = 1$ to $p = 2$. While such improvement for X-QAOA is common [7, 8, 10, 75], seeing the round-to-round improvement for G-QAOA had been observed only in [8] for the edge cover problem, which is less computationally challenging.

Table 1. Performance of X- and G-QAOA on IonQ Aria machines with 4 and 5 qubits. Results are summarized from figures 3 and 4. Shots-to-all-solutions: values are measured from 100 000 independent runs, and the indicated uncertainty is the 95% confidence interval. Shots-to-reject-fairness: values are taken from 10 000 trials. Bold values indicate the best-performing algorithm in each column. The metrics in each column are defined in the main text.

4-qubit	Solution % (noiseless)	Solution % (IonQ Aria)	Shots-to-all-solutions	Shots-to-reject-fairness
$p = 1$ X	52.0	49.70	12.53(3)	64
$p = 2$ X	83.7	72.20	11.42(3)	24
$p = 1$ G	56.1	45.35	12.16(2)	3780
$p = 2$ G	99.9	65.30	8.56(1)	426
5-qubit				
$p = 1$ X	38.1	28.18	35.51(7)	52
$p = 2$ X	68.9	51.25	18.61(4)	78
$p = 1$ G	43.1	31.48	26.79(5)	628
$p = 2$ G	98.7	46.40	18.09(3)	1278

5. Conclusion

Our research reports numerical evidence of the quadratic speedup of G-QAOA over random guessing for All-SAT and All-Max-SAT problems for thousands of tested 3-SAT instances. The sampled instances cover problems with 10–26 Boolean variables across various densities. Compared to the Grover algorithm, which also gives a quadratic speedup, the G-QAOA for 3-SAT problems requires less quantum resources in both circuit depth and auxiliary qubits. It is also more flexible and can be directly applied to Max-SAT problems. We do note that in contrast to random sampling, leading classical algorithms, as detailed in appendix F, may scale better than G-QAOA for typical instances; however, for hard instances in which the classical algorithms are likely to saturate the upper bound [76], we are optimistic on the performance of G-QAOA for the classically hard instances, since it does not rely in any obvious way on the structure of problems, neither show any dependence on the clause density. Furthermore, we execute X-QAOA and G-QAOA for small 4- and 5-variable instances on the IonQ Aria quantum computing platform, showing that the hardware maintains the uniform solution sampling property for G-QAOA relative to X-QAOA as quantified by reduced shots-to-all-solutions and other fairness metrics. Despite the longer circuit depth, we also note that the G-QAOA results obtained from demonstrations on the IonQ Aria machine improve from levels $p = 1$ to $p = 2$.

We also observe that a single-angle-pair constraint provides a simplified G-QAOA protocol that requires optimizing only two parameters while maintaining the quadratic speedup seen in the full G-QAOA. The faster optimization of the single-angle-pair approach allows us to obtain 26-qubit numerical results with comparable computational time required to obtain 20-qubit results for the full G-QAOA on classical simulators.

We also find clustering of the two parameters in the single-angle-pair G-QAOA, even for a large number of rounds $p \gtrsim 1000$. While parameter clustering has previously been seen in different QAOA variants, these focused on only few-round results [23, 24, 77]. Therefore, the findings in this work are promising for future large-scale applications of QAOA, and can significantly reduce the classical overhead on parameter training. For example, it possibly could circumvent the notorious barren plateaus challenge in variational searches [71, 78]. While quadratic speedups are unlikely to be sufficient for practical applications in the foreseeable future [79, 80], our results are likely to give insight into designing quantum algorithms more generally.

This research lays the groundwork for progress in several directions. These include understanding the observed efficacy of the single-angle-pair and clustering in the G-QAOA protocol, using parameter clustering to accelerate the optimization of the G-QAOA parameters further, applying G-QAOA to other problems in science and engineering, and exploring avenues for achieving more significant speedup. Given the fundamental role of SAT problems, it encourages future work to apply this single-angle-pair protocol to other combinatorial optimization problems, as different shapes of density of states may lead to different performance of G-QAOA [81], as well as different circuit complexity overhead [82]. We believe a fruitful area suggested by these results is to find simple, possibly analytical, ways of determining the G-QAOA angles from the problem structure, for example, statistical characteristics of its density of states. See [81, 83, 84] for related considerations. Finally, other questions are to what extent the run time may improve when we soften the single-angle-pair constraint differently from the naïve relaxation, or when use a warm-start [85, 86].

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.6084/m9.figshare.25221935>. Data will be available from 29 February 2024.

Acknowledgments

The authors thank Dr Nai-Hui Chia for providing valuable insights. G P, K R A H and Z Z acknowledge support from the Office of Naval Research (N00014-20-1-2695 and N00014-23-12665). K R A H and Z Z also acknowledge support from the National Science Foundation (PHY-1848304 and CMMI-2037545) and the Department of Energy (DE-SC0024301). G P also acknowledges support from the National Science Foundation (NSF CAREER Award No. PHY-2144910), the Army Research Office (W911NF22C0012), the Welch Foundation Grant No. C-2154, and the Office of Naval Research Young Investigator Program (N00014-22-1-2282). G P acknowledges that this material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under the Early Career Award No. DE-SC0023806. A K and D Q acknowledge support by NSF FET:Small No. 1907936, NSF CMMI-2037545, NSF CAREER award No. 2145629, Welch Foundation Grant #A22-0307, Microsoft Research Award, Amazon Research Award, and a Rice InterDisciplinary Excellence Award (IDEA). L D O and R P acknowledge support by NSF award CMMI-2037545.

Appendix A. Comparing with Grover algorithm

For 3-SAT problems, the Grover algorithm is more costly than the G-QAOA in ancilla number and circuit depth. Building up the Grover oracle for $\phi(m, n)$ takes $2m$ triple-controlled gates for the OR logic and two m -controlled gates for the AND logic, which can roughly be mapped to $10m$ Toffoli gates or $50m$ MS gates. One ancilla is needed for each clause as the register for OR logic, and m ancillas are needed to decompose the m -controlled gate, which gives a total of $2m$ ancillas. Thus, around the critical density $m \sim 4.26$, the QAOA problem unitary reduces the circuit depth and scalability by an order of magnitude. Meanwhile, the Grover algorithm requires $\frac{\pi}{8\sqrt{p}}$ iterations to boost the solution probability to $\frac{1}{2}$, where the constant is close to what we observe for G-QAOA in figures 1(a) and (b), i.e. at large $\frac{1}{p}$ it takes roughly $\frac{1}{3\sqrt{p}}$ rounds of G-QAOA to boost the solution probability to $\frac{1}{2}$.

The benefits of G-QAOA become more pronounced at higher densities, even though 3-SAT problems are less likely to possess satisfying solutions in such regimes. Nonetheless, identifying states that maximize the number of satisfied clauses (Max-SAT) is also a critical and hard problem. As the density m increases and solutions become sparse, the Grover algorithm may require relatively more resources due to the absence of solutions with large m . In these high-density regions, G-QAOA may most significantly outperform the traditional Grover algorithm in gate counts.

Additionally, the overhead put by satisfiability check is negligible in QAOA, as well as in most of the variational quantum algorithms. At the end of QAOA, checking whether a sampled instance is satisfiable is cheaper than calculating the objective (4), as the former only requires one measurement while the latter necessitates multiple measurements to reduce the error bar.

Appendix B. The Boolean satisfiability problem

A Boolean formula $\phi \equiv \phi(m, n)$ over a set X of n Boolean variables is in CNF when it is of the form

$$\phi = C_1 \wedge C_2 \wedge \cdots \wedge C_m, \quad (\text{B1})$$

where $\wedge$ denotes logical conjunction (AND), and every clause, C_i , $1 \leq i \leq m$, is of the form

$$C_i = l_{i_1} \vee l_{i_2} \vee \cdots \vee l_{i_{k_i}}, \quad (\text{B2})$$

where $\vee$ represents logical disjunction (OR), k_i is the length of clause C_i , and every literal l_{ij} is either equal to a variable $x \in X$ or its negation $\bar{x}$. We use s_{ij} to denote the sign of literal l_{ij} , i.e. $s_{ij} = 1$ for x_{ij} and $s_{ij} = -1$ for $\bar{x}_{ij}$.

An assignment of truth values to variables (or ‘configuration’), is a map $\tau: X \rightarrow \{0, 1\}$. For $x \in X$, we say that an assignment τ satisfies the literal x when $\tau(x) = 1$, and satisfies the literal $\bar{x}$ when $\tau(x) = 0$. An assignment τ satisfies clause C_i if τ satisfies at least one literal l_{ij} , $1 \leq j \leq k_i$, and τ satisfies formula ϕ when τ satisfies every clause C_i , $1 \leq i \leq m$. For convenience, we write $\phi(\tau) = 1$ when τ satisfies ϕ and write $\phi(\tau) = 0$, otherwise.

The Boolean satisfiability (SAT) problem is deciding whether an arbitrary Boolean formula ϕ expressed in CNF has at least one satisfying assignment. A *solution* is a satisfying assignment of the formula, and we use $R_\phi = \{\tau : \phi(\tau) = 1\}$ to denote the set of solutions of ϕ . The satisfying probability, i.e. the likelihood that a random assignment is a solution, is $P = |R_\phi|/2^n$.

Several crucial combinatorial optimization problems are variants of the Boolean satisfiability problem. If all the clauses have k literals, deciding whether there is a satisfying assignment is the k -SAT problem [30, 31]. Finding all the satisfying assignments is All-SAT [41]. Finding an assignment with the fewest unsatisfied clauses (which could be non-zero) is known as Max-SAT [59, 60]. Finding all assignments that have the fewest unsatisfied clauses is known as Max-All-SAT. All are important in theoretical computer science and have practical applications in science and engineering [87].

Appendix C. Optimizing QAOA angles

Optimization. To minimize the 3-SAT cost function, equation (4), for single-angle-pair G-QAOA, we first evaluate the cost function on a uniform grid in β and γ with grid spacing $\pi/180$ in each direction for problem size no larger than 20, while for larger problems we sample 4000 angles per instance in the (β, γ) -plane. We then locally optimize the (β, γ) using these values as initial parameters for *L-BFGS-B* as implemented in the *scipy.optimize* package. We use a stopping criterion of 10^{-6} relative changes of the cost function between consecutive iterations. For the G-QAOA results where the angles (β_p, γ_p) are allowed to depend on the round, we use the same initial seeds but allow the *L-BFGS-B* local relaxation to optimize all the angles.




Figure 5. Same 16-variable instances considered with G-QAOA in the main text for different target solution ratios for the final state, 25%, 50%, and 75%.

In generating random instances, for each independent n -variable sample, we uniformly select a value d within the density range, and generate nd random clauses according to the density. Should nd not result in an integer, appropriate rounding is applied. During the classical simulation, we leverage the G-QAOA's invariance under the relabeling of states within the same energy manifold. This simplification enables us to do the exact simulation into two steps. The first step is a classical pre-computation which requires counting the states in each distinct energy manifold, which takes $O(2^n)$ computational time for n Boolean variables. The second step involves exact classical simulation of the quantum circuit, but this simulation is not in the full Hilbert space of dimension 2^n , and instead is done in an effective space whose dimension is the number of distinct energy manifolds $\sim O(m)$ [52]. We are therefore able to simulate the G-QAOA for up to 26 qubits, with one of the main bottlenecks being the classical pre-computation.

Convergence. The choice of 50% solution probability at the end of the QAOA does not affect the quadratic trend in figure 1. For example, the 16-variable instances of 3-SAT are tested in figure 5 with two different ending solution ratios below and above 50%. As shown in the graph, different choices of ending ratios for the target states do not impact on the general scaling.

Landscape. To illustrate the optimization landscape of the single-angle-pair G-QAOA for 3-SAT, figure 6 shows the energy landscape on the (β, γ) surface. Figure 6(a) gives an example of a typical instance. Figure 6(b) averages the energy landscape over all 108 tested instances of 10-qubit. Typical landscapes are




Figure 6. (a) The single-angle-pair G-QAOA landscape (energy $\langle H \rangle$) according to (4) of a typical 10-qubit instance at $p = 10$. (b) Average landscape of all 10-qubit instances at $p = 10$.

peaked in a similar area to the average. These results indicate an area within the parameter space (β, γ) that yields the best performance, corresponding to the region of parameter clustering observed in figure 2.

Appendix D. Parameter clustering for Max-SAT

The parameter clustering (the same as figure 2) is also observed for all the Max-SAT instances in figure 1(b), as shown in figure 7.






Figure 7. Distribution of angle pairs (β, γ) [Max-SAT data points in figure 1(b) for 10 to 26 qubits] in the single-angle-pair protocol, color-marked with the number of clauses m , which is the product of the problem size and clause density.

Appendix E. Circuit compilation

The n -qubit Grover mixer [52] can be decomposed into multiple double-controlled gates, as shown in figure 8. Further mapping double-controlled gates onto native gates of trapped-ion machines, the Mølmer-Sørensen (MS) gates, is done as described in [88].

An illustrative problem unitary for clause $[+1, +2, +3]$ is encoded as shown in figure 9.

When running circuits on the Aria machines, to reduce the impact of possible local noises, we reassigned the order of encoding qubits every 500 shots during the demonstration on Aria.



Appendix F. Comparison with classical algorithms

Classical approaches to 3-SAT problems such as branch and bound [89] algorithms, tensor networks [90–94], and weighted model counting [91, 95], provide various important comparisons for the G-QAOA results. Among these, we analyze the performance of a bucket elimination algorithm *Elim-Count* [96] which finds all solutions. An upper bound on the time complexity of the *Elim-Count* algorithm is $\mathcal{O}(n \cdot \exp(w^*(d)))$, with n nodes and $w^*(d)$ induced treewidth, related to the treewidth w by $w^*(d) = w - 1$. *Elim-Count* shares the same exponential time and space complexity on the induced width of the problem graph as other bucket-elimination algorithms [76]. It can also aid in sampling instances, a task that is at least as hard as counting. Sampling through *Elim-Count* can be performed in time growing linearly with the number of possible assignments by accessing the underlying graph constructed during the bucket elimination process.

We also compare runtimes of our algorithm to results obtained from *TensorOrder* [91], a weighted model counting tool that operates by contracting tensor networks. *TensorOrder* can use different solvers to count the number of assignments that satisfy a given CNF. We use *FlowCutter* as the classical solver as it provides the best performance. *FlowCutter* optimizes cut size and undergoes network planning, tree construction, and tree contraction phases where the optimal contraction is sought with a time complexity of $\mathcal{O}(c|E|)$, with c the optimal cut size and $|E|$ the number of edges. It also enables us to calculate the treewidth for determining the upper bound of the scaling of *Elim-Count*.

Figure 10(a) illustrates the scaling of *Elim-Count* and *TensorOrder* on the 3-SAT instances used in figure 1(a), while figure 10(b) provides additional hard instances up to 30 variables around the critical density $d = 4.26$. We use *FlowCutter* to calculate the run time and treewidth of each instance, and supply the calculated treewidth using the upper bound on time complexity scaling of *Elim-Count*. By comparing the scaling of these classical algorithms with G-QAOA, both in complexity and runtime, we observe that *TensorOrder* gives better scaling than G-QAOA on typical random instances when focusing on counting but not enumerating solutions as *Elim-Count* does. However, for instances around the critical density in




figure 10(b), although at small $1/P$ the performance of *TensorOrder* is better than G-QAOA, the performance at large $1/P$ is comparable to G-QAOA, although *TensorOrder* does not sample solutions as G-QAOA ends up doing in its implementation.

ORCID iDs

Zewen Zhang <https://orcid.org/0000-0003-2258-613X>

Bhuvanesh Sundar <https://orcid.org/0000-0002-8867-360X>

Guido Pagano <https://orcid.org/0000-0002-6760-4015>

Kaden R A Hazzard <https://orcid.org/0000-0003-2894-7274>

References

- [1] Brassard G, Hoyer P, Mosca M and Tapp A 2002 Quantum amplitude amplification and estimation *Contemporary Mathematics: Quantum Computation and Information* (AMS Special Session Quantum Computation and Information vol 305) ed L Samuel and B Howard (American Mathematical Society) p 53
- [2] Kadowaki T and Nishimori H 1998 Quantum annealing in the transverse Ising model *Phys. Rev. E* **58** 5355
- [3] McClean J R, Romero J, Babbush R and Aspuru-Guzik A 2016 The theory of variational hybrid quantum-classical algorithms *New J. Phys.* **18** 023023
- [4] Hadfield S, Wang Z, O’Gorman B, Rieffel E G, Venturelli D and Biswas R 2019 From the quantum approximate optimization algorithm to a quantum alternating operator ansatz *Algorithms* **12** 34
- [5] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv:1411.4028)
- [6] Pagano G et al 2020 Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator *Proc. Natl. Acad. Sci.* **117** 25396
- [7] Harrigan M P et al 2021 Quantum approximate optimization of non-planar graph problems on a planar superconducting processor *Nat. Phys.* **17** 332
- [8] Zhu Y, Zhang Z, Sundar B, Green A M, Alderete C H, Nguyen N H, Hazzard K R and Linke N M 2022 Multi-round QAOA and advanced mixers on a trapped-ion quantum computer *Quantum Sci. Technol.* **8** 015007
- [9] Moses S et al 2023 A race track trapped-ion quantum processor (arXiv:2305.03828)
- [10] Shaydulin R and Pistoia M 2023 QAOA with $N \cdot p \geq 200$ (arXiv:2303.02064)
- [11] Shaydulin R et al 2023 Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem (arXiv:2308.02342)
- [12] Lubinski T, Coffrin C, McGeoch C, Sathe P, Apanavicius J and Neira D E B 2023 Optimization applications as quantum performance benchmarks (arXiv:2302.02278)
- [13] Dupont M et al 2023 Quantum enhanced greedy solver for optimization problems (arXiv:2303.05509)
- [14] Maciejewski F B et al 2023 Design and execution of quantum circuits using tens of superconducting qubits and thousands of gates for dense Ising optimization problems (arXiv:2308.12423)
- [15] Wurtz J and Lykov D 2021 Fixed-angle conjectures for the quantum approximate optimization algorithm on regular MaxCut graphs *Phys. Rev. A* **104** 052419
- [16] Kremenetski V, Apte A, Hogg T, Hadfield S and Tubman N M 2023 Quantum alternating operator Ansatz (QAOA) beyond low depth with gradually changing unitaries (arXiv:2305.04455)
- [17] Yang Z-C, Rahmani A, Shabani A, Neven H and Chamon C 2017 Optimizing variational quantum algorithms using Pontryagin’s minimum principle *Phys. Rev. X* **7** 021027
- [18] Sack S H and Serbyn M 2021 Quantum annealing initialization of the quantum approximate optimization algorithm *Quantum* **5** 491
- [19] Brady L T, Baldwin C L, Bapat A, Kharkov Y and Gorshkov A V 2021 Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems *Phys. Rev. Lett.* **126** 070505
- [20] Wurtz J and Love P J 2022 Counterdiabaticity and the quantum approximate optimization algorithm *Quantum* **6** 635

- [21] Wu M, Liu Z and Chen H 2023 Adiabatic-passage based parameter setting method for quantum approximate optimization algorithm on 3-SAT problem (arXiv:2312.00077)
- [22] Scriva G, Astrakhantsev N, Pilati S and Mazzola G 2023 Challenges of variational quantum optimization with measurement shot noise (arXiv:2308.00044)
- [23] Streif M and Leib M 2020 Training the quantum approximate optimization algorithm without access to a quantum processing unit *Quantum Sci. Technol.* **5** 034008
- [24] Moussa C, Wang H, Bäck T and Dunjko V 2022 Unsupervised strategies for identifying optimal parameters in quantum approximate optimization algorithm *EPJ Quantum Technol.* **9** 11
- [25] Farhi E, Goldstone J, Gutmann S and Zhou L 2022 The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size *Quantum* **6** 759
- [26] Galda A, Gupta E, Falla J, Liu X, Lykov D, Alexeev Y and Saftro I 2023 Similarity-based parameter transferability in the quantum approximate optimization algorithm *Front. Quantum Sci. Technol.* **2** 1200975
- [27] Shaydulin R, Lotshaw P C, Larson J, Ostrowski J and Humble T S 2023 Parameter transfer for quantum approximate optimization of weighted MaxCut *ACM Trans. Quantum Comput.* **4** 1
- [28] Kapit E, Barton B A, Feeney S, Grattan G, Patnaik P, Sagal J, Carr L D and Oganessian V 2023 On the approximability of random-hypergraph MAX-3-XORSAT problems with quantum algorithms (arXiv:2312.06104)
- [29] Yu Y, Cao C, Wang X-B, Shannon N and Joynt R 2023 Solution of SAT problems with the adaptive-bias quantum approximate optimization algorithm *Phys. Rev. Res.* **5** 023147
- [30] Cook S A 2023 The complexity of theorem-proving procedures *Logic, Automata and Computational Complexity: The Works of Stephen A. Cook* (Association for Computing Machinery) pp 143–52
- [31] Karp R M 2010 Reducibility among combinatorial problems *50 Years of Integer Programming 1958–2008: From the Early Years to the State-of-the-Art* ed M Jünger, T M Lieblich, D Naddef, G L Nemhauser, W R Pulleyblank, G Reinelt, G Rinaldi and L A Wolsey (Springer) pp 219–41
- [32] Lucas A 2014 Ising formulations of many NP problems *Front. Phys.* **2** 5
- [33] Khurshid S and Marinov D 2004 TestEra: specification-based testing of Java programs using SAT *Autom. Softw. Eng.* **11** 403
- [34] Vizek Y, Weissenbacher G and Malik S 2015 Boolean satisfiability solvers and their applications in model checking *Proc. IEEE* **103** 2021
- [35] Dutra R, Loeferer K, Bachrach J and Sen K 2018 Efficient sampling of SAT solutions for testing *Proc. 40th Int. Conf. on Software Engineering* (ACM) pp 549–59
- [36] Gaber L, Hussein A I, Mahmoud H, Mabrook M M and Moness M 2020 Computation of minimal unsatisfiable subformulas for SAT-based digital circuit error diagnosis *J. Ambient Intell. Humaniz. Comput.* **13** 3693–371
- [37] Lafitte F, Nakahara J Jr and Van Heule D 2014 Applications of SAT solvers in cryptanalysis: finding weak keys and preimages *J. Satisf. Boolean Model. Comput.* **9** 1–25
- [38] Lynce I and Marques-Silva J 2006 Efficient haplotype inference with Boolean satisfiability *National Conf. on Artificial Intelligence (AAAI)* (AAAI Press)
- [39] Paredes R, Dueñas-Osorio L, Meel K S and Vardi M Y 2019 Principled network reliability approximation: a counting-based approach *Reliab. Eng. Syst. Saf.* **191** 106472
- [40] Marques-Silva J 2008 Practical applications of Boolean satisfiability *2008 9th Int. Workshop on Discrete Event Systems* (IEEE) pp 74–80
- [41] Yu Y, Subramanyan P, Tsiskaridze N and Malik S 2014 All-SAT using minimal blocking clauses *2014 27th Int. Conf. on VLSI Design and 2014 13th Int. Conf. on Embedded Systems* (IEEE) pp 86–91
- [42] Valiant L G 1979 The complexity of computing the permanent *Theor. Comput. Sci.* **8** 189
- [43] Battaglia D A, Santoro G E and Tosatti E 2005 Optimization by quantum annealing: lessons from hard satisfiability problems *Phys. Rev. E* **71** 066707
- [44] Azinović M, Herr D, Heim B, Brown E and Troyer M 2017 Assessment of quantum annealing for the construction of satisfiability filters *SciPost Phys.* **2** 013
- [45] Ayanzadeh R, Halem M and Finin T 2020 Reinforcement quantum annealing: a hybrid quantum learning automata *Sci. Rep.* **10** 7952
- [46] Cheng S-T and Tao M-H 2007 Quantum cooperative search algorithm for 3-SAT *J. Comput. Syst. Sci.* **73** 123
- [47] Alasow A, Jin P and Perkowski M 2022 Quantum algorithm for variant maximum satisfiability *Entropy* **24** 1615
- [48] Varmantchaonala C M, Fendji J L K E, Njafa J P T and Atemkeng M 2023 Quantum hybrid algorithm for solving sat problem *Eng. Appl. Artif. Intell.* **121** 106058
- [49] Grover L K 1996 A fast quantum mechanical algorithm for database search *Proc. 28th Annual ACM Symp. on Theory of Computing (STOC '96)* (Association for Computing Machinery) pp 212–9
- [50] Jiang Z, Rieffel E G and Wang Z 2017 Near-optimal quantum circuit for Grover's unstructured search using a transverse field *Phys. Rev. A* **95** 062317
- [51] Yang W-L, Wei H, Zhou F, Chang W-L and Feng M 2009 Solution to the satisfiability problem using a complete Grover search with trapped ions *J. Phys. B: At. Mol. Opt. Phys.* **42** 145503
- [52] Sundar B, Paredes R, Damanik D T, Duenas-Osorio L and Hazzard K R 2019 A quantum algorithm to count weighted ground states of classical spin Hamiltonians (arXiv:1908.01745)
- [53] Bartschi A and Eidenbenz S 2020 Grover mixers for QAOA: shifting complexity from mixer design to state preparation *Proc. - IEEE Int. Conf. on Quantum Computing and Engineering, QCE 2020* p 72
- [54] Golden J, Bartschi A, O'Malley D and Eidenbenz S 2023 The quantum alternating operator ansatz for satisfiability problems (arXiv:2301.11292)
- [55] Mandl A, Barzen J, Bechtold M, Leymann F and Wild K 2023 Amplitude amplification-inspired QAOA: improving the success probability for solving 3SAT (arXiv:2303.01183)
- [56] Golden J, Bartschi A, O'Malley D and Eidenbenz S 2023 Numerical evidence for exponential speed-up of qaoa over unstructured search for approximate constrained optimization *2023 IEEE Int. Conf. on Quantum Computing and Engineering (QCE) vol 1* (IEEE) pp 496–505
- [57] Kirkpatrick S and Selman B 1994 Critical behavior in the satisfiability of random boolean expressions *Science* **264** 1297
- [58] Mézard M, Parisi G and Zecchina R 2002 Analytic and algorithmic solution of random satisfiability problems *Science* **297** 812
- [59] Asano T and Williamson D P 2002 Improved approximation algorithms for MAX SAT *J. Algorithms* **42** 173

- [60] Nannicini G 2019 Performance of hybrid quantum-classical variational heuristics for combinatorial optimization *Phys. Rev. E* **99** 013304
- [61] Mézard M and Zecchina R 2002 Random k-satisfiability problem: from an analytic solution to an efficient algorithm *Phys. Rev. E* **66** 056126
- [62] Chen Y, Safarpour S, Veneris A and Marques-Silva J 2009 Spatial and temporal design debug using partial MaxSAT *Proc. 19th ACM Great Lakes Symposium on VLSI* pp 345–50
- [63] Dimitrova R, Ghasemi M and Topcu U 2018 Maximum realizability for linear temporal logic specifications *Automated Technology for Verification and Analysis: 16th Int. Symp., ATVA 2018 (Los Angeles, CA, USA, October 7–10, 2018) (Proc. 16)* (Springer) pp 458–75
- [64] Mallouf D and Meel K S 2018 Mlic: a MaxSAT-based framework for learning interpretable classification rules *Int. Conf. on Principles and Practice of Constraint Programming* (Springer) pp 312–27
- [65] Berg O J, Hyttinen A J and Jarvisalo M J 2019 Applications of MaxSAT in data analysis *Proc. of Pragmatics of SAT 2015 and 2018 Proc. of Pragmatics of SAT 2015 and 2018* vol 59 pp 50–64
- [66] Guerra J and Lynce I 2012 Reasoning over biological networks using maximum satisfiability *Int. Conf. on Principles and Practice of Constraint Programming* (Springer) pp 941–56
- [67] Martins R 2017 Solving RNA alignment with MaxSAT MaxSAT Evaluation 2017 MaxSAT Evaluation 2017: Solver and Benchmark Descriptions vol B-2017-2 p 29–30 (available at: <https://hdl.handle.net/10138/233112>)
- [68] Matsuda Y, Nishimori H and Katzgraber H G 2009 Ground-state statistics from annealing algorithms: quantum versus classical approaches *New J. Phys.* **11** 073021
- [69] Bittel L and Kliesch M 2021 Training variational quantum algorithms is NP-hard *Phys. Rev. Lett.* **127** 120502
- [70] Shaydulin R, Safro I and Larson J 2019 Multistart methods for quantum approximate optimization 2019 *IEEE High Performance Extreme Computing Conf. (HPEC)* (IEEE) pp 1–8
- [71] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 4812
- [72] Sørensen A and Mølmer K 1999 Quantum computation with ions in thermal motion *Phys. Rev. Lett.* **82** 1971
- [73] Zhu Y, Green A M, Nguyen N H, Huerta Alderete C, Mossman E and Linke N M 2023 Pairwise-parallel entangling gates on orthogonal modes in a trapped-ion chain *Adv. Quantum Technol.* **6** 2300056
- [74] Golden J, Bärttschi A, O'Malley D and Eidenbenz S 2022 Fair sampling error analysis on NISQ devices *ACM Trans. Quantum Comput.* **3** 1
- [75] Moses S A et al 2023 A race-track trapped-ion quantum processor *Phys. Rev. X* **13** 041052
- [76] Dechter R 1999 Bucket elimination: a unifying framework for reasoning *Artif. Intell.* **113** 41
- [77] Akshay V, Rabinovich D, Campos E and Biamonte J 2021 Parameter concentrations in quantum approximate optimization *Phys. Rev. A* **104** L010401
- [78] Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L and Coles P J 2021 Noise-induced barren plateaus in variational quantum algorithms *Nat. Commun.* **12** 6961
- [79] Babbush R, McClean J R, Newman M, Gidney C, Boixo S and Neven H 2021 Focus beyond quadratic speedups for error-corrected quantum advantage *PRX Quantum* **2** 010103
- [80] Beverland M E, Murali P, Troyer M, Svore K M, Hoefler T, Kliuchnikov V, Low G H, Soeken M, Sundaram A and Vaschillo A 2022 Assessing requirements to scale to practical quantum advantage (arXiv:2211.07629)
- [81] Headley D and Wilhelm F K 2023 Problem-size-independent angles for a Grover-driven quantum approximate optimization algorithm *Phys. Rev. A* **107** 012412
- [82] Bakó B, Glos A, Salehi O and Zimborás Z 2022 Near-optimal circuit design for variational quantum optimization (arXiv:2209.03386)
- [83] Bridi G A and de Lima Marquezino F 2024 Analytical results for the quantum alternating operator ansatz with Grover mixer (arXiv:2401.11056)
- [84] Xie N, Xu J, Chen T, Saito Y, Asai N and Cai D 2024 Performance upper bound of the grover-mixer quantum alternating operator ansatz (arXiv:2405.03173)
- [85] Egger D J, Mareček J and Woerner S 2021 Warm-starting quantum optimization *Quantum* **5** 479
- [86] Tate R, Moondra J, Gard B, Mohler G and Gupta S 2023 Warm-started QAOA with custom mixers provably converges and computationally beats Goemans-Williamson's Max-Cut at low circuit depths *Quantum* **7** 1121
- [87] Duenas-Osorio L, Meel K, Paredes R and Vardi M 2017 Counting-based reliability estimation for power-transmission grids *Proc. AAAI Conf. on Artificial Intelligence* vol 31
- [88] Maslov D 2017 Basic circuit compilation techniques for an ion-trap quantum machine *New J. Phys.* **19** 023035
- [89] Li C-M and Quan Z 2010 An efficient branch-and-bound algorithm based on maxsat for the maximum clique problem *Proc. AAAI Conf. on Artificial Intelligence* vol 24 p 128
- [90] Gray J and Kourtis S 2021 Hyper-optimized tensor network contraction *Quantum* **5** 410
- [91] Dudek J M and Vardi M Y 2020 Parallel weighted model counting with tensor networks (arXiv:2006.15512)
- [92] Tamaki H 2019 Positive-instance driven dynamic programming for treewidth *J. Comb. Optim.* **37** 1283
- [93] Hamann M and Strasser B 2018 Graph bisection with pareto optimization *J. Exp. Algorithmics* **23** 1
- [94] Hicks I V 2002 Branchwidth heuristics *Congressus Numerantium* 31
- [95] Nagy S A, Paredes R, Dudek J M, Dueñas-Osorio L and Vardi M Y 2022 Ising model partition function computation as a weighted counting problem (arXiv:2212.12812)
- [96] Dechter R, Kask K, Bin E and Emek R 2002 Generating random solutions for constraint satisfaction problems *AAAI/IAAI*