

High-Round QAOA for MAX k -SAT on Trapped Ion NISQ Devices

Elijah Pelofske, Andreas Bäertschi, John Golden, Stephan Eidenbenz

CCS-3 Information Sciences

Los Alamos National Laboratory

Los Alamos, NM 87544, USA

{epelofske, baertschi, golden, eidenbenz}@lanl.gov

Abstract—The Quantum Alternating Operator Ansatz (QAOA) is a hybrid classical-quantum algorithm that aims to sample the optimal solution(s) of discrete combinatorial optimization problems. We present optimized QAOA circuit constructions for sampling MAX k -SAT problems, specifically for $k = 3$ and $k = 4$. The novel 4-SAT QAOA circuit construction we present uses measurement based uncomputation, followed by classical feed forward conditional operations. The QAOA circuit parameters for 3-SAT are optimized via exact classical (noise-free) simulation, using HPC resources to simulate up to 20 rounds on 10 qubits. In order to explore the limits of current NISQ devices we execute these optimized QAOA circuits for random 3-SAT test instances with clause-to-variable ratio 4 on four trapped ion quantum computers: Quantinuum H1-1 (20 qubits), IonQ Harmony (11 qubits), IonQ Aria 1 (25 qubits), and IonQ Forte (30 qubits). The QAOA circuits that are executed include $n = 10$ up to $p = 20$, and $n = 22$ for $p = 1$ and $p = 2$. The high round circuits use upwards of 9,000 individual gate instructions, making these some of the largest QAOA circuits executed on NISQ devices. Our main finding is that current NISQ devices perform best at low round counts (i.e., $p = 1, \dots, 5$) and then – as expected due to noise – gradually start returning satisfiability truth assignments that are no better than randomly picked solutions as the number of QAOA rounds are further increased.

Index Terms—Quantum Alternating Operator Ansatz, QAOA, Boolean Satisfiability, quantum circuit, quantum computing, MAX k -SAT, MAX 3-SAT, MAX 4-SAT, combinatorial optimization

I. INTRODUCTION

The Quantum Approximate Optimization Algorithm (QAOA) [15], [16], which has subsequently been generalized to the Quantum Alternating Operator Ansatz [25] to include a wider variety of mixers and state preparation algorithms [58], [7], [21], [22], [9], [33], [29], under the same acronym of QAOA, is a quantum algorithm that is intended to sample optimal solutions of discrete combinatorial optimization problems.

QAOA has been applied to a number of types of combinatorial optimization problems, both numerically and experimentally on near term quantum computers, including on boolean k -SAT problems [63], [5], [23], maximum cut [52], [30], [27], [40], and the Sherrington-Kirkpatrick model [27], [17]. How QAOA scales with respect to problem size and number of rounds [59], [64], how learn-able the optimal (or near optimal) QAOA parameters are in the presence of noise [56], and many other related questions are topics of interest for investigating

problem-specific QAOA settings for potential advantageous quantum computation [18], [49]. The primary challenge is managing the relatively high error rates on current hardware.

While the suitability of QAOA for the NISQ era [44] is thus an ongoing debate, particularly because in general QAOA requires high p in order to be effective [14], [13], we instead propose to study QAOA performance as a benchmark-style test for NISQ devices, where we expect solutions to initially improve as we increase the rounds of QAOA before hitting a peak and then degrading to random solutions. The higher number of QAOA rounds a NISQ device can maintain or even increase the average quality of solutions sampled, the better its performance.

NISQ QAOA Performance Measures. To be more precise and as illustrated conceptually in Figure 1, we define the quality achieved in a QAOA NISQ experiment as the average approximation ratio of all solutions sampled by the device, where the approximation ratio of a solution s with objective function value $C(s)$ is defined as $\frac{C(s)}{C_{opt}}$ (for a maximization problem, and inverted for a minimization problem) with C_{opt} being the optimal objective function value across all feasible solutions. This is under the assumption that good angles, but not necessarily optimal angles, can be learned efficiently for each round p ; in other words the conceptual picture of NISQ QAOA in Figure 1 is not intended to capture the difficulty of the angle finding problem (especially at very high rounds). We plot approximation ratios on the y -axis, and put the length of the QAOA calculation, crudely measured in number of rounds p on the x -axis. On a noise-free or error-corrected quantum computing device (see red line), we expect the approximation ratio to gradually increase with p until it eventually hits 1.0 when QAOA has put all amplitude into the optimum solution(s). On the other hand, if we have a device that samples purely randomly because of its overwhelming noise level, it would sample random solutions perhaps mostly within a few percentile points above the average random solution (marked by the dashed horizontal lines in Figure 1). We expect a NISQ device (dark blue line) to initially improve approximation ratios to reach a peak value at the point labeled QAOA Volume at p_{max} , which is a measure defined as $n \times p$, where n is the number of qubits used. As we further increase p beyond the peak, we expect noise to gradually become more dominant until the solution quality becomes indistinguishable from

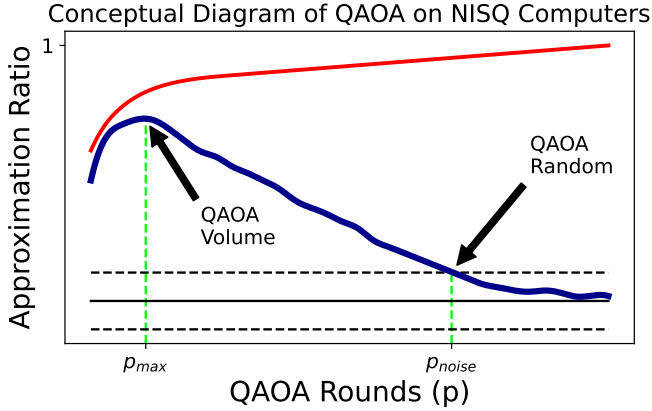


Fig. 1: Conceptual diagram of using QAOA to characterize the performance of NISQ devices for general combinatorial optimization problems. The horizontal black lines show the random sampling distribution, the red line shows the ideal (zero noise) quantum computation, and the blue line is a conceptual sketch of the expected QAOA performance when noise is present in the computation. The points p_{\max} and p_{noise} are marked with vertical dashed green lines.

random samples and we have thus lost all (quantum) signal. Let us define a point at p_{noise} , where the NISQ performance reaches the upper envelope of the 60th percentile of standard deviation from random sampling, where we have chosen the number 60 in an arbitrary but fixed fashion. In particular, the motivation for this tight percentile choice is based on the sampling characteristics of the 3-SAT problems that we will be sampling, for other combinatorial optimization problems it may make sense to use other reasonable percentiles. We call this point “QAOA Random”. The pair $(p_{\max}, p_{\text{noise}})$ characterizes QAOA performance on a NISQ device and it is of particular interest to study these components as we vary problem instance sizes.

For an error-corrected device, we would define p_{\max} to be the smallest p that reaches an approximation ratio that we wish to achieve, ideally we would hope that it is small and grows at most polynomially with the problem instance size as it becomes algorithmically defined by QAOA performance. In contrast, on a noisy device, we hope for p_{\max} to be as large as possible as it serves much more as an indicator that the NISQ device still improves signal at high round counts. The QAOA Random point with p_{noise} would be at infinity for an error-corrected device. In the reality of NISQ devices, we hope for p_{noise} to be as large as possible.

In this paper, we test how well NISQ devices match our proposed framework of QAOA performance. We choose the well-known MAX 3-SAT problem as our object of study. We are specifically interested in observing what happens to the QAOA computation on NISQ computers when pushed to very high rounds, consisting of several thousand single and two qubit gates. To this end, we experimentally utilize four trapped

ion based quantum computers, IonQ Harmony [60], IonQ Aria, IonQ Forte [10], and Quantinuum H1-1 [43] in order to push these QAOA computations to very high rounds (up to $p = 20$), where each round has a high gate depth (approximate gate depth of 200 per round).

Optimized k -SAT QAOA Circuits. We present optimized k -SAT transverse field QAOA circuit construction algorithms, in particular for 3-, and 4-SAT, targeting all-to-all connectivities with native Rzz two-qubit gates. The 4-SAT implementation makes use of mid-circuit measurement based AND uncomputation with mid circuit qubit resets in order to re-use ancilla qubits for the computation. Therefore, implementing the 4-SAT QAOA circuits can additionally provide a future benchmark on the capability of quantum computers combined with classical co-processing (and mid circuit measurement) for feed forward conditioning. The 4-SAT implementation requires even more gate operations and gate depth per round compared to 3-SAT, and puts it well out of reach of the state-of-the-art quantum computers at this moment. However, we believe that such 4-SAT QAOA circuits may be within reach for new quantum computers in the coming years, and therefore we provide these circuit construction algorithms as a benchmark for future quantum computers, with the intent being that very high round QAOA will measure the computational reach for very high depth quantum circuits.

Formally, in the problem of Maximum k -Satisfiability (or Max k -SAT), we are given a logical formula in Conjunctive Normal Form (CNF) defined on n binary truth variables in the variable vector $x := x_0, \dots, x_{n-1}$ with $x_j = \{0, 1\}$ for $j = 0, \dots, n-1$. The CNF formula can be expressed as a conjunction of disjunctions (more commonly known as clauses), where each clause is defined on k literals, which are instances or negated instances of the variables. Furthermore we require each variable to appear at most once per clause. The goal is to find a truth value assignment for the variables such that a maximum number of clauses are satisfied; thus, we define an objective function $C(x): \{1, 0\}^n \rightarrow \mathbb{N}$, where $C(x)$ is the number of clauses satisfied by truth assignment x . Max k -SAT is NP-hard for $k \geq 2$, and no classical polynomial-time approximation algorithm can guarantee to find a solution that is within $\frac{2^k-1}{2^k} + \epsilon$ of the optimum [28] (for any arbitrary but fixed $\epsilon > 0$); note that the decision version of 2-SAT can be solved classically polynomially. Consider an example of MAX 3-SAT with $n = 3$ (though in general $n \neq k$):

$$C(x) = (\overline{x_0} \vee x_1 \vee x_2) + (\overline{x_0} \vee \overline{x_1} \vee x_2) + (\overline{x_0} \vee \overline{x_1} \vee \overline{x_2}) + (x_0 \vee x_1 \vee x_2)$$

We have liberally mixed logical notation (e.g., the \vee for logical OR) and arithmetic notation to best serve our purpose. The first clause,¹ $(\overline{x_0} \vee x_1 \vee x_2)$, is satisfiable (i.e. equals 1), when at least one of the literals $\overline{x_0}, x_1, x_2$ evaluates to True. Equivalently, the only unsatisfiable assignment is $(x_0, x_1, x_2) \leftarrow (\text{True}, \text{False}, \text{False}) = (1, 0, 0)$. Thus, of the $2^3 = 8$ possible truth assignments, the four

¹which we will also use to illustrate the circuit design in Section II-B.

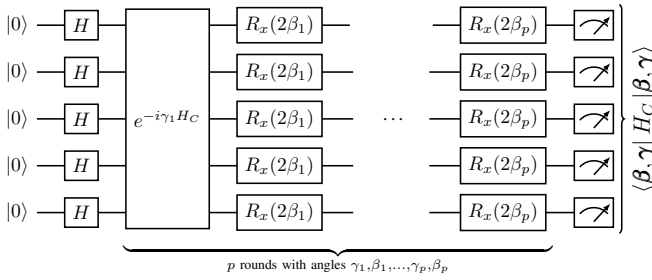


Fig. 2: Main QAOA subroutine: State Preparation, p rounds of Phase Separator and Mixer applications, Measurement. Both state preparation and each Transverse Field mixer can be implemented with single-qubit gates, $H|0\rangle = |+\rangle$ and $R_x(2\beta) = e^{-i\beta X}$.

assignments $(0, 0, 1), (0, 1, 0), (0, 1, 1), (1, 0, 1)$ result in all four clauses being satisfied, e.g. $C((0, 1, 0)) = 4$, while $(1, 0, 0), (1, 1, 0), (0, 0, 0), (1, 1, 1)$ result in only three satisfied clauses (where each solution vector is (x_0, x_1, x_2)).

Related work. With respect to qubit count, QAOA has been applied to short depth QAOA circuits on IBMQ hardware lattices up to 127 qubits [41]. However, with respect to circuit depth and overall gate count, the results we present here are significantly larger than previous studies. QAOA has been implemented on Quantinuum quantum computers in several studies [52], [38], [39], [34], however the high round and large problem size QAOA circuits evaluated in this paper are the largest QAOA experimental implementations yet presented with respect to gate counts and depths of the circuits that the authors are aware of. We note in particular, that the impressive work [52] on MaxCut on 3-regular graphs on Quantinuum H1-1 device achieves a higher QAOA volume than our work, where QAOA volume is defined as $n \cdot p$ with n being the number of qubits; however, the phase separator for MaxCut on 3 regular graphs consist of $1.5n$ two-qubit gates (or $3n$ qubit gates if only using CNOT gates), whereas our 3-SAT phase separator requires $20n$ two-qubit gates (details in later sections).

There are a growing number of QAOA variants, for example using different mixers and phase separators [58], [7], [21], [22], [33], [29], [24], [65], however for the circuit construction algorithm and NISQ computer experiments we present, the standard transverse field mixer is used.

II. METHODS

In Section II-A, we briefly overview QAOA. Section II-B details a general algorithm for constructing QAOA circuits for k -SAT problems. Section II-C describes the hardware implementations on IonQ Harmony and Quantinuum H1-1.

A. QAOA overview

Given an discrete variable assignment combinatorial optimization problem, we can define a an objective function of the form $C(x): \{1, 0\}^n \rightarrow \mathbb{N}$ which provides a real number evaluation for any variable assignments. In this case, where

the combinatorial optimization problem is maximization, we aim to find variable assignments that provide the maximum possible objective function evaluation over all possible inputs. The QAOA algorithm intends to sample a combinatorial optimization of this form. The ingredients of QAOA are as follows:

- an initial state $|\psi\rangle$, here the equal superposition over all computational basis states $|\psi\rangle = |+\rangle^n$,
- a classical phase separating Cost Hamiltonian $H_C = \sum_{x \in \{0,1\}^n} C(x) |x\rangle \langle x|$,
- a mixing Hamiltonian, here the Pauli-X transverse field $H_M = \sum_{i=1}^n X_i$,
- an integer $p \geq 1$, the number of subroutine rounds (also referred to as *layers*),
- two vectors composed of real numbers $\gamma = (\gamma_1, \dots, \gamma_p)$ and $\beta = (\beta_1, \dots, \beta_p)$, each vector having a length of p .

The quantum subroutine of QAOA, depicted in Figure 2, then prepares a state

$$|\beta, \gamma\rangle = e^{-i\beta_p H_M} e^{-i\gamma_p H_C} \dots e^{-i\beta_1 H_M} e^{-i\gamma_1 H_C} |\psi\rangle \quad (1)$$

and samples solutions from it, with an expectation value of $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$. To utilize QAOA, the classical outer loop routine is to find these real vectors $\vec{\gamma}$ and $\vec{\beta}$, typically referred to as the *angles*, such that the expectation value is maximized (or minimized). Importantly at low rounds (e.g. $p = 1, 2$) it is not expected that we will be able to reliably sample the optimal solutions, especially for very large problems. In order to increase the probability of sampling good solutions, we must increase the number of rounds; as we increase the number of rounds, assuming we are able to learn good angles at each round, we expect a consistent increase of the ideal computation of the approximation ratio $\langle \beta, \gamma | H_C | \beta, \gamma \rangle / C_{\text{opt}}$ of the sampled problem. QAOA is considered a type of variational quantum algorithm [8], or a hybrid quantum-classical algorithm, because in the absence of knowing good angles β and γ a-priori, it is likely the case that the computation would need evaluate the quality of some angle combination using a quantum computer, and then update those angles using machine learning on a classical co-processor. Errors in the quantum computation make this angle finding task considerably more difficult however, and in the case of the experiments we present, we classical learn high quality angles to use for the QAOA computation on quantum computer. In this respect, this study is not aiming to determine the learnability of high round QAOA angles, but more so the ideal and the NISQ QAOA scaling of approximation ratios with respect to problem size and increasing rounds.

B. QAOA Circuit Construction for MAX k -SAT

To design a circuit that prepares and samples from the state $|\beta, \gamma\rangle$ in Equation (1), we need to take care of the state preparation of $|+\rangle^n$, mixers of the form $e^{-i\beta \sum_i X_i}$, and phase separators of the form $e^{-i\gamma H_C}$. As illustrated in Figure 2, the first two are achieved with simple one-qubit gates.

To implement the phase separator for MAX k -SAT, we write first write the cost function $C(x)$ as a sum of m clause




Fig. 3: Phase Separator for a single 3-SAT clause: The logical OR of three binary variables corresponds to the negated AND of the negated variables. We neglect the global phase of $e^{-i\gamma \cdot 1}$ and phase the only non-satisfiable assignment $|100\rangle$ by $e^{+i\gamma}$. The 2-controlled phase gate $\text{CCPhase}(\gamma) = \text{diag}(1, 1, 1, 1, 1, 1, 1, e^{i\gamma})$ we decompose with 5 two-qubit gates (4 CNOT, 1 RZZ), analogous to standard Toffoli decompositions [2], [47], which yields the optimal two-qubit gate count [62].




Fig. 4: Phase Separator for a single 4-SAT clause (focusing on negative literals): The number of quantum controls on the phase shift gate can be decreased by a pairwise AND computation into an ancilla qubit, followed by the smaller controlled phase shift and an AND^\dagger uncomputation. The AND can be implemented with only 3 CNOTs using a relative phase Toffoli implementation [36], [54] given by Qiskit’s Margolus gate RCCX and a consecutive S^\dagger .

The AND^\dagger can be implemented with measurement based uncomputation [20], utilizing mid circuit measurement (followed by qubit reset) and a classical feed forward condition for a single CZ two-qubit gate. Furthermore, this also allows to replace the last CNOT of the 2-controlled phase shift from Figure 3 to be moved to a classically controlled single-qubit gate.

Thus we used 3 CNOTs and 1 classically controlled CZ to intermittently decrease the number of quantum controls on the phase shift by one, plus another 3 CNOTs and 1 RZZ Gate to implement the phase shift, for a total of 8 two-qubit gates. For a k -SAT clause with $k > 3$, recursively nesting this construction generalizes to a decomposition using $k - 3$ ancilla qubits and $4 \cdot (k - 3) + 4 = 4k - 8$ two-qubit gates.

functions $C_j(x): \{0, 1\}^n \rightarrow \{0, 1\}$, each selecting k *distinct* variables x_{j1}, \dots, x_{jk} variables of the input with positive or negative literal (i.e. either $\ell_{ji}(x_{ji}) = x_{ji}$ or $\ell_{ji}(x_{ji}) = \overline{x_{ji}}$):

$$C(x) = \sum_{j=1}^m C_j(x)$$

$$C_j(x) = (\ell_{j1}(x_{j1}) \vee \ell_{j2}(x_{j2}) \vee \dots \vee \ell_{jk}(x_{jk}))$$

We can then rewrite the cost Hamiltonian H_C as

$$H_C = \sum_{x \in \{0,1\}^n} \sum_{j=1}^m C_j(x) |x\rangle \langle x|$$

$$= \sum_{j=1}^m \underbrace{\sum_{x \in \{0,1\}^n} C_j(x) |x\rangle \langle x|}_{=: H_{C_j}}$$

Since H_C is diagonal in the computational basis and the terms commute, we can implement $e^{-i\gamma H_C}$ as $\prod_{j=1}^m e^{-i\gamma H_{C_j}}$, meaning we implement the phasing for each clause individually. Furthermore, only k variables appear in each clause, making H_{C_j} k -local. We now describe efficient circuits for $k = 3, 4$ in detail using Figures 3 and 4. We note that our construction easily extends to arbitrary k -SAT clauses.

Clause Phase Separators: Our goal is to compile $e^{-i\gamma H_{C_j}}$ efficiently with the respect to our NISQ hardware. In particular, we assume all-to-all connectivity and native parameterizable $\text{RZZ}(2\gamma) = e^{-i\gamma Z \otimes Z}$ two-qubit gates. The

main cost in both execution time and source of error on NISQ devices is the number of two-qubit gates, hence we aim to minimize their use per clause. A single CNOT gate can always be implemented (up to single-qubit gates) with a single RZZ gate, but *not vice versa* (depending on the angle γ). Hence our approach utilizes both CNOT and RZZ where applicable.

For a k -SAT clause C_j , note that applying $e^{-i\gamma H_{C_j}}$ amounts to phase all $2^k - 1$ satisfiable variable assignments by $e^{-i\gamma}$. This is—up to a global phase of $e^{-i\gamma}$ —equivalent to instead phasing the *only* unsatisfiable assignment by $e^{+i\gamma}$. We illustrate this for $C_j(x) = (\overline{x_0} \vee x_1 \vee x_2)$, see Figure 3:

$$e^{-i\gamma H_{C_j}} |x_0 x_1 x_2\rangle = \begin{cases} e^{-i\gamma} |x_0 x_1 x_2\rangle & \text{if } x \neq (1, 0, 0) \\ 1 |x_0 x_1 x_2\rangle & \text{otherwise,} \end{cases}$$

$$= e^{-i\gamma} \cdot \begin{cases} e^{+i\gamma} |x_0 x_1 x_2\rangle & \text{if } x = (1, 0, 0) \\ 1 |x_0 x_1 x_2\rangle & \text{otherwise.} \end{cases}$$

Evaluating whether the clause C_j is unsatisfied amounts to evaluate whether the AND of its negated literals is satisfied, $1 - (\overline{x_0} \vee x_1 \vee x_2) = (x_0 \wedge \overline{x_1} \wedge \overline{x_2})$. The latter is true if and only if (i) **after negating** variables x_1 and x_2 (which in C_j have positive literals), we have (ii) $(x_0 \wedge x_1 \wedge x_2) = \text{True}$.

Circuit Examples and Gate Counts: Implementing (i) is done by simply conjugating (ii) with X -gates on qubits q_1 and q_2 . In Figure 3, we then show how to implement the 3-SAT clause phase separator $e^{-i\gamma(x_0 \wedge x_1 \wedge x_2)}$. This operation corresponds to a 2-controlled $\text{CCPhase}(\gamma) = \text{diag}(1, 1, 1, 1, 1, 1, 1, e^{i\gamma})$ gate, which we decompose using 4




Fig. 5: Example QAOA circuit for $p = 1$ on a single random $n = 6$ 3-SAT instance, with clause density of 4. This circuit is defined in terms of the gateset h, x, cx, rzz, rz .




Fig. 6: Compiled quantum circuit drawing for $p = 1$ on a single random $n = 6$ 3-SAT instance, with clause density of 4. Compiled connectivity is all-to-all, and the target gateset is rzz, rz, ry, rz . This circuit is the (Qiskit) compiled form of Figure 5.

CNOTs and 1 Rzz gate, analogous to standard Toffoli decompositions [2], [47]. Since 5 two-qubit gates are necessary for 2-controlled phase shift gates [62], this decomposition is essentially optimal.

For k -SAT with $k > 3$, (ii) amounts to a $(k - 1)$ -controlled $C^{k-1}\text{Phase}(\gamma)$ gate. We can then iteratively decrease the number of controls by computing a pairwise AND into a $|0\rangle$ -initialized ancilla qubit. This is illustrated for 4-SAT in Figure 4: Each such step reduces the number of controls by one, using 3 CNOTs and one ancilla. As our base case we start with the $CC\text{Phase}$ gate from Figure 3. Finally, all ancilla need to have an AND^\dagger uncomputation. To further decrease the number of two-qubit gates, we make use of measurement-based uncomputations for k -SAT with $k > 3$, which needs mid-circuit measurements, qubit reset, and classical feed-forward control. This trade-off is beneficially for ion-traps, which have high measurement fidelities. Measurement-based uncomputation amounts to 1 CZ gate per ancilla (which as an added benefit also allows us to move the last CNOT of our base case into a classically controlled single-qubit gate).

Overall we need $k - 3$ ancilla qubits and $4k - 8$ two-qubit gates for any k -SAT clause with $k > 3$ and 5 two-qubit gates per clause for 3-SAT. Assuming a clause density of 4 for 3-SAT (justified later), we get a total of $20 \cdot n \cdot p$ two-qubit gates for a p -round QAOA circuit for a 3-SAT problem on n qubits, see Figures 5 and 6.

C. Quantum Hardware Implementation

The QAOA circuits are described using Qiskit [46], then transpiled using the Qiskit compiler in order to adapt the gateset to be as near to the quantum computer instruction set as possible. The final optimized compiled circuits that have had their instructions adapted are submitted to Quantinuum

H1-1 via the Quantinuum API with the circuits described as Open QASM [12] strings. The corresponding circuits are also submitted to IonQ Harmony via the Qiskit-IonQ plugin. The QAOA circuits are submitted to IonQ Aria 1 via the Python 3 Amazon Braket package. The QAOA circuits were executed on IonQ Forte by the internal IonQ team, with the QAOA circuits having been represented as OpenQASM files [12]. Table I shows the qubit count on the 4 quantum computers; the device qubit counts define the maximum 3-SAT problem size that can be executed on that device.

The circuits targeting Quantinuum H1-1 are compiled to have the gateset of rzz, rz, ry, rx . All four of these gates are single angle parameterized gates. The circuits targeting IonQ Harmony, IonQ Aria, and IonQ Forte are compiled to have the gateset of cx, rz, ry, rx . For circuit compilation to all 4 quantum computers, the arbitrary (e.g. all-to-all) connectivity is used since all of the devices use the ion trap architecture which offers an all-to-all connectivity graph. An example $p = 1$ 3-SAT QAOA circuit that has been compiled to the rzz, rz, ry, rx gateset is shown in Figure 6.

Figure 7 shows the scaling of the gate depth and single and two qubit gate counts for the (Qiskit compiled) QAOA circuits as a function of p and n , for the gateset of rzz, rz, ry, rx . The gate depth plots in Figure 7 are computed using Qiskit [46] (specifically `QuantumCircuit.depth`), and includes both single and two qubit gates. Interestingly, Figure 7 shows that as n increases for a fixed p , the gate depth does not substantially increase for larger n , but rather seems oscillate within a few percentage points around an average. This is notable because it implies that these QAOA circuits have fairly constant circuit depth per round even as we increase n ; this is due to our choice of constant clause densities (in our case 4).

QPU name	Number of Qubits	Avg. 1 qubit gate RB fidelity	Avg. 2 qubit gate RB fidelity	Avg. SPAM fidelity
IonQ Harmony	11	0.9958	0.9652	0.99752
Quantinuum H1-1	20	0.99996	0.998	0.997
IonQ Aria 1	25	0.981	0.9817	0.994
IonQ Forte	30	0.9998	0.99536	0.995

TABLE I: NISQ processor hardware summary used to execute the trained 3-SAT QAOA circuits.

For reference with respect to machine calibrations and features, all experiments were executed between April 2023 and June 2023. Quantinuum device backend parameters were set to default, which means that the pytket [53], [11] optimization level of 2 (which is the highest possible optimization level in pytket) is applied server side. IonQ backend parameters are also all set to default for all of the IonQ device runs.

All QAOA parameter combinations, e.g. for each p and n , are sampled using at least 40 shots (samples) on each quantum computer. This means that there will be shot noise present in the results.

For reference with respect to the expected device computation performance, we provide here the vendor-measured randomized benchmarking [35], [26] calibration metrics for single qubit gates, two qubit gates, and SPAM (State Preparation and Measurement) error rates in Table I. Table I also gives the QPU name and the number of qubits available on that device. Because all 4 of these quantum computers are trapped-ion based, their connectivity is all-to-all. The device calibration data for Quantinuum H1-1 was obtained from the device documentation during the timeframe in which the circuits were run. The device calibration data for IonQ Harmony and IonQ Aria 1 was obtained from the Amazon Braket device page within the time frame that the circuits were run. The device calibration data for IonQ Forte was measured using direct randomized benchmarking [45], and described in ref. [10].

D. Experimental Design

The essential question of “how well does QAOA perform?” is perhaps best thought of as two separate questions: “assuming one has good (or optimal) angles, what kind of approximation ratio can QAOA give?” and “how difficult is it to find good angles?”. This work is particularly interested in the first question, with an emphasis on real-world performance on current NISQ devices. In this section we give details as to how we found the angles which are then tested on hardware, without significantly commenting on the difficulty (i.e. algorithmic scaling) of the angle finding.

The combinatorial optimization problems that will be sampled on the NISQ computers are maximization random 3-SAT instances, with a clause density of 4. The selection of clause density is important because we need to have a balance between having too few clauses, which makes the optimization problem easy to solve, but each additional clause comes with substantial additional gate operations and higher circuit depths which means the computation will have more errors. We chose a clause density of 4 based on ref. [23], which shows that a

clause density of 4 is right at the start of the plateau of the 3-SAT problems being hard, which makes the problems hard in general to solve exactly and minimizes the number of circuit operations.

For a given MAX 3-SAT instance and number of rounds p , we selected QAOA angles via a classical optimization algorithm with (exact, error free and shot noise free) QAOA simulations. In particular, our classical angle-finding scheme relies on combining a global optimization algorithm known as basin-hopping [55] with local gradient descent. We found that 10 basin-hopping iterations produced high quality angles (though not necessarily optimal) up to $p = 20$. This QAOA angle optimization code is implemented using the Julia [3] programming language. More discussion of our angle-finding approach, and comparison with other angle-finding schemes, can be found in refs. [22], [23].

For large problem sizes and high rounds, simulating the QAOA circuits is quite computationally intensive and requires HPC resources. Furthermore, moving forward it is clear that evaluating QAOA circuits via real hardware, rather than simulation, will be necessary for this style of angle-finding. However, we note that in general the scalability of real time angle-finding for a hybrid quantum-classical feedback loop is not clear for variational quantum algorithms, and in particular is expected to be quite hard especially with the presence of noise [57], [4]. On the other hand, there is increasing evidence that classes of problems exhibit pronounced parameter concentration of the optimal or near optimal angles [19], [51], [1], [6], [61], [17], [48], [41], which in some cases could lead to scalable heuristics for good angle finding.

III. RESULTS

Figure 8 shows the scaling of mean approximation ratio as a function of increasing the number of QAOA rounds p , for $n = 10$, $n = 8$, and $n = 6$, respectively. The angles for each round were found using HPC classical computation in order to get very good, but not necessarily optimal, angles γ and β . The classically learned angles allow the QAOA computation, without noise, to increase in approximation ratio monotonically, as shown by the red lines in Figure 8. The $n = 6$ plot ends after $p = 12$ because the ideal computation of the approximation ratio reaches above 0.999, at which point the angle finding is terminated since the ideal computation has essentially plateaued and further rounds would not improve the computation substantially.

Because each additional round significantly adds to the total gate count, it is expected that the computation will accumulate

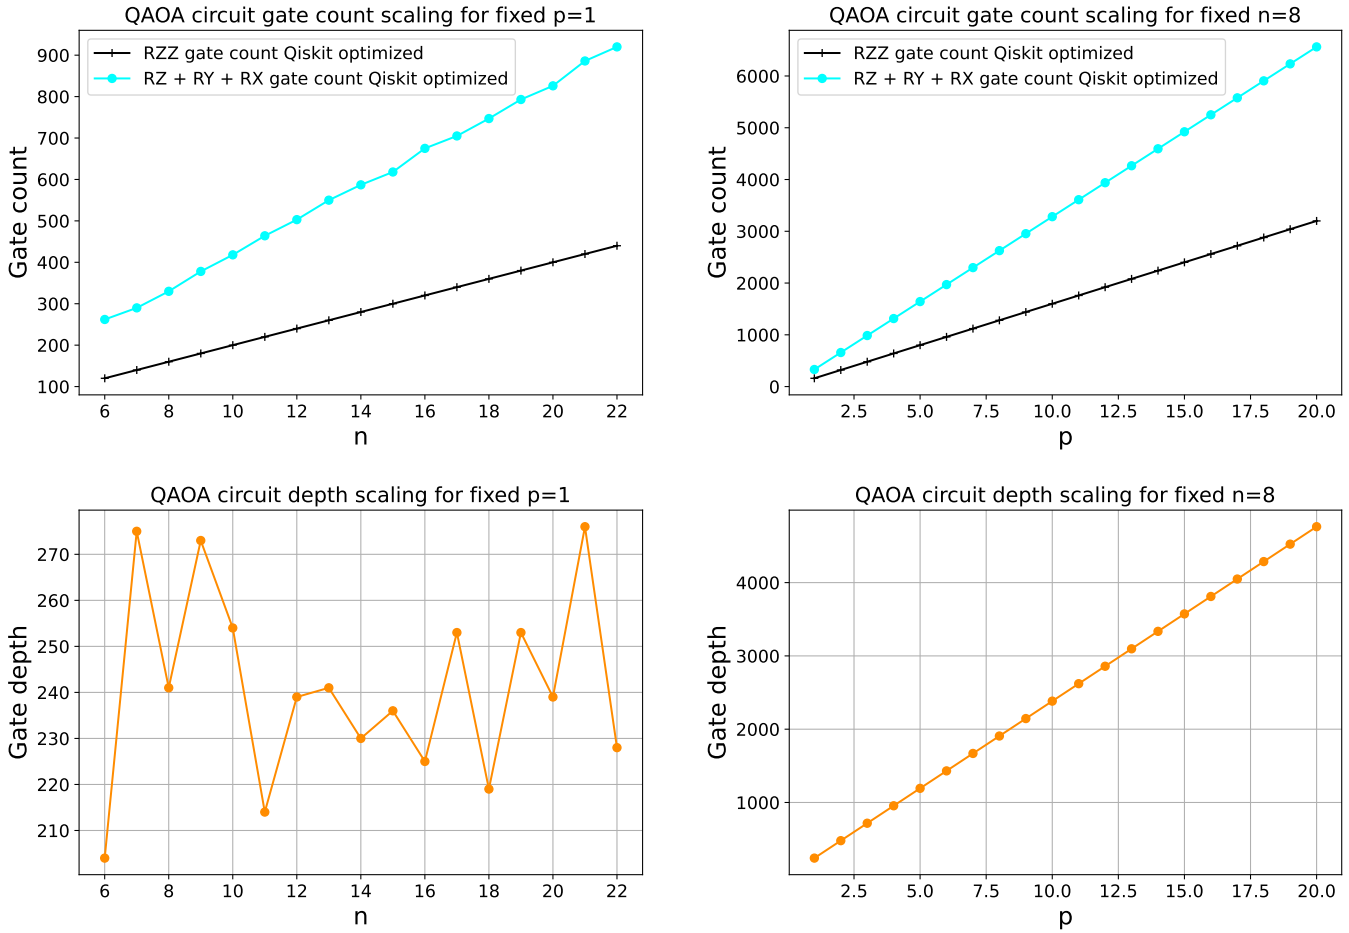


Fig. 7: Top right: plotting single and two qubit gate counts in the generated QAOA circuits as function of increasing p for a fixed $n = 8$ 3-SAT instance. The lower right plot shows the circuit depth (including both single and two qubit gates) for the same (increasing p) circuits as the top right plot. Top left: plotting the single and two qubit gate counts as a function of increasing the 3-SAT problem size (n) for a fixed $p = 1$. The lower left plot shows the circuit depth for the same increasing n circuits from the top left plot.

errors at very high rounds and converge to random sampling. This is indeed what is observed in Figure 8, in particular for $n = 10$ and $n = 8$. Notably, rounds $p = 1, 2, 3, 4$ for all three 3-SAT problem instances on Quantinuum H1-1 are relatively stable in their mean approximation ratios, not immediately decreasing, but also not increasing substantially as is expected for the ideal QAOA computation. In Figure 8 IonQ Harmony has a lower mean approximation ratio compared to Quantinuum H1-1, and in general there is not a clear signal that the sampled mean approximation ratio was measurably different than random sampling. The IonQ Aria samples in Figure 8 perform similarly to IonQ Harmony. There is little IonQ Forte data in Figure 8, but the data that is present shows at low rounds ($p = 1$ and $p = 2$) the sample quality is comparable to Quantinuum H1-1.

Figure 9 shows the scaling of mean approximation ratios as a function of increasing the 3-SAT problem size, for a fixed number of rounds $p = 1$, $p = 2$, and $p = 10$, respectively.

Note that Figure 9 goes up to $n = 20$ for both $p = 1$ and $p = 2$, but only up to $n = 19$ for $p = 10$: the reason for the missing datapoint at $n = 20$ for 10 rounds is that the classical angle finding computation was too time intensive to compute in a reasonable amount of time. Interestingly, on Quantinuum H1-1 we observe that there remains a constant separation between the ideal approximation ratio and the NISQ computer sampled approximation ratio as a function of n for both $p = 1$ and $p = 2$, suggesting that the computation does not significantly degrade as more qubits (e.g. variables) are added. Figure 7 shows that as n increases the gate depth does not significantly change, which in part explains why there remains this constant approximation ratio gap where Quantinuum H1-1 still performs reasonably well up to $n = 20$ for $p = 1, 2$. Figure 9 again shows that IonQ Harmony has a lower mean approximation ratio compared to the measured distribution from Quantinuum H1-1. The samples from IonQ Aria in Figure 9 are comparable to IonQ Harmony. The data

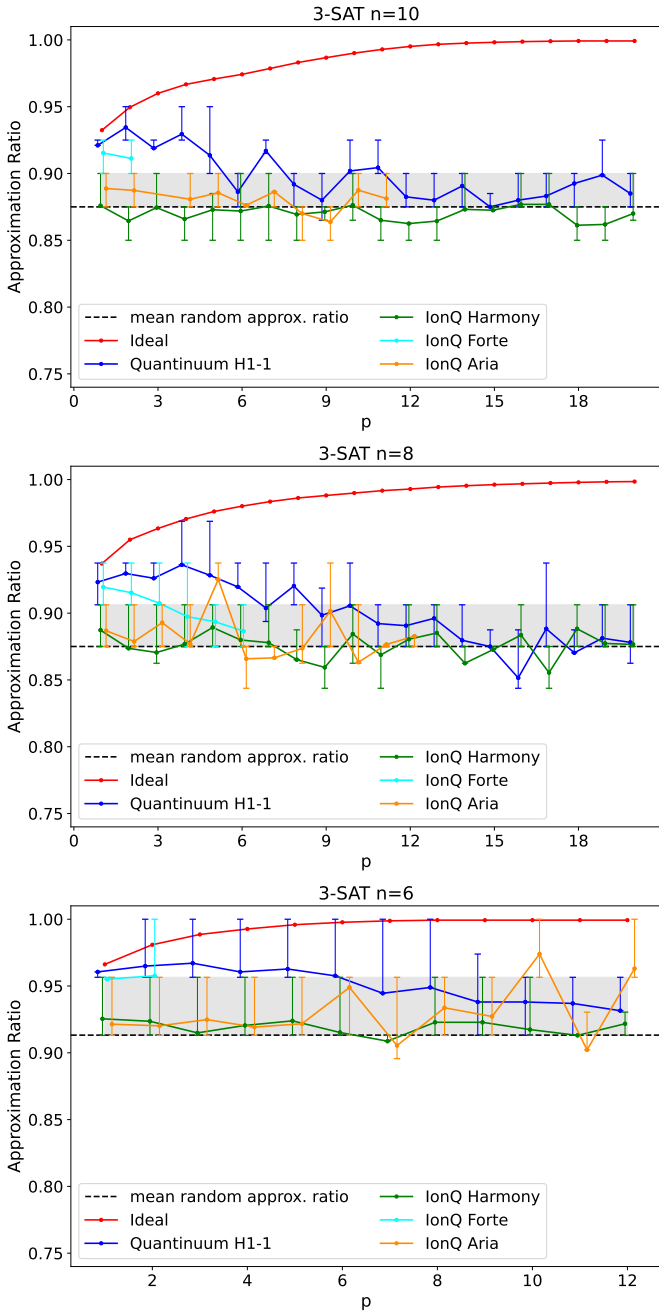


Fig. 8: Approximation ratio (y-axis) vs the number of rounds (x-axis) when executing the 3-SAT $n = 10$, $n = 8$, $n = 6$ QAOA circuits on the 4 trapped ion quantum computers. The mean approximation ratios are plotted as solid blue and green lines. The exact ideal approximation ratios for the classically computed angles (with no shot noise) are plotted in red. The 40th and 60th percentiles of random samples are shown by the shaded light grey region, the mean of the random sample distribution is plotted as the black dashed horizontal line.

from IonQ Forte in Figure 9 for $p = 1$ and $p = 2$ shows that the mean approximation ratios decrease to the noisy regime approximately at $n = 14$, but at small n IonQ Forte samples

within the same regime of Quantinuum H1-1 (e.g. slightly below the ideal QAOA approximation ratios).

Error bars are plotted in Figure 8 and 9 for all of the experiments. The error bars are asymmetric, and they are defined as the 40th and 60th percentiles of the sampled approximation ratio distributions so that the random sampling (light grey) shaded region can be compared against. The maximum ranges of the 40th and 60th percentile error bars are capped at the mean of the distribution for visually consistency of the plots. The error bars are visually slightly horizontally offset in order to make overlap of the distributions clear.

The random sample distribution used in Figure 8 and Figure 9 (plotted as the light grey regions and the dashed black lines) are generated using 100,000 (uniform) random variable assignment vectors, the approximation ratios of which are then evaluated.

For both Figure 8 and Figure 9, the number of shots used for IonQ Forte is 2500 for the tested QAOA parameters and problem sizes. For the other three quantum computers the number of shots used was 40 for each problem size and parameter combination, with the exception of the $n = 6, 8, 10$ plots in Figure 8 where a total of 140 shots were used for $p = 1, 2, 3, 4, 5$ (on IonQ Harmony, Quantinuum H1-1, and IonQ Aria) in order to obtain lower variance measurements at these low rounds in order to better determine whether the mean measured approximation ratio was clearly changing (increasing or decreasing) or staying constant. 140 shots were also sampled for $n = 8$, at $p = 6, 7, 8$ rounds on IonQ Harmony, Quantinuum H1-1, and IonQ Aria. Examining the behavior of the solution quality at lower rounds is important because the accumulation of errors per round has a clear impact (performance degradation) on the actual execution of the algorithm of the algorithm even at low rounds.

The relative performance difference between Quantinuum H1-1, IonQ Harmony, IonQ Forte, and IonQ Aria are consistent with the vendor provided calibration metrics listed in Table I. Note that the incomplete data from IonQ Aria and IonQ Forte (incomplete in that there are datapoints missing that could have been executed given the device qubit count, see Table I) is in part due to a current backend error that prevents extremely high gate count circuits from being executed on those two devices.

Let us return to our initial hypothesis of QAOA NISQ performance from Fig. 1. In both Figure 8 and Figure 9, we have added light grey shaded regions that show the range between the 40th and 60th percentiles when randomly sampling the 3-SAT instances to define p_{noise} . Our conjectured peak point p_{max} is somewhat hard to identify due to shot noise, nevertheless it would be $p_{\text{max}} = 5, 4, 2$ for $n = 6, 8, 10$, respectively, which overall is not as much of a dependence on instance size n . The round counts p_{noise} at which QAOA results become random are at $p_{\text{noise}} = 6, 7, 6$ for $n = 6, 8, 10$, respectively, as read off from Figure 8 with little or no dependence on n due the somewhat constant circuit depth. Therefore, a comparison against the random distribution becomes more meaningful when the problem size is significantly

IV. DISCUSSION AND CONCLUSION

Experimentally, these high round QAOA circuits are executed on the 4 trapped ion quantum computers, IonQ Harmony, Quantinuum H1-1, IonQ Aria, and IonQ Forte. These QAOA circuits contain thousands of gate instructions, and constitute the largest experimental evaluation of QAOA circuits, with respect to gate depth and gate count, to date. These results show that very high round QAOA computations degrade into being random noise due the accumulation of errors, however we also observed that was not pronounced solution quality degradation for increasing the problem size n for fixed small rounds $p = 1$ and $p = 2$. These experiments present a very strenuous evaluation of current quantum computing hardware, largely due to the random 3-SAT instances having significantly higher QAOA gate depth per round (and total instruction count) compared to previous QAOA experiments on problems such as maximum cut [52] which can have much lower gate depths per QAOA round.

Importantly, in these experiments the problem of angle finding on noisy quantum computers has been removed (e.g. offloaded to HPC simulations) with the intention of investigating the scaling of QAOA (with respect to problem size and number of rounds) on NISQ computers assuming reasonably good angles can be efficiently computed. However, future research could examine the angle learning aspect of these k -SAT circuits, along with associated questions such as what parameter concentrations can be observed, and whether those parameter concentrations offer a capability of transfer learning for k -SAT problems in QAOA as has been demonstrated for problems such as weighted maximum cut [50]. Sampling tasks with QAOA are sensitive to shot noise, therefore future larger scale QAOA investigations should also aim to increase the shot count significantly. Lastly, evaluating unexplored variants of QAOA mixers and state preparation algorithms applied to these k -SAT circuits would also be valuable future research.

The experiments we present here did not use all existing trapped-ion quantum computers that are publicly accessible. Future research could utilize other existing trapped-ion quantum computers for even larger system sizes such as Quantinuum H2-1 [37]. The primary difficulty for running this particular experiment type (e.g. heavily optimized QAOA circuits) on these larger systems, such as the 32 qubit Quantinuum H2-1 [37], is that the classical angle finding is approaching significant classical computational cost for an increasing number of QAOA rounds on high qubit count QAOA circuits. Therefore, future research on this topic will also need to include improved angle finding strategies, and optimized quantum circuit simulations for HPC systems.

V. ACKNOWLEDGMENTS

All figures were generated using quantikz [32], matplotlib [31] and Qiskit [46] in Python 3.

We acknowledge the use of IonQ cloud and access to IonQ Harmony provided by IonQ. The authors thank the Quantinuum team, the IonQ team, and the Amazon Braket team for their technical support.




Fig. 9: Approximation ratio (y-axis) vs the number of variables n (x-axis) in the 3-SAT problems for $p = 1$, $p = 2$, and $p = 10$ QAOA circuits executed on the 4 trapped ion quantum computers. The remaining figure details are the same as in Figure 8.

larger, e.g. the $n = 15$ and greater distributions in Figure 9. As we would expect, the NISQ computers return mostly noise at $p = 10$, however, when there is signal e.g., at $p = 2$, this signal is almost independent of problem size, at least for Quantinuum H1-1.

The data generated in these experiments and the 3-SAT problem instances are available as a public dataset [42].

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. The Oak Ridge Leadership Computing Facility provided access to the Quantinuum H1-1 computer. This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). Research presented in this article was supported by the NNSA’s Advanced Simulation and Computing Beyond Moore’s Law Program at Los Alamos National Laboratory. This research used resources provided by the Darwin testbed at Los Alamos National Laboratory (LANL) which is funded by the Computational Systems and Software Environments subprogram of LANL’s Advanced Simulation and Computing program (NNSA/DOE). We acknowledge the support of the LANL Information Science and Technology Institute. This work has been assigned LANL technical report number LA-UR-23-24237.

REFERENCES

- [1] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte. Parameter concentrations in quantum approximate optimization. *Physical Review A*, 104(1), jul 2021. URL: <https://doi.org/10.1103/2Fphysreva.104.1010401>, doi:10.1103/physreva.104.1010401.
- [2] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, Nov 1995. [arXiv:quant-ph/9503016](https://arxiv.org/abs/quant-ph/9503016), doi:10.1103/PhysRevA.52.3457.
- [3] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017. URL: <https://pubs.siam.org/doi/10.1137/141000671>, doi:10.1137/141000671.
- [4] Lennart Bittel and Martin Kliesch. Training variational quantum algorithms is NP-hard. *Physical Review Letters*, 127(12), sep 2021. URL: <https://doi.org/10.1103/2Fphysrevlett.127.120502>, doi:10.1103/physrevlett.127.120502.
- [5] Sami Boulebnane and Ashley Montanaro. Solving boolean satisfiability problems with the quantum approximate optimization algorithm, 2022. [arXiv:2208.06909](https://arxiv.org/abs/2208.06909).
- [6] Fernando G. S. L. Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances, 2018. [arXiv:1812.04170](https://arxiv.org/abs/1812.04170).
- [7] Andreas Bärttschi and Stephan Eidenbenz. Grover Mixers for QAOA: Shifting Complexity from Mixer Design to State Preparation. In *IEEE International Conference on Quantum Computing & Engineering QCE’20*, pages 72–82, 2020. [arXiv:2006.00354](https://arxiv.org/abs/2006.00354), doi:10.1109/QCE49297.2020.00020.
- [8] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, aug 2021. URL: <https://doi.org/10.1038/2Fs42254-021-00348-9>, doi:10.1038/s42254-021-00348-9.
- [9] Nicholas Chancellor. Domain wall encoding of discrete variables for quantum annealing and qaoa. *Quantum Science and Technology*, 4(4):045004, 2019. doi:10.1088/2058-9565/ab33c2.
- [10] Jwo-Sy Chen, Erik Nielsen, Matthew Ebert, Volkan Inlek, Kenneth Wright, Vandiver Chaplin, Andrii Maksymov, Eduardo Pérez, Amrit Poudel, Peter Maunz, and John Gamble. Benchmarking a trapped-ion quantum computer with 29 algorithmic qubits, 2023. [arXiv:2308.05071](https://arxiv.org/abs/2308.05071).
- [11] Alexander Cowtan, Silas Dilkes, Ross Duncan, Alexandre Krajenbrink, Will Simmons, and Seyon Sivarajah. On the Qubit Routing Problem. In Wim van Dam and Laura Mancinska, editors, *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:32, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/10397>, doi:10.4230/LIPIcs.TQC.2019.5.
- [12] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta. Open quantum assembly language, 2017. [arXiv:1707.03429](https://arxiv.org/abs/1707.03429).
- [13] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: A typical case, 2020. [arXiv:2004.09002](https://arxiv.org/abs/2004.09002).
- [14] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples, 2020. [arXiv:2005.08747](https://arxiv.org/abs/2005.08747).
- [15] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A Quantum Approximate Optimization Algorithm. *arXiv e-prints*, Nov 2014. [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [16] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, 2015. [arXiv:1412.6062](https://arxiv.org/abs/1412.6062).
- [17] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The quantum approximate optimization algorithm and the sherrington-kirkpatrick model at infinite size. *Quantum*, 6:759, jul 2022. URL: <https://doi.org/10.22331/q-2022-07-07-759>, doi:10.22331/q-2022-07-07-759.
- [18] Edward Farhi and Aram W Harrow. Quantum supremacy through the quantum approximate optimization algorithm, 2019. [arXiv:1602.07674](https://arxiv.org/abs/1602.07674).
- [19] Alexey Galda, Xiaoyuan Liu, Danylo Lykov, Yuri Alexeev, and Ilya Safro. Transferability of optimal qaoa parameters between random graphs. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 171–180, 2021. doi:10.1109/QCE52317.2021.00034.
- [20] Craig Gidney. Verifying Measurement Based Uncomputation, August 2019. URL: <https://algassert.com/post/1903>.
- [21] John Golden, Andreas Bärttschi, Daniel O’Malley, and Stephan Eidenbenz. Threshold-based quantum optimization. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, oct 2021. URL: <https://doi.org/10.1109/2Fqce52317.2021.00030>, doi:10.1109/qce52317.2021.00030.
- [22] John Golden, Andreas Bärttschi, Stephan Eidenbenz, and Daniel O’Malley. Numerical evidence for exponential speed-up of qaoa over unstructured search for approximate constrained optimization, 2023. [arXiv:2202.00648](https://arxiv.org/abs/2202.00648).
- [23] John Golden, Andreas Bärttschi, Daniel O’Malley, and Stephan Eidenbenz. The quantum alternating operator ansatz for satisfiability problems, 2023. [arXiv:2301.11292](https://arxiv.org/abs/2301.11292).
- [24] Stuart Hadfield, Tad Hogg, and Eleanor G Rieffel. Analytical framework for quantum alternating operator ansätze. *Quantum Science and Technology*, 8(1):015017, dec 2022. URL: <https://doi.org/10.1088/2F2058-9565/2Faca3ce>, doi:10.1088/2058-9565/aca3ce.
- [25] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms*, 12(2):34, 2019. [arXiv:1709.03489](https://arxiv.org/abs/1709.03489), doi:10.3390/a12020034.
- [26] Robin Harper, Ian Hincks, Chris Ferrie, Steven T. Flammia, and Joel J. Wallman. Statistical analysis of randomized benchmarking. *Physical Review A*, 99(5), may 2019. URL: <https://doi.org/10.1103/2Fphysreva.99.052350>, doi:10.1103/physreva.99.052350.
- [27] Matthew P Harrigan, Kevin J Sung, Matthew Neeley, Kevin J Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C Bardin, Rami Barends, Sergio Boixo, et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3):332–336, 2021. doi:10.1038/s41567-020-01105-y.
- [28] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [29] Zichang He, Ruslan Shaydulin, Shouvanik Chakrabarti, Dylan Herman, Changhao Li, Yue Sun, and Marco Pistoia. Alignment between initial state and mixer improves qaoa performance for constrained portfolio optimization, 2023. [arXiv:2305.03857](https://arxiv.org/abs/2305.03857).

- [30] Rebekah Herrman, Lorna Treffert, James Ostrowski, Phillip C Lotshaw, Travis S Humble, and George Siopsis. Impact of graph structures for qaoa on maxcut. *Quantum Information Processing*, 20(9):289, 2021. doi:10.1007/s11128-021-03232-8.
- [31] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi:10.1109/MCSE.2007.55.
- [32] Alastair Kay. Quantikz: A tikz library to typeset quantum circuit diagrams. Tutorial on the Quantikz Package. arXiv:1809.03842, doi:10.17637/rh.7000520.
- [33] Ryan LaRose, Eleanor Rieffel, and Davide Venturelli. Mixer-phaser ansätze for quantum optimization with hard constraints. *Quantum Machine Intelligence*, 4(2):17, 2022. doi:10.1007/s42484-022-00069-x.
- [34] Phillip C. Lotshaw, Hanjing Xu, Bilal Khalid, Gilles Buchs, Travis S. Humble, and Arnab Banerjee. Simulations of frustrated ising hamiltonians using quantum approximate optimization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 381(2241), dec 2022. URL: <https://doi.org/10.1098/rsta.2021.0414>, doi:10.1098/rsta.2021.0414.
- [35] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. Characterizing quantum gates via randomized benchmarking. *Physical Review A*, 85(4), apr 2012. URL: <https://doi.org/10.1103/PhysRevA.85.042311>, doi:10.1103/PhysRevA.85.042311.
- [36] Dmitri Maslov. Advantages of using relative-phase toffoli gates with an application to multiple control toffoli optimization. *Physical Review A*, 93(2):022311, 2016. arXiv:1508.03273, doi:10.1103/PhysRevA.93.022311.
- [37] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. Campora III au2, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen, A. Chernoguzov, E. Chertkov, J. Colina, J. P. Curtis, R. Daniel, M. DeCross, D. Deen, C. Delaney, J. M. Dreiling, C. T. Ertsgaard, J. Esposito, B. Estey, M. Fabrikant, C. Figgatt, C. Foltz, M. Foss-Feig, D. Francois, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, E. Glynn, A. Hall, A. M. Hankin, A. Hansen, D. Hayes, B. Higashi, I. M. Hoffman, B. Horning, J. J. Hout, R. Jacobs, J. Johansen, L. Jones, J. Karcz, T. Klein, P. Lauria, P. Lee, D. Liefer, C. Lytle, S. T. Lu, D. Lucchetti, A. Malm, M. Matheny, B. Mathewson, K. Mayer, D. B. Miller, M. Mills, B. Neyenhuis, L. Nugent, S. Olson, J. Parks, G. N. Price, Z. Price, M. Pugh, A. Ransford, A. P. Reed, C. Roman, M. Rowe, C. Ryan-Anderson, S. Sanders, J. Sedlacek, P. Shevchuk, P. Siegfried, T. Skripka, B. Spaun, R. T. Sprenkle, R. P. Stutz, M. Swallows, R. I. Tobey, A. Tran, T. Tran, E. Vogt, C. Volin, J. Walker, A. M. Zolot, and J. M. Pino. A race track trapped-ion quantum processor. 2023. arXiv:2305.03828.
- [38] Pradeep Niroula, Ruslan Shaydulín, Romina Yalovetzky, Pierre Minssen, Dylan Herman, Shaohan Hu, and Marco Pistoia. Constrained quantum optimization for extractive summarization on a trapped-ion quantum computer. *Scientific Reports*, 12(1):17171, 2022. doi:10.1038/s41598-022-20853-w.
- [39] Siyuan Niu and Aida Todri-Saniai. Multi-programming cross platform benchmarking for quantum computing hardware, 2022. URL: <https://arxiv.org/abs/2206.03144>, doi:10.48550/ARXIV.2206.03144.
- [40] Julian Obst, Johanna Barzen, Martin Beisel, Frank Leymann, Marie Salm, and Felix Truger. Comparing quantum service offerings: A case study of qaoa for maxcut, 2023. arXiv:2304.12718.
- [41] Elijah Pelofske, Andreas Bärtzchi, and Stephan Eidenbenz. Quantum annealing vs. QAOA: 127 qubit higher-order ising problems on NISQ computers. In *Lecture Notes in Computer Science*, pages 240–258. Springer Nature Switzerland, 2023. URL: https://doi.org/10.1007/978-3-031-32041-5_13, doi:10.1007/978-3-031-32041-5_13.
- [42] Elijah Pelofske, Andreas Bärtzchi, John Golden, and Stephan Eidenbenz. Dataset for High-Round QAOA for MAX k-SAT on Trapped Ion NISQ Devices, August 2023. doi:10.5281/zenodo.8201693.
- [43] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, apr 2021. URL: <https://doi.org/10.1038/2Fs41586-021-03318-4>, doi:10.1038/s41586-021-03318-4.
- [44] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. doi:10.22331/q-2018-08-06-79.
- [45] Timothy J. Proctor, Arnaud Carignan-Dugas, Kenneth Rudinger, Erik Nielsen, Robin Blume-Kohout, and Kevin Young. Direct randomized benchmarking for multiqubit devices. *Phys. Rev. Lett.*, 123:030503, Jul 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.030503>, doi:10.1103/PhysRevLett.123.030503.
- [46] Qiskit contributors. Qiskit: An open-source framework for quantum computing, 2023. doi:10.5281/zenodo.2573505.
- [47] Norbert Schuch and Jens Siewert. Programmable networks for quantum algorithms. *Physical Review Letters*, 91(2):027902, 2003. arXiv:quant-ph/0303063, doi:10.1103/PhysRevLett.91.027902.
- [48] Ruslan Shaydulín and Yuri Alexeev. Evaluating quantum approximate optimization algorithm: A case study. In *2019 Tenth International Green and Sustainable Computing Conference (IGSC)*, pages 1–6, 2019. doi:10.1109/IGSC48788.2019.8957201.
- [49] Ruslan Shaydulín, Changhao Li, Shouvanik Chakrabarti, Matthew De-Cross, Dylan Herman, Niraj Kumar, Jeffrey Larson, Danylo Lykov, Pierre Minssen, Yue Sun, Yuri Alexeev, Joan M. Dreiling, John P. Gaebler, Thomas M. Gatterman, Justin A. Gerber, Kevin Gilmore, Dan Gresh, Nathan Hewitt, Chandler V. Horst, Shaohan Hu, Jacob Johansen, Mitchell Matheny, Tanner Mengle, Michael Mills, Steven A. Moses, Brian Neyenhuis, Peter Siegfried, Romina Yalovetzky, and Marco Pistoia. Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem, 2023. arXiv:2308.02342.
- [50] Ruslan Shaydulín, Phillip C. Lotshaw, Jeffrey Larson, James Ostrowski, and Travis S. Humble. Parameter transfer for quantum approximate optimization of weighted MaxCut. *ACM Transactions on Quantum Computing*, 4(3):1–15, apr 2023. URL: <https://doi.org/10.1145/3584706>, doi:10.1145/3584706.
- [51] Ruslan Shaydulín, Kunal Marwaha, Jonathan Wurtz, and Phillip C. Lotshaw. Qaoakit: A toolkit for reproducible study, application, and verification of the qaoa. In *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)*, pages 64–71, 2021. doi:10.1109/QCS54837.2021.00011.
- [52] Ruslan Shaydulín and Marco Pistoia. Qaoa with $n \cdot p \geq 200$, 2023. URL: <https://arxiv.org/abs/2303.02064>, doi:10.48550/ARXIV.2303.02064.
- [53] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. t—ket_q: a retargetable compiler for nisq devices. *Quantum Science and Technology*, 6(1):014003, 2020. doi:10.1088/2058-9565/ab8e92.
- [54] Guang Song and Andreas Klappenecker. The simplified toffoli gate implementation by margolus is optimal, 2003. arXiv:quant-ph/0312225.
- [55] David J. Wales and Jonathan P. K. Doye. Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116, 1997. arXiv:https://doi.org/10.1021/jp970984n, doi:10.1021/jp970984n.
- [56] Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1), nov 2021. URL: <https://doi.org/10.1038/2Fs41467-021-27045-6>, doi:10.1038/s41467-021-27045-6.
- [57] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature communications*, 12(1):6961, 2021. doi:10.1038/s41467-021-27045-6.
- [58] Zhihui Wang, Nicholas C. Rubin, Jason M. Dominy, and Eleanor G. Rieffel. XY mixers: Analytical and numerical results for the quantum alternating operator ansatz. *Physical Review A*, 101(1):012320, 2020. arXiv:1904.09314, doi:10.1103/PhysRevA.101.012320.
- [59] Johannes Weidenfeller, Lucia C. Valor, Julien Gacon, Caroline Tornow, Luciano Bello, Stefan Woerner, and Daniel J. Egger. Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware. *Quantum*, 6:870, December 2022. doi:10.22331/q-2022-12-07-870.
- [60] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pistenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducre, A. Blinov, S. M. Kreikemeier,

- V. Chaplin, M. Keesan, C. Monroe, and J. Kim. Benchmarking an 11-qubit quantum computer. *Nature Communications*, 10(1), nov 2019. URL: <https://doi.org/10.1038/s41467-019-13534-2>, doi: [10.1038/s41467-019-13534-2](https://doi.org/10.1038/s41467-019-13534-2).
- [61] Jonathan Wurtz and Danylo Lykov. Fixed-angle conjectures for the quantum approximate optimization algorithm on regular maxcut graphs. *Phys. Rev. A*, 104:052419, Nov 2021. URL: <https://link.aps.org/doi/10.1103/PhysRevA.104.052419>, doi: [10.1103/PhysRevA.104.052419](https://doi.org/10.1103/PhysRevA.104.052419).
- [62] Nengkun Yu, Runyao Duan, and Mingsheng Ying. Five two-qubit gates are necessary for implementing the toffoli gate. *Physical Review A*, 88(1):010304, 2013. [arXiv:1301.3372](https://arxiv.org/abs/1301.3372), doi: [10.1103/PhysRevA.88.010304](https://doi.org/10.1103/PhysRevA.88.010304).
- [63] Bingzhi Zhang, Akira Sone, and Quntao Zhuang. Quantum computational phase transition in combinatorial problems. *npj Quantum Information*, 8(1), jul 2022. URL: <https://doi.org/10.1038/s41534-022-00596-2>, doi: [10.1038/s41534-022-00596-2](https://doi.org/10.1038/s41534-022-00596-2).
- [64] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.021067>, doi: [10.1103/PhysRevX.10.021067](https://doi.org/10.1103/PhysRevX.10.021067).
- [65] Yingyue Zhu, Zewen Zhang, Bhuvanesh Sundar, Alaina M. Green, C. Huerta Alderete, Nhung H. Nguyen, Kaden R. A. Hazzard, and Norbert M. Linke. Multi-round qaoa and advanced mixers on a trapped-ion quantum computer, 2022. [arXiv:2201.12335](https://arxiv.org/abs/2201.12335), doi: [10.1088/2058-9565/ac91ef](https://doi.org/10.1088/2058-9565/ac91ef).