

# Evidence of Scaling Advantage for the Quantum Approximate Optimization Algorithm on a Classically Intractable Problem

Ruslan Shaydulin,<sup>1,\*</sup> Changhao Li,<sup>1</sup> Shouvanik Chakrabarti,<sup>1</sup> Matthew DeCross,<sup>2</sup> Dylan Herman,<sup>1</sup> Niraj Kumar,<sup>1</sup> Jeffrey Larson,<sup>3</sup> Danylo Lykov,<sup>1,4</sup> Pierre Minssen,<sup>1</sup> Yue Sun,<sup>1</sup> Yuri Alexeev,<sup>4</sup> Joan M. Dreiling,<sup>2</sup> John P. Gaebler,<sup>2</sup> Thomas M. Gatterman,<sup>2</sup> Justin A. Gerber,<sup>2</sup> Kevin Gilmore,<sup>2</sup> Dan Gresh,<sup>2</sup> Nathan Hewitt,<sup>2</sup> Chandler V. Horst,<sup>2</sup> Shaohan Hu,<sup>1</sup> Jacob Johansen,<sup>2</sup> Mitchell Matheny,<sup>2</sup> Tanner Mengle,<sup>2</sup> Michael Mills,<sup>2</sup> Steven A. Moses,<sup>2</sup> Brian Neyenhuis,<sup>2</sup> Peter Siegfried,<sup>2</sup> Romina Yalovetzky,<sup>1</sup> and Marco Pistoia<sup>1</sup>

<sup>1</sup>*Global Technology Applied Research, JPMorgan Chase, New York, NY 10017, USA*

<sup>2</sup>*Quantinuum, Broomfield, CO 80021, USA*

<sup>3</sup>*Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA*

<sup>4</sup>*Computational Science Division, Argonne National Laboratory, Lemont, IL 60439, USA*

The quantum approximate optimization algorithm (QAOA) is a leading candidate algorithm for solving optimization problems on quantum computers. However, the potential of QAOA to tackle classically intractable problems remains unclear. In this paper, we perform an extensive numerical investigation of QAOA on the Low Autocorrelation Binary Sequences (LABS) problem. The rapid growth of the problem's complexity with the number of spins  $N$  makes it classically intractable even for moderately sized instances, with the best-known heuristics observed to fail to find a good solution for problems with  $N \gtrsim 200$ . We perform noiseless simulations with up to 40 qubits and observe that out to this system size, the runtime of QAOA with fixed parameters and a constant number of layers scales better than branch-and-bound solvers, which are the state-of-the-art exact solvers for LABS. The combination of QAOA with quantum minimum-finding on an idealized quantum computer gives the best empirical scaling of any algorithm for the LABS problem. We demonstrate experimental progress in compiling and executing QAOA for the LABS problem using an algorithm-specific error detection scheme on Quantinuum trapped-ion processors. Our results provide evidence for the utility of QAOA as an algorithmic component when executed on an idealized quantum computer.

## INTRODUCTION

Quantum computers have been shown to have the potential to speed up the solution of optimization problems. At the same time, only a small number of algorithmic primitives are known that provide broadly applicable speedups. These include amplitude amplification<sup>1,2</sup> and quantum walks more generally,<sup>3,4</sup> as well as the recently introduced short path algorithm.<sup>5,6</sup>

The dearth of provable speedups in quantum optimization motivates the development of heuristics. A leading candidate for demonstrating a heuristic speedup in quantum optimization is the quantum approximate optimization algorithm (QAOA).<sup>7,8</sup> QAOA uses two operators applied in alternation  $p$  times to prepare a quantum state such that, upon measuring it, a high-quality solution to the problem is obtained with high probability. A pair of such operators is commonly referred to as one QAOA “layer.” The state is evolved with a diagonal Hamiltonian encoding the optimization problem by the first operator and with a non-diagonal transverse-field Hamiltonian by the second operator. In this work, we consider the evolution times to be hyperparameters that are set by using a fixed, predetermined rule, analogously to the choice of a schedule in simulated annealing.

While QAOA has been studied extensively,<sup>9–12</sup> little is known about its potential to provide a scaling advantage over classical solvers. A recent numerical study<sup>11</sup> of random  $k$ -SAT with  $N \leq 20$  variables has shown that the time-to-solution (TTS) of QAOA with fixed param-

eters and constant depth grows as  $1.23^N$ . When QAOA is combined with amplitude amplification, the quantum TTS grows as  $1.11^N$ ,<sup>11</sup> whereas the best classical heuristic has TTS that grows as  $1.25^N$ .<sup>11</sup> Our work is motivated by this preliminary numerical evidence on small instances, which indicates that QAOA may potentially scale better than classical solvers when executed on an idealized quantum computer.

We study the scaling of QAOA TTS with the problem size on the Low Autocorrelation Binary Sequences (LABS) problem,<sup>13,14</sup> also known as the Bernasconi model in statistical physics.<sup>15,16</sup> The LABS problem has applications in communications engineering, where the low autocorrelation sequences are used for designing radar pulses.<sup>13,17</sup> To solve LABS, one has to produce a sequence of  $N$  bits that minimizes a specific quartic objective.

We choose LABS to study the scaling of QAOA TTS for the following three reasons. First, the complexity of LABS grows rapidly, with optimal solutions known only for  $N \leq 66$  and the best heuristics producing approximate solutions of quality decaying with  $N$  for  $N \gtrsim 200$ .<sup>18,19</sup> This makes it a promising candidate problem, since only a few hundred qubits are required to tackle classically intractable instances. Second, the performance of classical solvers for LABS has been benchmarked<sup>18,19</sup> in terms of the scaling of their TTS with problem size. We reproduce these results and observe that the scaling of classical solvers at  $N \leq 40$  matches the behavior at large  $N$  reported in the literature. This provides evidence that the scaling we observe

for QAOA at  $N \leq 40$  will similarly extrapolate to large  $N$ . Third, LABS has only one instance per problem size  $N$ . Combined with the hardness of LABS, this makes it possible to reliably study the scaling of QAOA at large problem sizes, where simulating tens or hundreds of random instances would be computationally infeasible.

We obtain the scaling by performing noiseless exact simulation of QAOA with fixed schedules. Our results are enabled by a custom algorithm-specific GPU simulator,<sup>20</sup> which we execute using up to 1,024 GPUs per simulation on the Polaris supercomputer accessed through the Argonne Leadership Computing Facility. We find that the TTS of QAOA with number of layers  $p = 12$  grows as  $1.46^N$ , which is improved to  $1.21^N$  if combined with quantum minimum-finding. This scaling is better than that of the best classical heuristic, which has a TTS that grows as  $1.34^N$ . We note that QAOA is a general quantum optimization heuristic with broad applicability, and no specific modifications have been done to adapt it to the LABS problem.

Our numerical evidence indicates that the proposed quantum algorithm scales better than the best classical heuristic in an idealized setting. However, we do not claim that QAOA is the best theoretically possible algorithm for the LABS problem. In particular, it may be possible to quadratically accelerate the best-known classical heuristic (Memetic Tabu<sup>21</sup>) by applying ideas similar to those used in quantum simulated annealing.<sup>3,22,23</sup> Nonetheless, our results highlight the potential of QAOA to act as a useful algorithmic component that enables super-Grover quantum speedups.

As a first step toward execution of QAOA for the LABS problem, we implement QAOA on Quantinuum trapped-ion quantum processors<sup>24,25</sup> on problems with up to  $N = 18$ . We further implement an algorithm-specific error detection scheme inspired by Pauli error detection<sup>26,27</sup> and demonstrate that it can reduce the impact of noise on solution quality by up to 65%. Our experiments highlight the continuing improvements to quantum computing hardware and the potential of algorithm-specific techniques to reduce the overhead of error detection and correction.

## PROBLEM STATEMENT

We begin by formally defining the LABS problem, discussing the state of the art of classical LABS solvers, and describing how QAOA is applied to solve the problem.

For a given sequence of spins  $s_i \in \{\pm 1\}$ , the autocorrelation is given as

$$\mathcal{A}_k(\mathbf{s}) = \sum_{i=1}^{N-k} s_i s_{i+k}. \quad (1)$$

The goal of the LABS problem is to find a sequence of

spins that minimizes the so-called ‘‘sidelobe’’ energy,

$$\mathcal{E}_{\text{sidelobe}}(\mathbf{s}) = \sum_{k=1}^{N-1} \mathcal{A}_k^2(\mathbf{s}), \quad (2)$$

or, equivalently, maximizes the merit factor

$$\mathcal{F}(\mathbf{s}) = \frac{N^2}{2\mathcal{E}_{\text{sidelobe}}(\mathbf{s})}. \quad (3)$$

The time-to-solution (TTS) is defined as the time a solver takes to produce this sequence. The energy  $\mathcal{E}_{\text{sidelobe}}(\mathbf{s})$  is a polynomial containing terms of degree 2 and 4, visualized in Fig. 1a. It can be encoded on qubits by the following Hamiltonian:

$$H_C = 2 \sum_{i=1}^{N-3} z_i \sum_{t=1}^{\lfloor \frac{N-i-1}{2} \rfloor} \sum_{k=t+1}^{N-i-t} z_{i+t} z_{i+k} z_{i+k+t} + \sum_{i=1}^{N-2} z_i \sum_{k=1}^{\lfloor \frac{N-i}{2} \rfloor} z_{i+2k}, \quad (4)$$

where  $z_j$  is a Pauli  $z$  operator acting on qubit  $j$ .

The runtimes of state-of-the-art classical solvers for the LABS problem scale exponentially, with clear exponential scaling present at  $N \leq 40$  as shown in Fig. 1b. The best-known exact solvers are branch-and-bound methods that have a running time that scales as  $1.73^N$ .<sup>19</sup> The best-known heuristic for general LABS is tabu search initialized with a memetic algorithm (Memetic Tabu)<sup>21</sup>, and has a running time that scales as  $1.34^N$ .<sup>28</sup>

To see why LABS is harder to solve than other commonly studied problems such as MaxCut, we can examine the correlation between the Hamming distance to the optimal solution and the objective. The comparison is shown in Fig. 1c. This correlation is one example of problem structure used by both classical and quantum heuristics to solve the problem quickly.<sup>7</sup> The absence of this correlation highlights the hardness of LABS compared with other commonly considered problems such as MaxCut.

As a consequence of the exponential scaling, the LABS problem becomes classically intractable at moderate sizes. Specifically, the value of the best-known merit factor decreases significantly for high  $N$ , whereas the asymptotic limit predicts that the merit factor should stay approximately constant.<sup>29</sup> This failure of state-of-the-art heuristics has been observed for  $N > 200$ .<sup>18,19</sup> The clear failure of the classical method to obtain high-quality solutions even at small sizes makes LABS an appealing candidate problem for quantum optimization heuristics.<sup>29</sup>

In this work, we tackle the LABS problem using QAOA. As shown in the circuit diagram Fig. 1d, QAOA solves optimization problems by preparing a parameterized state

$$|\boldsymbol{\beta}, \boldsymbol{\gamma}\rangle = \prod_{l=1}^p e^{-i\beta_l \sum_{j=1}^N x_j} e^{-i\gamma_l H_C} |+\rangle^{\otimes N}, \quad (5)$$

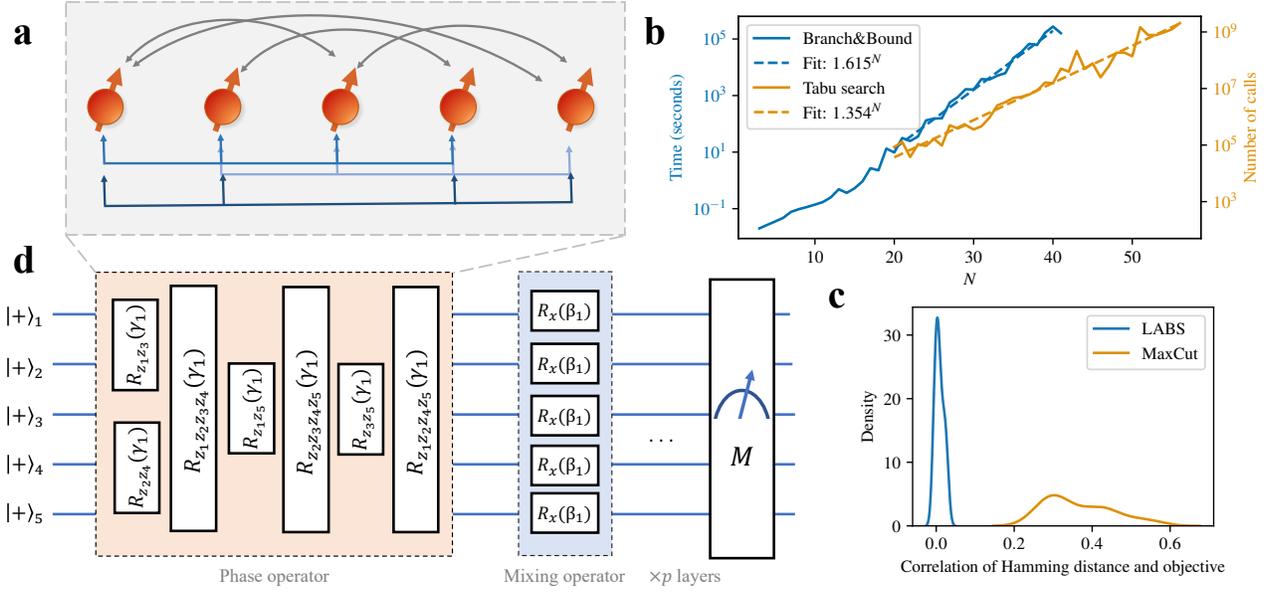


FIG. 1. **Classical and quantum algorithms applied to the LABS problem.** **a**, Diagram of the LABS problem (with example of  $N = 5$ ). The problem involves non-local two-body (black lines) and four-body (blue lines) interactions. **b**, Time-to-solution (TTS) of classical solvers. For the sizes considered, we observe clear exponential scaling with exponents matching their asymptotic values reported in the literature (see Table I). **c**, The distribution over  $21 \leq N \leq 31$  (for LABS) and 34 random instances (for MaxCut on random 3-regular graphs with 20 nodes) of Pearson product-moment correlation coefficients relating the Hamming distance of bitstrings from the optimal solution with the objective value of the bitstring. LABS has a much lower correlation between the Hamming distance and objective, indicating that it is much harder than the commonly considered MaxCut problem. **d**, Diagram of QAOA circuit for a 5-qubit example. Starting from a uniform superposition of the computational basis states, we apply  $p$  layers of phase and mixing operators, followed by measurement in the computational basis.

where  $|+\rangle^{\otimes N}$  is a uniform superposition over computational basis states,  $H_C$  is the diagonal Hamiltonian encoding the problem, and  $X_j$  is a Pauli  $x$  operator acting on qubit  $j$ . The operator  $e^{-i\gamma H_C}$  is commonly referred to as the phase operator and  $e^{-i\beta \sum_{j=1}^N X_j}$  as the mixing operator. The evolution times  $\beta, \gamma$  are hyperparameters chosen to maximize some figure of merit, such as the expected quality of the measurement outcomes or the probability of measuring the optimal solution. While  $\beta, \gamma$  can be optimized independently for each problem size, we consider them to be hyperparameters and use one fixed set of parameters for the LABS problem with a given QAOA depth  $p$  regardless of size. The fixed set of parameters is obtained by optimizing  $\beta, \gamma$  numerically for a number of small problem sizes and introducing an averaging and rescaling procedure to extrapolate parameters to any problem size (see the Methods section).

When choosing the parameters  $\beta, \gamma$  and evaluating the quality of the solution obtained by QAOA, two figures of merit are commonly considered. The first one is the expected merit factor of the sampled binary strings, given by

$$\langle C \rangle_{\text{MF}} = \langle \beta, \gamma | \frac{N^2}{2H_C} | \beta, \gamma \rangle = \sum_{s \in \{0,1\}^N} \text{Pr}(s) \mathcal{F}(s). \quad (6)$$

We will refer to  $\langle C \rangle_{\text{MF}}$  as the ‘‘QAOA energy’’ as a short-

hand. The second figure of merit is the probability of sampling the exact optimal solution, denoted by  $p^{\text{opt}}$  and equal to the sum of squared absolute values of amplitudes of basis states corresponding to exactly optimal solutions.

In the numerical experiments below, we follow the protocol of Ref. 11 and focus on scaling of the QAOA TTS with problem size  $N$  as the QAOA depth  $p$  is held constant. QAOA TTS is defined as  $\frac{1}{p^{\text{opt}}}$ , i.e. the expected number of measurements required to obtain an optimal solution from the QAOA state. Ref. 11 rigorously shows that, for random  $k$ -SAT, the runtime of constant-depth QAOA grows exponentially with  $N$  at any fixed  $p$ , with the scaling exponent depending on  $p$ . While the nature of the LABS problem makes it difficult to obtain analytical results analogous to Ref. 11, our numerical results also show clear exponential scaling of TTS. We note that, in practice, TTS of QAOA is  $\Theta(N^2) \frac{1}{p^{\text{opt}}}$ , where the  $\Theta(N^2)$  prefactor comes from the cost of implementing the LABS phase oracle.<sup>30</sup> However, we do not include it in our analysis because it does not affect the scaling exponent.

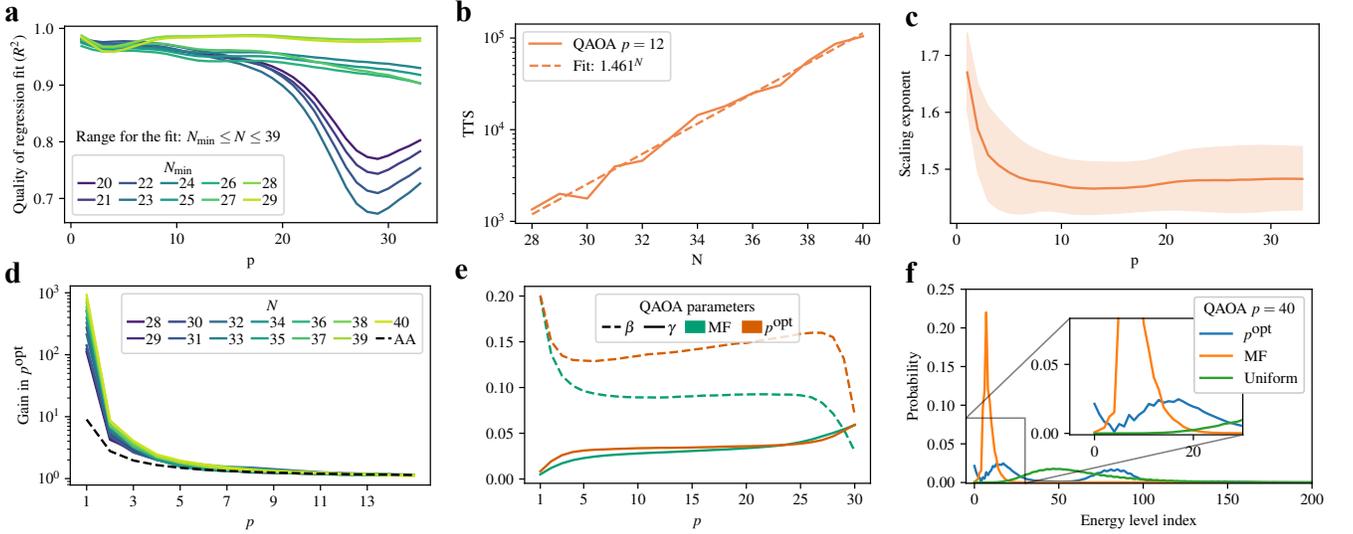


FIG. 2. **QAOA runtime scaling and dynamics under different parameter schedules.** **a**, The quality of the exponential fit for different choices of minimum  $N$  to include in the fit.  $N \geq 28$  results in a robust fit, the quality of which does not deteriorate with  $p$ .  $N = 40$  is omitted as it was only simulated up to  $p = 22$ . **b**, TTS of QAOA at  $p = 12$ . Clear exponential scaling is observed. **c**, The scaling exponent of QAOA runtime for different QAOA depths  $p$ . Shaded area shows 95% confidence interval. Increasing  $p$  beyond  $p \approx 12$  does not lead to better scaling. **d**, The gain in success probability  $p^{\text{opt}}$  from applying step  $p$  of QAOA and amplitude amplification (AA). The gain is defined as  $p_{\text{at step } p}^{\text{opt}}/p_{\text{at step } (p-1)}^{\text{opt}}$ . The gain at  $p = 1$  is over the random guess. Only one line is plotted for amplitude amplification since the lines for the values of  $N$  considered are visually indistinguishable. For small  $p$ , a QAOA layer gives orders of magnitude larger gain than a step of AA. **e**, Fixed QAOA parameters for  $p = 30$  chosen with respect to the QAOA energy  $\langle C \rangle_{\text{MF}}$  (“MF”) and probability of sampling the optimal solution (“ $p^{\text{opt}}$ ”). Different choice of optimization objective gives different resulting parameters. **f**, Probability of obtaining a binary string corresponding to a given energy level of the LABS problem (the zeroth energy level is the ground state or optimal solution; lower is better). When parameters are optimized with respect to the expected merit factor (labeled “MF”), the QAOA output state is concentrated around the mean and fails to obtain a high overlap with the ground state. On the other hand, when parameters are optimized with respect to  $p^{\text{opt}}$  (labeled “ $p^{\text{opt}}$ ”), the QAOA state has a high overlap with both the ground state and higher energy states. The probability of obtaining the ground state is 27.3 times greater for QAOA with parameters optimized with respect to  $p^{\text{opt}}$  at  $p = 40$ .

## SCALING OF QUANTUM TIME-TO-SOLUTION FOR LABS PROBLEM

We now present the numerical results demonstrating the scaling of TTS of QAOA and QAOA augmented with quantum minimum-finding (“QAOA+QMF”). The results are summarized in Table I. Throughout this section, we present the numerical results obtained using exact noiseless simulations. The runtime scaling is obtained by evaluating QAOA once with fixed parameters  $\beta, \gamma$  (i.e., with no overhead of parameter optimization) and computing the value  $p^{\text{opt}}$  with high precision. We discuss the parameter setting procedure and the details of simulation in the Methods section.

We are interested in the scaling of the runtime of QAOA for large problem sizes  $N$ . An important question to address is the choice of the smallest  $N$  to include in the scaling analysis, since the algorithm’s behavior at small sizes may not be predictive of its behavior at large sizes. Note that the largest  $N$  we include is limited by the capability of the classical simulator. We use the quality of the fit as the criterion for the choice of the cutoff on

$N$ . Figure 2a shows that if we set the cutoff at  $N \geq 28$ , we obtain a robust high-quality fit ( $R^2 > 0.94$ ), with the quality of the fit remaining stable as  $p$  grows. On the other hand, if smaller  $N$  are included, the quality of fit begins to decay with  $p$ . Therefore we include only  $N \geq 28$ , obtaining the fit presented in Fig. 2b. We observe that TTS of QAOA grows as  $1.46^N$  with problem size at constant QAOA depth  $p = 12$ . We present evidence that the scaling exponent for QAOA at  $p = 12$  is not sensitive to the choice of  $N_{\text{min}}$  in Supplementary Information.<sup>29</sup>

As a quantum optimization heuristic with constant depth, on a fault-tolerant quantum computer the QAOA performance can be improved by using amplitude amplification<sup>11,30</sup> or, more specifically, quantum minimum-finding<sup>31</sup> (see Methods). The resulting scaling of TTS of QAOA augmented with quantum minimum-finding (“QAOA+QMF”) is  $1.21^N$ .

We observe that, beyond a certain value ( $p \approx 12$ ), increasing QAOA depth does not lead to better scaling of TTS. This behavior is demonstrated in Fig. 2c. Consequently, running QAOA with  $p$  higher than 12 does

	QAOA+QMF	QAOA	Memetic Tabu		Branch-and-bound		
			Reproduced	<sup>21,28</sup>	Reproduced TTS	TTO	TTO <sup>19</sup>
Fit	<b>1.21</b>	1.46	1.35	1.34	1.62	1.76	1.73
CI	(1.19, 1.23)	(1.42,1.50)	(1.33,1.38)	N/A	(1.57,1.66)	(1.72,1.79)	N/A

TABLE I. Scaling exponents for quantum and classical algorithms. Confidence intervals (CIs) are 95%. The reported asymptotic exponential scaling of classical state-of-the-art solvers is reproduced at  $N \leq 40$ . For branch-and-bound, we include both the time to obtain a certificate of optimality (TTO reported in Ref. 19) as well as the much shorter time to find an optimal solution (TTS). We observe that QAOA with constant depth of  $p = 12$  augmented with quantum minimum-finding (“QAOA+QMF”) has better time-to-solution scaling than the best known classical heuristics.

not give any scaling advantage over amplitude amplification. This behavior is illustrated in Fig. 2d, which shows the increase in the success probability  $p^{\text{opt}}$  from applying a given step of QAOA and amplitude amplification. For amplitude amplification, at step  $p$  we have  $p^{\text{opt}} = (\sin((2p+1)\arcsin\sqrt{p_0}))^2$ , where  $p_0 = \frac{8}{2^N}$  is the initial (random guess) success probability.<sup>32</sup> Note that the 8 in the numerator is a consequence of a dihedral group symmetry, namely,  $D_4$ . While asymptotically equivalent, amplitude amplification performs better than a realistic generalized minimum-finding algorithm,<sup>31</sup> as the formula used here considers the scenario where we know which states to amplify (i.e., the optimal merit factor is known). We observe that for small  $p$ , a step (layer) of QAOA gives orders of magnitude larger increase in success probability than does a step of amplitude amplification, implying an even larger improvement over direct application of quantum minimum-finding. We provide additional details on comparison between QAOA and amplitude amplification in the Supplementary Information.<sup>29</sup>

We observe that the QAOA dynamics with parameters optimized for expected solution quality  $\langle C \rangle_{\text{MF}}$  and success probability  $p^{\text{opt}}$  are different. We plot the optimized parameters in Fig. 2e. We note that the parameters optimized with respect to one metric give performance that is far from optimal with respect to the other metric. This can be seen in Fig. 2f, which plots the energy distribution (with respect to the cost Hamiltonian) of the states appearing in the QAOA wavefunction weighted by probability. With the parameters optimized for  $\langle C \rangle_{\text{MF}}$ , the QAOA output distribution is concentrated around its mean, and the overlap with the ground state or  $p^{\text{opt}}$  is very small. On the other hand, when the parameters are optimized with respect to  $p^{\text{opt}}$ , the wavefunction is not concentrated and has large probability weight on the target ground state (i.e., high  $p^{\text{opt}}$ ). This comes at the cost of significant overlap with high-energy states, which leads to poor expected solution quality. In the Supplementary Information<sup>29</sup>, we discuss the behavior of QAOA with parameters optimized with respect to different objectives.

## EXPERIMENTS ON TRAPPED-ION SYSTEM

We now present the experimental results demonstrating the algorithmic and hardware progress toward the practical implementation of QAOA. Implementation of the phase operator is especially challenging for currently available quantum processors. It requires a large number of geometrically nonlocal two-qubit gates, demanding high gate fidelity.

Recent progress in trapped-ion platforms based on the QCCD architecture<sup>24,25,29</sup> has led to a rapid increase in the number of qubits while maintaining high fidelity, enabling large-scale QAOA demonstrations.<sup>33,34</sup> These systems implement two-qubit gates between arbitrary pairs of qubits by transporting ions into physically separate gate zones, resulting in high-fidelity two-qubit gates with low crosstalk. We leverage this progress to execute QAOA circuits for the LABS problem on Quantinuum H-series trapped-ion systems.

To implement the QAOA circuit shown in Fig. 1d, we have to implement the phase operator. The four-body terms in the phase operator are decomposed into CNOT gates and the native  $R_{zz}(\theta) = e^{-i\frac{\theta}{2}zz}$  rotation as shown in Fig. 3a. To reduce the cost of implementing both the two-qubit and four-qubit interaction terms, we optimize the circuit by greedily canceling CNOT gates (for algorithm details and gate count reduction see the Supplementary Information<sup>29</sup>). The resulting circuit containing CNOTs and  $R_{zz}$ s is then transpiled into the two-qubit  $R_{zz}$  gates and single-qubit gates that can be natively implemented by the trapped-ion system.<sup>29</sup> We remark that the number of two-qubit gates is  $\approx 10^3$  at  $N = 18$ , putting our experiments among the largest quantum optimization demonstrations on quantum hardware to date.<sup>33-37</sup>

In this work we execute QAOA circuits with  $p = 1$  using parameters  $\beta, \gamma$  optimized in noiseless simulation, followed by a projective measurement in the computational basis. In Fig. 3b, we show the energy probability distribution of measured bitstrings for  $N = 13$ . We observe a broad distribution due to the limited number of layers and experimental imperfections. Nevertheless, even at high  $N$ , where two-qubit gate count is high and the gate errors can be significant, we observe a clear signal that indicates that QAOA is outperforming random guess. This is shown in Fig. 3c, which presents the ex-

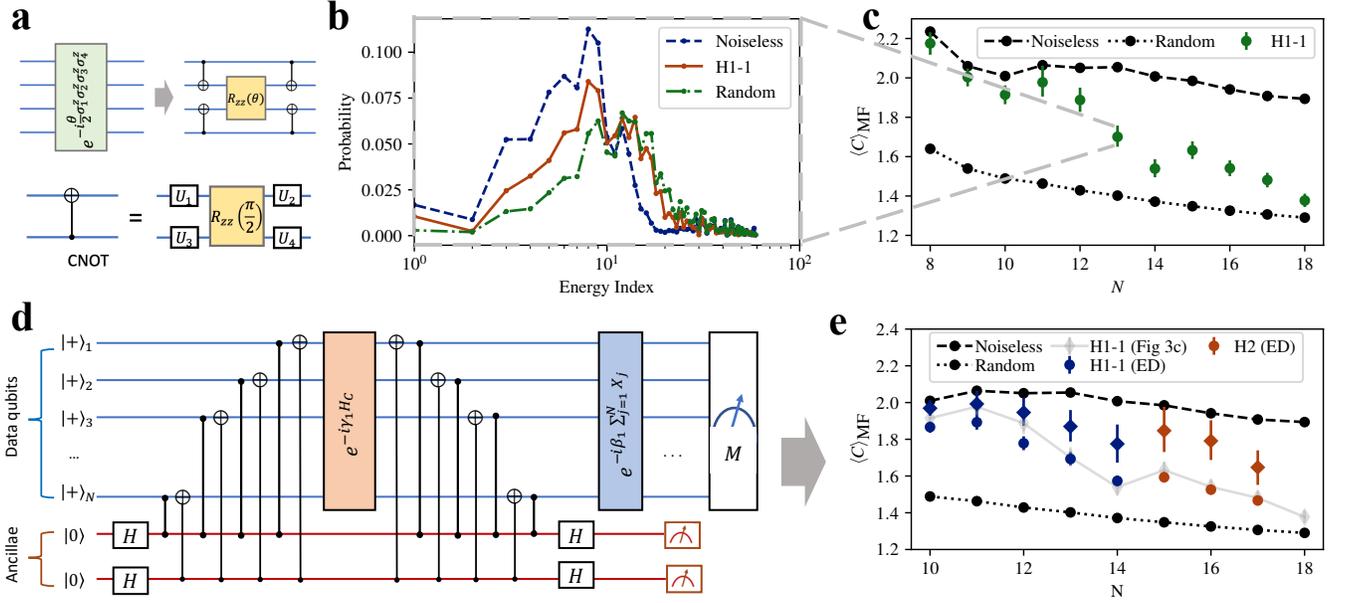


FIG. 3. **Experimental results on trapped-ion system.** **a**, Decomposition of four-body interaction terms into a two-body  $R_{zz}$  gate and four two-body CNOT gates, which can be realized via native  $R_{zz}$  gates. **b**, Energy density plot from experimental measured bitstrings for  $N=13$ . Energy index is arranged in energy ascending order. As a comparison, the distributions for noiseless  $p = 1$  QAOA simulation and random guess (assuming uniform distribution of all possible bitstrings) are shown. **c**, Experimental results up to 18 qubits on a trapped-ion quantum device (H1-1) with QAOA layer  $p = 1$  with optimized QAOA parameters. The error bars are calculated with 99% confidence intervals hereafter. **d**, Illustration of parity check circuit. The Z and X parities of states are mapped to ancillary qubits after implementation of full (or part of) phase operators via CZ and CNOT gates, respectively, followed by mid-circuit measurement on the ancillary qubits to extract the parity syndrome result. **e**, Experimental results for circuit with parity check. Three mid-circuit z-parity and x-parity checks were performed using six ancillary qubits. The ancillae can also be reused after appropriate reset during the circuit. The red data points were run on the Quantinuum H2 hardware while the blue data were from the H1-1 device. Data run on the H1-1 device without any ancillary qubits are shown in grey. Circles (diamonds) are the data without (with) post-selection. The abbreviation ED refers to the error detection via the parity checks. Number of mid-circuit parity checks is fixed to be two for  $N = 10, 11$  and three for all other  $N$ . Improvement in expected merit factor after post-selection according to parity syndrome measurement is observed.

perimentally obtained expected merit factors for various problem size up to  $N = 18$ . We note that the merit factor drops quickly for larger  $N$  and is approaching random guess because of experimental imperfections. We also note that at this scale LABS is easy for classical heuristics, which obtain optimal merit factors in  $< 1$  second.<sup>29</sup> Implementing QAOA for LABS instances that are hard for classical solvers would likely require error correction as the current implementation leads to an estimated two-qubit gate count of  $\approx 7.5 \times 10^5$  already at  $N = 67$  and  $p = 12$ .<sup>29</sup>

To improve the performance in the presence of noise, we implement an algorithm-specific error detection scheme. Since only the phase operator requires two-qubit gates, we focus on detecting errors that occur in the corresponding part of the circuit. Our scheme is based on the Pauli sandwiching error-detecting procedure of Ref. 26, which uses pairs of parity checks to detect some but not necessarily all errors that occur in a given part of the circuit. Following Refs. 38 and 39, we use the symmetries of the optimization problem to construct the parity checks. Specifically, we note that the

LABS Hamiltonian preserves both Z and X parities, that is,  $[H_C, \otimes_i^N z_i] = [H_C, \otimes_i^N x_i] = 0$ . We compute the parities onto ancillary qubits and perform mid-circuit measurement to determine whether an odd number of z- or x-flip errors occur during the circuit execution. The circuit with one check is shown in Fig. 3d. In the hardware experiments shown in Fig. 3e, we use up to three parity checks and observe consistent improvements in QAOA performance after postselecting on their outcomes. After postselection, the difference of merit factor between experimental results and noiseless simulation is reduced by 54% on average and up to 65% for specific  $N$ . In the Supplementary Information<sup>29</sup> we present additional details on the error-detecting scheme performance, including how performance improves with the number of parity checks and the reduction in the algorithm runtime. We note that while error detection does not directly give samples with better merit factors, the potential improvement in runtime can be translated into performance gains at the algorithm level, for example by being able to take more samples within a given time budget.<sup>29</sup> In our experiments, in all but two cases the optimal bitstring could

be found within the post-selected sample, and in all cases within the total sample.<sup>29</sup>

## DISCUSSION

Our main finding is that quantum minimum-finding enhanced with QAOA scales better than the best known classical heuristics for the LABS problem. This provides evidence for the potential of QAOA to act as a building block that provides algorithmic speedups on an idealized fault-tolerant quantum computer. We envision QAOA being used in a variety of algorithmic settings, similarly to how amplitude amplification acts as a subroutine in quantum algorithms for backtracking, branch-and-bound and so on.

We take the first step toward the execution of QAOA for the LABS problem by implementing an algorithm-specific error-detection scheme on a trapped-ion quantum processor. However, further improvements in quantum error correction and hardware are necessary to implement the quantum minimum-finding augmented with QAOA. In particular, the overheads of fault-tolerance<sup>40</sup> must be significantly reduced to realize the quantum speedup.

## METHODS

### Quantum minimum-finding enhanced with QAOA

In this work, we present the scaling results for QAOA combined with amplitude amplification (AA), or, more specifically, with quantum minimum-finding (“QAOA+QMF” in Table I). This reduces the scaling exponent by half as compared to directly sampling QAOA output. We now discuss in detail how QAOA is combined with the generalized quantum minimum-finding algorithm of Ref. 31 to obtain the stated scaling.

We begin by noting that standard AA is not sufficient. This is because the LABS problem is framed as optimization and not search, i.e. there is no oracle for marking a global minimum. The trick for handling optimization is to perform a standard reduction from optimization to feasibility. The reduction is performed by introducing a threshold on the cost as a constraint and performing a binary search using AA as a subroutine. The oracle used by AA marks the elements below the current threshold. This reduction was first introduced by Dürr and Høyer (DH).<sup>1</sup> However, the quantum minimum-finding algorithm of Dürr and Høyer utilizes standard Grover search, i.e. it requires the initial state to be the uniform superposition. A modification to it is required to leverage the improved success probability afforded by QAOA.

Ref. 31 provided a simple extension of DH that allows arbitrary initial states, with the overall cost scaling inversely with overlap between the initial state and state encoding the optimal solution. We leverage this

extension in our quantum algorithm. We use constant-depth QAOA to prepare the initial state for the quantum minimum-finding algorithm. As QAOA state has overlap with the optimal state that is much larger than that of uniform superposition<sup>29</sup> and scales more favorably, we obtain better performance than the direct minimum-finding of Dürr and Høyer. Specifically, we provide numerical evidence that our algorithm obtains a super-Grover speedup over exhaustive search for the LABS problem, and scales better than the best known classical heuristics. We present our modification to include QAOA for outputting an optimal solution  $x^*$  to the LABS problem in Algorithm 1 below. It is based on the generalized minimum-finding procedure outlined in Lemma 48 of Ref. 31. To keep the current work self-contained, we include the analysis of the algorithm below. We will use the following standard quantum subroutine based on Grover search that searches for an element with unknown probability in a quantum state.

**Lemma 1** (Exponential Quantum Search, Ref. 41). *Let  $|\psi\rangle = U|0\rangle^{\otimes N}$  be a quantum state in a  $2^N$ -dimensional Hilbert space with computational basis elements indexed by  $N$ -bit bitstrings, and  $m: \{0, 1\}^N \rightarrow \{0, 1\}$  be a marking function such that  $\sum_{\{x|m(x)=1\}} |\langle\psi|x\rangle|^2 \geq p$ . There exists a quantum algorithm **EQSearch**( $U, m, \delta$ ) that outputs an element  $x^*$  such that  $m(x^*) = 1$  with probability at least  $\delta$  using  $O\left(\frac{1}{\sqrt{p}} \log\left(\frac{1}{\delta}\right)\right)$  applications of  $U$  and  $m$ .*

---

### Algorithm 1 QAOA Enhanced with Quantum Minimum-Finding

---

**Require:** Unitary  $U_{\text{QAOA}}$  acting on  $\mathbb{C}^{2^N}$  such that  $|\langle x^*|U_{\text{QAOA}}|0\rangle^{\otimes N}| \geq 1/\sqrt{p_{\text{opt}}}$  for unknown  $p_{\text{opt}}$ ,  $V_{\text{LABS}}$  for computing  $\mathcal{E}_{\text{sidelobe}}$  into a register, and  $\delta \in (0, 1)$ , positive number  $M \leq 2^N$ ,  $C$  is the constant corresponding to the  $O(\cdot)$  in Lemma 1

**Ensure:** If  $M$  is greater than  $1/\sqrt{p_{\text{opt}}}$ , output  $x^*$  with  $\geq 1 - \delta$  probability using  $O(\log(1/\delta)M)$  calls to  $U_{\text{QAOA}}$  and  $V_{\text{LABS}}$  (and their inverses).

$x_{\text{res}}$  is set to an empty list.

**for**  $i \leftarrow 1$  **to**  $\lceil \log(1/\delta) \rceil$  **do**

$t \leftarrow 0$ ;  $s_0 \leftarrow \infty$

**while** number of calls to  $U_{\text{QAOA}}$  &  $V_{\text{LABS}}$  is  $< 3CMN$  **do**

$t \leftarrow t + 1$

Define  $m_t: \{0, 1\}^N \rightarrow \{0, 1\}$  such that  $m_t(x) = 1$  if and only if  $\mathcal{E}_{\text{sidelobe}}(x) < s_{t-1}$ . Note that  $m_t$  can be coherently evaluated using one query to  $V_{\text{LABS}}$ .

Set  $s_t = \mathbf{EQSearch}(U_{\text{QAOA}}, m_t, 1/(6 \cdot 2^N))$ .

**end while**

Append  $s_t$  to  $x_{\text{res}}$ .

**end for**

Output minimum of  $x_{\text{res}}$ .

---

**Theorem 1.** *Suppose a constant-depth QAOA circuit  $U_{\text{QAOA}}$  prepares a state  $|\psi\rangle = U_{\text{QAOA}}|0\rangle^{\otimes N}$  with  $N \geq 3$ , such that we have  $|\langle x^*|\psi\rangle| \geq 1/\sqrt{p_{\text{opt}}}$ , where  $|x^*\rangle$*

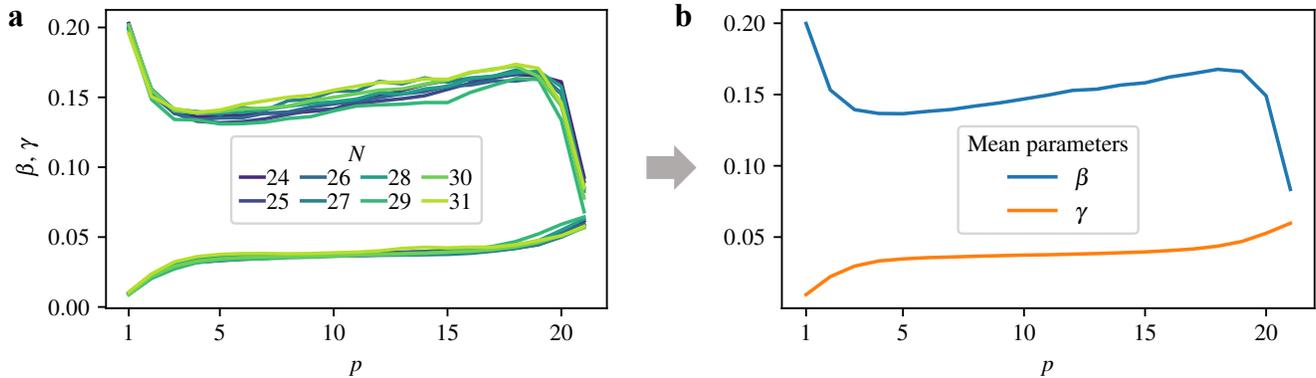


FIG. 4. **Visualization of how the fixed parameters are obtained.** **a**, Optimized QAOA parameters  $\beta$  (top lines) and  $\gamma$  (bottom lines) for  $p = 21$ .  $\gamma$  is multiplied by  $N/24$  (constant factor of  $\frac{1}{24}$  added for figure readability in both subfigures). **b**, Fixed parameters obtained by taking the arithmetic mean over the optimized parameters.

encodes an optimal solution to the  $N$ -bit LABS problem in a computational basis state, and we assume that  $p_{\text{opt}} \geq 1/N$ . Then, running Algorithm 1 with parameters  $M \geq 1/\sqrt{p_{\text{opt}}}$  and failure probability  $\delta$ , runs with a gate complexity of  $O(\text{poly}(N)\log(1/\delta)M)$  and finds  $x^*$  with probability at least  $1 - \delta$ .

*Proof.* See Supplementary Information.<sup>29</sup>  $\square$

### Choice of QAOA parameters $\beta, \gamma$

Our strategy for setting the QAOA parameters  $\beta, \gamma$  used in our experiments is twofold. First, we optimize QAOA parameters for small  $N$  using the FOURIER reparameterization scheme of Ref. 9. Second, we use the optimized parameters for small  $N$  to compute fixed QAOA parameters that are then used for larger  $N$ . To apply the fixed parameters to an instance with a given size  $N$ , we rescale the parameters  $\gamma$  by  $N$ .<sup>29</sup> We discuss the parameter optimization scheme and the parameter rescaling in the Supplementary Information.<sup>29</sup> We note that the results presented above can be improved if better parameter setting strategies are used.

The procedure for obtaining the set of fixed QAOA parameters is visualized in Fig. 4. Specifically, we optimize QAOA parameters for a set of small instances with sizes  $\{N_j\}_{j=1}^M$  attainable in simulation and set the fixed parameters to be the mean over the optimized parameters:

$$\beta^{\text{Fixed}} = \frac{1}{M} \sum_{j=1}^M \beta_{N_j}^*, \quad (7)$$

$$\gamma^{\text{Fixed}} = \frac{1}{M} \sum_{j=1}^M N_j \gamma_{N_j}^*, \quad (8)$$

where  $\beta_{N_j}^*, \gamma_{N_j}^*$  are the QAOA parameters optimized for the LABS instance of size  $N_j$  and  $M$  is the number of optimized instances. Then the parameters used in QAOA

for size  $N$  are given by  $\beta^{\text{Fixed}}, \frac{\gamma^{\text{Fixed}}}{N}$ . We use  $24 \leq N_j \leq 31$  ( $M = 8$ ).

### Error detection by symmetry verification

The error detection scheme relies on the symmetry of phase operator defined by Eq. 2. As it commutes with both  $\otimes_i^N Z_i$  and  $\otimes_i^N X_i$  operators, one can measure the value of these operators and perform postselection on the measurement outcomes. That is, the state after the phase operator should have the same  $Z$  and  $X$  parity as before it. In the presence of an odd number of bit flip or phase flip errors that occur during the implementation of phase operators, the resulting state will not be in the  $+1$  eigenspace of the two syndrome operators.

Experimentally, we divide the whole phase operator into  $m$  splits such that each split has approximately the same number of two-qubit gates, and we perform syndrome checks at the end of each split to detect errors. The syndrome operators are mapped to ancillary qubits via sequential controlled- $X$  or controlled- $Z$  gates and Hadamard gates applied before and after the partial phase operator. Since the number of two-qubit gates for the phase operators is higher than the number of gates used for the mapping by 2–3 orders of magnitude, additional errors introduced by ancillae are negligible. Furthermore, the crosstalk error probability during mid-circuit measurements is on the order of  $10^{-5}$ , considerably lower than the typical two-qubit-gate infidelity of  $2 \times 10^{-3}$  for the trapped-ion systems we used.<sup>25</sup> As a result, our error detection scheme leads to large improvements in QAOA performance on hardware at the cost of the number of repetitions growing exponentially with the number of checks.<sup>26</sup> We note that the performance of the error detection scheme can be further improved by implementing parity checks using fault-tolerant constructions.<sup>42</sup>

## Scaling of classical solvers

All scaling coefficients are obtained by fitting a least-squares linear regression on the logarithm of TTS. The confidence intervals on the scaling coefficients are obtained by using the Student's  $t$ -distribution and are reported with 95% confidence.

We use commercial state-of-the-art branch-and-bound solvers in numerical experiments. Figure 1b and Table I show results obtained using Gurobi,<sup>43</sup> although we obtain similar results for CPLEX<sup>44</sup> (see the Supplementary Information). The use of commercial branch-and-bound solvers is motivated by the observation that their scaling closely matches that reported in Ref. 19. Specifically, we observe that for both solvers the time to produce a certificate of optimality (TTO) scales with an exponent within a 95% confidence interval of the 1.73 exponent reported in Ref. 19. We note that unlike the solver presented in Ref. 19, commercial solvers are not parallelizable and can take advantage of only one CPU with at most tens of cores. Since QAOA is a heuristic and does not guarantee optimality, we additionally run branch-and-bound solvers until a solution with an exactly optimal merit factor is found, at which point the execution is stopped. The resulting TTS scales more favorably: for Gurobi, the scaling is  $1.615^N$ , with a 95% confidence interval of (1.571, 1.659). All the numbers reported correspond to the mean CPU time, with the mean taken over 100 random seeds for  $N \leq 32$  and 10 random seeds for  $N > 32$ . We present additional details of classical solver benchmarking in the Supplementary Information.<sup>29</sup>

Branch-and-bound algorithms are the best-known exact solvers for the LABS problem. In the regime where proving optimality is out of reach and the goal is simply to efficiently obtain sequences with high merit factors, heuristic algorithms are preferable. The best runtimes and runtime scaling reported in the literature<sup>21</sup> are from an algorithm known as Memetic Tabu. Memetic Tabu is a *memetic algorithm*, that is, an evolutionary algorithm augmented by local search. Specifically, an evolutionary algorithm is used to find initializations for tabu search, a metaheuristic that augments local neighborhood search with a data structure (known as the tabu list) that filters possible local moves if the potential solutions have been recently visited or diversification rules are violated.<sup>45</sup> In terms of the runtime required to find optimal solutions in the regime where exact solutions have been found using branch-and-bound methods,<sup>19</sup> Memetic Tabu has been observed to outperform both non-evolutionary methods as well as memetic algorithms that use simpler neighborhood search schemes such as steepest descent. To verify the scaling of tabu search on the regime of interest for comparison with QAOA, we use the implementation of Memetic Tabu in Ref. 18. For each length, we average the runtime over 50 random seeds, obtaining the scaling of the time-to-solution of  $1.35^N$  with a 95% confidence interval of (1.33, 1.38). This scaling closely matches the one reported in Ref. 28. We also note that solvers based on

self-avoiding random walks<sup>18</sup> have been shown to be competitive with or outperform Memetic Tabu when the task is to find skew-symmetric<sup>29</sup> sequences with the lowest autocorrelation. These solvers are specialized to search for skew-symmetric sequences and do not naturally extend to the unrestricted LABS problem.

## High-performance simulation of QAOA

Our numerical results are enabled by a custom scalable high-performance algorithm-specific QAOA simulator. We briefly describe the simulator here; for additional details and benchmarks comparing the developed simulator with the state-of-the-art methods for simulating QAOA the reader is referred to Ref. 20.

In this work, the main goal of the numerical simulation of QAOA is to evaluate the expectation of the cost Hamiltonian  $\langle C \rangle_{\text{MF}}$  and the success probability  $p^{\text{opt}}$ . Since  $p^{\text{opt}}$  is exponentially small, it has to be evaluated with high precision. While many techniques can be leveraged for exact simulation, we opt to directly simulate the full quantum state as it is propagated through the QAOA circuit. We note in particular that tensor network techniques do not provide a benefit in this case since the circuit we simulate is deep and fully connected (see Ref. 20 for detailed comparison).

First, we leverage the observation that the cost Hamiltonian and hence the phase-separation operator are diagonal. This allows us to precompute the cost function evaluated at every binary input and multiply the exponentiated costs elementwise with the statevector to simulate the application of the phase-separation operator. This operation can be easily parallelized since it is an elementwise operation local to each element in the statevector. The same precomputed vector of cost function values is used to compute  $\langle C \rangle_{\text{MF}}$  by taking the inner product with the final QAOA state. The cost of precomputation is amortized over the large number of objective evaluations performed during parameter optimization and is thereby negligible.

Second, we note that the mixing operator consists of an application of a uniform  $x$  rotation applied on each qubit. Therefore, each rotation operation can be computed by multiplying a fixed  $2 \times 2$  unitary matrix with a  $2 \times 2^{n-1}$  matrix constructed from reshaping the statevector. This step is parallelized by grouping the pairs of indices on which the  $2 \times 2$  unitary is applied.

We perform the simulations on the Polaris supercomputer located in Argonne Leadership Computing Facility. We distribute the simulation to 256 Polaris nodes with four NVIDIA A100 GPUs on each node and one AMD EPYC CPU. The CPU is used to manage the communication and the assembly of final results. Each GPU hosts a chunk of the full statevector and a chunk of the integer cost operator vector. Application of the cost operator does not require any communication since it is local to each element. The grouping in the mixing oper-

ator depends on index  $i$  of the operator  $x_i$  analogous to the grouping in the fast Walsh–Hadamard transform.<sup>46</sup> For  $i \leq n - \log_2(1024) = 29$  the pairing is local within each GPU. For  $i > 29$  we use CUDA-enabled MPI to distribute full chunks between nodes, which requires space to be reserved for two statevector chunks on each GPU.

## AUTHOR CONTRIBUTIONS

R. Shaydulin devised the project. J. Larson, N. Kumar, and R. Shaydulin implemented QAOA parameter optimization and the parameter setting schemes. R. Shaydulin and Y. Sun implemented the single-node version of the QAOA simulator. D. Lykov implemented the distributed version of the QAOA simulator and executed the large-scale simulations on Polaris. R. Shaydulin analyzed the simulation results. D. Herman and S. Hu developed the circuit optimization pipeline. C. Li implemented and analyzed the error detection scheme. M. DeCross and D. Herman executed the experiments on trapped-ion hardware, and C. Li analyzed the results. S. Chakrabarti, D. Lykov and P. Minssen benchmarked classical solvers. S. Chakrabarti, D. Herman and R. Shaydulin analyzed the generalized quantum minimum-finding enhanced with QAOA. J. Dreiling, J.P. Gaebler, T.M. Gatterman, J.A. Gerber, K. Gilmore, D. Gresh, N. Hewitt, C.V. Horst, J. Johansen, M. Matheny, T. Mengle, M. Mills, S.A. Moses, B. Neyenhuis, and P.

Siegfried built, optimized, and operated the trapped-ion hardware. M. Pistoia led the overall project. All authors contributed to technical discussions and the writing of the manuscript.

## ACKNOWLEDGMENTS

The authors thank their colleagues at the Global Technology Applied Research center of JPMorgan Chase for support and helpful discussions. Special thanks are also due to Tony Uttley and Jenni Strabley from Quantinuum for their continued support throughout the project. This material is based upon work supported in part by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357 and the Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing program.

## DATA AVAILABILITY

The data presented in this paper can be found at <https://doi.org/10.5281/zenodo.8190275>.

## CODE AVAILABILITY

The code used to produce the results in this paper can be found at <https://github.com/jpmorganchase/QOKit>.

---

\* Corresponding author. Email: [ruslan.shaydulin@jpmchase.com](mailto:ruslan.shaydulin@jpmchase.com)

<sup>1</sup> Christoph Dürr and Peter Høyer, [arXiv:quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014) (1996).

<sup>2</sup> S. Chakrabarti, P. Minssen, R. Yalovetzky, and M. Pistoia, [arXiv:2210.03210](https://arxiv.org/abs/2210.03210) (2022).

<sup>3</sup> R. D. Somma, S. Boixo, H. Barnum, and E. Knill, *Physical Review Letters* **101** (2008).

<sup>4</sup> P. Wocjan and A. Abeyesinghe, *Physical Review A* **78** (2008).

<sup>5</sup> M. B. Hastings, *Quantum* **2**, 78 (2018).

<sup>6</sup> A. M. Dalzell, N. Pancotti, E. T. Campbell, and F. G. Brandão, in *Proceedings of the ACM Symposium on Theory of Computing* (2023).

<sup>7</sup> T. Hogg and D. Portnov, *Information Sciences* **128**, 181 (2000).

<sup>8</sup> E. Farhi, J. Goldstone, and S. Gutmann, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).

<sup>9</sup> L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, *Physical Review X* **10**, 021067 (2020).

<sup>10</sup> J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, *Proceedings of the Conference on the Theory of Quantum Computation, Communication and Cryptography* **7**, 1 (2022).

<sup>11</sup> S. Boulebnane and A. Montanaro, [arXiv:2208.06909](https://arxiv.org/abs/2208.06909) (2022).

<sup>12</sup> S. H. Sureshababu, D. Herman, R. Shaydulin, J. Basso, S. Chakrabarti, Y. Sun, and M. Pistoia, [arXiv:2305.15201](https://arxiv.org/abs/2305.15201) (2023).

<sup>13</sup> A. Boehmer, *IEEE Transactions on Information Theory* **13**, 156 (1967).

<sup>14</sup> M. Schroeder, *IEEE Transactions on Information Theory* **16**, 85 (1970).

<sup>15</sup> J. Bernasconi, *Journal de Physique* **48**, 559 (1987).

<sup>16</sup> S. Mertens and C. Bessenrodt, *Journal of Physics A: Mathematical and General* **31**, 3731 (1998).

<sup>17</sup> M. Golay, *IEEE Transactions on Information Theory* **23**, 43 (1977).

<sup>18</sup> B. Bošković, F. Brglez, and J. Brest, “A GitHub Archive for Solvers and Solutions of the LABS problem,” (2016).

<sup>19</sup> T. Packebusch and S. Mertens, *Journal of Physics A: Mathematical and Theoretical* **49**, 165001 (2016).

<sup>20</sup> D. Lykov, R. Shaydulin, Y. Sun, Y. Alexeev, and M. Pistoia, “Fast simulation of high-depth QAOA,” In preparation.

<sup>21</sup> J. E. Gallardo, C. Cotta, and A. J. Fernández, *Applied Soft Computing* **9**, 1252 (2009).

<sup>22</sup> J. Lemieux, B. Heim, D. Poulin, K. Svore, and M. Troyer, *Quantum* **4**, 287 (2020).

<sup>23</sup> S. Boixo, G. Ortiz, and R. Somma, *The European Physical Journal Special Topics* **224**, 35 (2015).

- <sup>24</sup> 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, *Nature* **592**, 209 (2021).
- <sup>25</sup> S. A. Moses *et al.*, [arXiv:2305.03828](https://arxiv.org/abs/2305.03828) (2023).
- <sup>26</sup> A. Gonzales, R. Shaydulin, Z. H. Saleem, and M. Suchara, *Scientific Reports* **13** (2023).
- <sup>27</sup> D. M. Debroy and K. R. Brown, *Physical Review A* **102** (2020).
- <sup>28</sup> B. Bošković, F. Brglez, and J. Brest, *Applied Soft Computing* **56**, 262 (2017).
- <sup>29</sup> Supplementary Information: Evidence of Scaling Advantage for the Quantum Approximate Optimization Algorithm on a Classically Intractable Problem.
- <sup>30</sup> Y. R. Sanders, D. W. Berry, P. C. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, *PRX Quantum* **1** (2020).
- <sup>31</sup> J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, *Quantum* **4**, 230 (2020).
- <sup>32</sup> M. Boyer, G. Brassard, P. Høyer, and A. Tapp, *Fortschritte der Physik* **46**, 493 (1998).
- <sup>33</sup> R. Shaydulin and M. Pistoia, [arXiv:2303.02064](https://arxiv.org/abs/2303.02064) (2023).
- <sup>34</sup> Z. He, R. Shaydulin, S. Chakrabarti, D. Herman, C. Li, Y. Sun, and M. Pistoia, [arXiv:2305.03857](https://arxiv.org/abs/2305.03857) (2023).
- <sup>35</sup> E. Pelofske, A. Bärtschi, J. Golden, and S. Eidenbenz, [arXiv:2306.03238](https://arxiv.org/abs/2306.03238) (2023).
- <sup>36</sup> E. Pelofske, A. Bärtschi, and S. Eidenbenz, in *Lecture Notes in Computer Science* (Springer Nature Switzerland, 2023) pp. 240–258.
- <sup>37</sup> P. Niroula, R. Shaydulin, R. Yalovetzky, P. Minssen, D. Herman, S. Hu, and M. Pistoia, *Scientific Reports* **12** (2022).
- <sup>38</sup> R. Shaydulin and A. Galda, in *IEEE International Conference on Quantum Computing and Engineering* (2021).
- <sup>39</sup> A. Kakkar, J. Larson, A. Galda, and R. Shaydulin, in *IEEE International Conference on Quantum Computing and Engineering* (2022).
- <sup>40</sup> R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, *PRX Quantum* **2** (2021).
- <sup>41</sup> G. Brassard, P. Høyer, M. Mosca, and A. Tapp, *Contemporary Mathematics*, 53 (2002).
- <sup>42</sup> C. N. Self, M. Benedetti, and D. Amaro, [arXiv:2211.06703](https://arxiv.org/abs/2211.06703) (2022).
- <sup>43</sup> Gurobi Optimization, [www.gurobi.com](http://www.gurobi.com).
- <sup>44</sup> IBM ILOG CPLEX, International Business Machines Corporation.
- <sup>45</sup> F. Glover and M. Laguna, *Tabu Search* (Kluwer Academic Publishers, Norwell, MA, USA, 1997).
- <sup>46</sup> Fino and Algazi, *IEEE Transactions on Computers* **C-25**, 1142 (1976).

## DISCLAIMER

This paper was prepared for informational purposes with contributions from the Global Technology Applied Research center of JPMorgan Chase & Co., Argonne National Laboratory and Quantinuum LLC. This paper is not a product of the Research Department of JPMorgan Chase & Co. or its affiliates. Neither JPMorgan Chase & Co. nor any of its affiliates makes any explicit or implied representation or warranty and none of them accept any liability in connection with this position paper, including, without limitation, with respect to the completeness, accuracy, or reliability of the information contained herein and the potential Legal, compliance, tax, or accounting effects thereof. This document is not intended as investment research or investment advice, or as a recommendation, offer, or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

The submitted manuscript includes contributions from UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <http://energy.gov/downloads/doe-public-access-plan>.

**Supplementary Information:**  
**Evidence of Scaling Advantage for the Quantum Approximate  
Optimization Algorithm on a Classically Intractable Problem**

Ruslan Shaydulin,<sup>1,\*</sup> Changhao Li,<sup>1</sup> Shouvanik Chakrabarti,<sup>1</sup> Matthew DeCross,<sup>2</sup> Dylan Herman,<sup>1</sup>  
Niraj Kumar,<sup>1</sup> Jeffrey Larson,<sup>3</sup> Danylo Lykov,<sup>1,4</sup> Pierre Minssen,<sup>1</sup> Yue Sun,<sup>1</sup> Yuri Alexeev,<sup>4</sup> Joan M. Dreiling,<sup>2</sup>  
John P. Gaebler,<sup>2</sup> Thomas M. Gatterman,<sup>2</sup> Justin A. Gerber,<sup>2</sup> Kevin Gilmore,<sup>2</sup> Dan Gresh,<sup>2</sup> Nathan Hewitt,<sup>2</sup>  
Chandler V. Horst,<sup>2</sup> Shaohan Hu,<sup>1</sup> Jacob Johansen,<sup>2</sup> Mitchell Matheny,<sup>2</sup> Tanner Mengle,<sup>2</sup> Michael Mills,<sup>2</sup>  
Steven A. Moses,<sup>2</sup> Brian Neyenhuis,<sup>2</sup> Peter Siegfried,<sup>2</sup> Romina Yalovetzky,<sup>1</sup> and Marco Pistoia<sup>1</sup>

<sup>1</sup>*Global Technology Applied Research, JPMorgan Chase, New York, NY 10017, USA*

<sup>2</sup>*Quantinuum, Broomfield, CO 80021, USA*

<sup>3</sup>*Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA*

<sup>4</sup>*Computational Science Division, Argonne National Laboratory, Lemont, IL 60439, USA*

**CONTENTS**

SI. Background on the LABS problem	2
SII. QAOA as an exact and approximate optimization algorithm	2
SIII. Details of numerical studies	5
A. Optimized QAOA parameters for LABS change with $N$	5
B. QAOA parameter optimization with FOURIER scheme	5
C. Evidence of the success of the FOURIER reparameterization heuristic	6
D. Procedure for obtaining the fixed parameters	6
E. Evidence of the success of the fixed parameter scheme	7
F. Scaling coefficient of QAOA TTS is not sensitive to the choice of $N_{\min}$	7
G. Scaling coefficient of QAOA TTS does not follow power law	8
H. Comparison of performance between QAOA and amplitude amplification	8
I. Proof of Theorem 1	10
J. Details of the classical solver scaling	12
SIV. Experiments on trapped-ion systems	13
A. Experimental system	13
B. Circuit compilation and optimization	14
C. Summary of the error-detection scheme	16
D. Performance of the error-detection scheme	17
References	18

## SI. BACKGROUND ON THE LABS PROBLEM

The problem of finding sequences with low sidelobe energies attracted attention in the 1960s and 1970s due to its applications to the reduction of the peak power of radar pulses.<sup>1,2</sup> The merit factor  $F$  was first introduced by Golay<sup>3</sup> and defined as the ratio of central to sidelobe energy of a binary sequence. Improved merit factors were obtained over the years by exhaustive<sup>4,5</sup> and non-exhaustive<sup>6</sup> search methods. Explicit sequences asymptotically achieving the merit factor of  $\mathcal{F} \approx 6.34$  are known.<sup>7</sup> The conjectured asymptotic limit of  $\mathcal{F} \rightarrow 12.32$  as  $N \rightarrow \infty$  was derived using arguments from statistical mechanics in Ref. 8. Bernasconi reframed the LABS problem as a spin model with long-range 4-spin interactions to apply simulated annealing to it.<sup>9</sup> However, simulated annealing failed to obtain high-quality solutions, with the failure attributed to the “golf-course type” energy landscape.<sup>9</sup> Bernasconi further conjectured that this property of the landscape will prevent stochastic search procedures from obtaining high-quality solutions for long sequences.<sup>9</sup> This conjecture has held up so far.

A commonly considered class of sequences are those exhibiting skew-symmetry, which for sequences of odd length  $N = 2k - 1$  are defined as  $s_{k+l} = (-1)^l s_{k-l}$ ,  $l \in \{1, \dots, k-1\}$ . Skew-symmetric sequences are known to be optimal for many odd  $N$  instances. Since skew-symmetry reduces the search space from  $2^N$  to  $2^{\frac{N}{2}}$ , only searching this subspace leads to better runtime scaling. Therefore many algorithms are restricted to only searching this subspace. The best-known heuristic for skew-symmetric LABS uses a sequence of self-avoiding walk segments when searching and has a running time that scales as  $1.15^N$ .<sup>10</sup> In this work, we target the general LABS problem and therefore we do not consider solvers that are only capable of tackling the skew-symmetric instances.

## SII. QAOA AS AN EXACT AND APPROXIMATE OPTIMIZATION ALGORITHM

We now provide additional numerical results highlighting the differences in QAOA behavior with parameters optimized for approximate and exact solutions. In this work, we use QAOA as an exact solver, with time to solution as the target metric. However, QAOA is typically used as an approximation algorithm,<sup>11</sup> with most theoretical results focusing on the expected solution quality obtained by QAOA. On the LABS problem, we observe that QAOA can provide poor approximations in polynomial time while still offering speedups as an exact exponential-time solver.

We begin by discussing the parameters themselves. QAOA parameters are typically chosen with respect to QAOA energy<sup>11–18</sup> or the probability of sampling the optimal solution.<sup>19</sup> Figure S1a shows the difference in QAOA parameters in the two scenarios for the specific case of the LABS problem. First, we note that QAOA parameters that maximize  $\langle C \rangle_{\text{MF}}$  are significantly different from those maximizing  $p^{\text{opt}}$ . Specifically, while the values of  $\gamma$  are similar, the value of  $\beta$  for  $p^{\text{opt}}$  is much larger than that for  $\langle C \rangle_{\text{MF}}$ . Qualitatively, this implies that the probability amplitudes are allowed to “mix” more, making QAOA state not concentrated with respect to Hamming distance. This behavior is seen when examining the QAOA output distribution, which is shown for both parameter schedules in Fig. S2 and discussed in detail below.

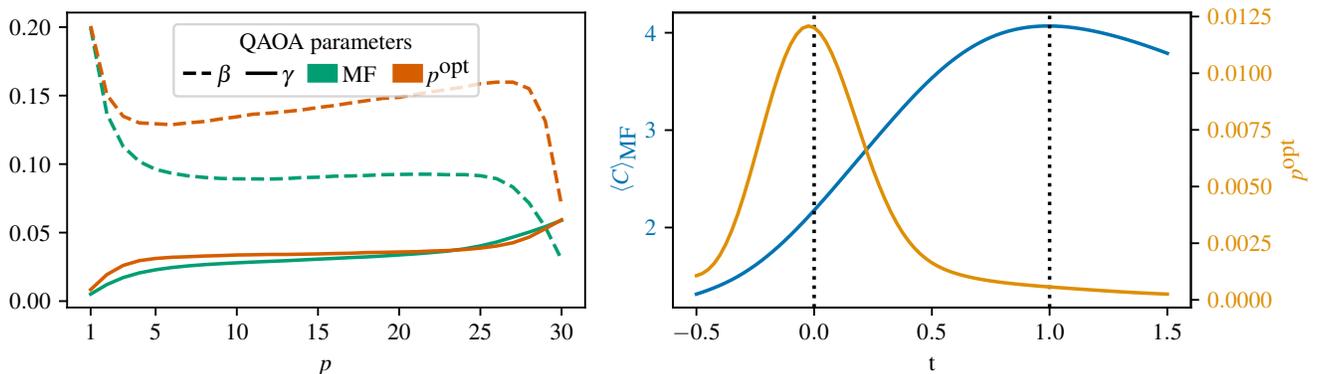


FIG. S1. **a**, Fixed QAOA parameters for  $p = 30$  chosen with respect to the QAOA energy  $\langle C \rangle_{\text{MF}}$  (“MF”) and probability of sampling the optimal solution (“ $p^{\text{opt}}$ ”). When the parameters are optimized with respect to  $p^{\text{opt}}$ , the value of  $\beta$  is significantly larger throughout QAOA evolution. Subfigure reproduced from the main text. **b**, QAOA performance for  $N = 25, p = 30$  with parameters linearly extrapolated between fixed parameters for  $p^{\text{opt}}$  ( $t = 0$ ) and  $\langle C \rangle_{\text{MF}}$  ( $t = 1$ ). QAOA parameters optimized for  $\langle C \rangle_{\text{MF}}$  give very poor values of  $p^{\text{opt}}$  and vice versa.

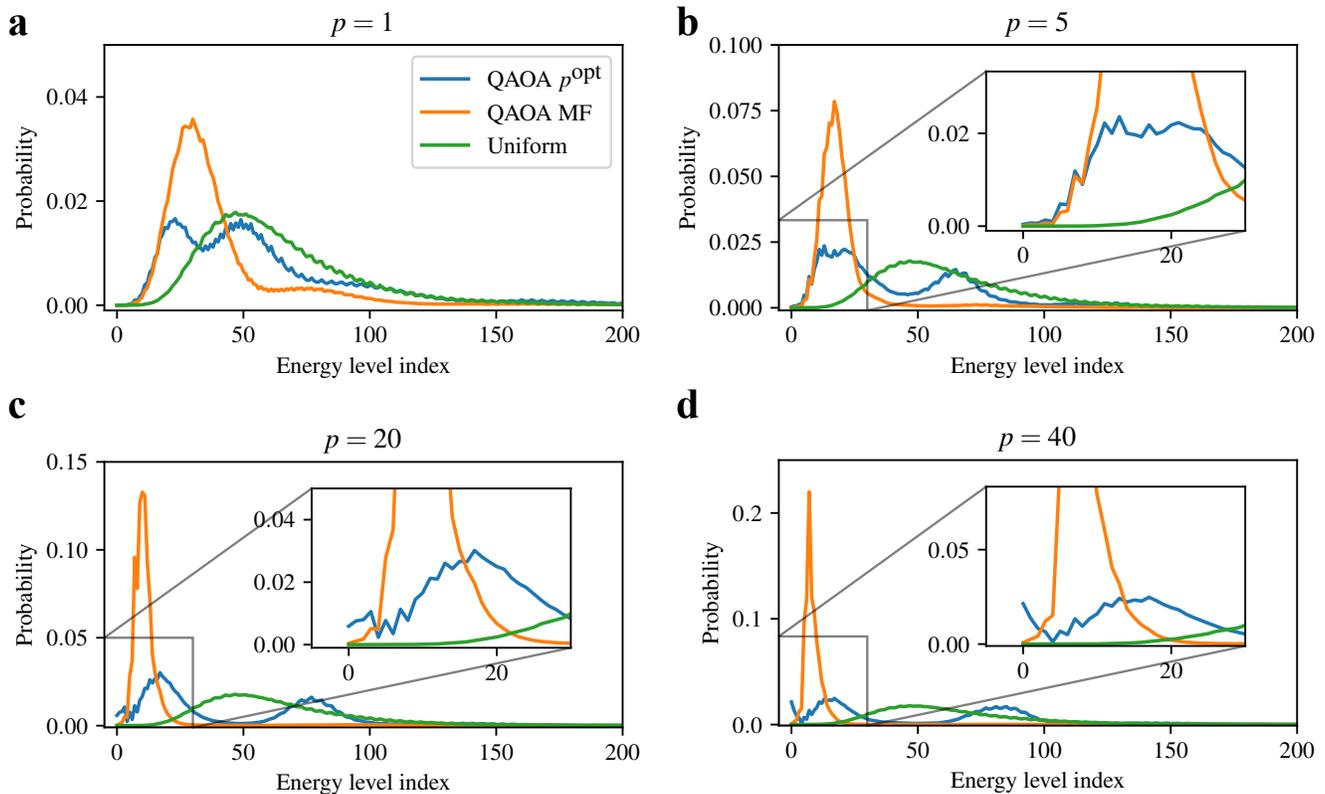


FIG. S2. The probability of obtaining a binary string corresponding to a given energy level of the LABS problem (the zeroth energy level is the ground state or optimal solution, lower is better) for varying  $p$  (a-d). When parameters are optimized with respect to the expected merit factor (labeled “QAOA MF”), the QAOA output state is concentrated around the mean and fails to obtain a high overlap with the ground state. On the other hand, when parameters are optimized with respect to  $p^{\text{opt}}$  (labeled “QAOA  $p^{\text{opt}}$ ”), the QAOA state has a high overlap with both ground state and higher energy states. The probability of obtaining the ground state is 27.3 times greater for QAOA with parameters optimized with respect to  $p^{\text{opt}}$  at  $p = 40$  (d).

Second, QAOA parameters that give good performance with respect to one metric are far from optimal with respect to the other metric. Fig. S1b shows QAOA performance with parameters linearly extrapolated between the parameters  $(\beta^{(C)}_{\text{MF}}, \gamma^{(C)}_{\text{MF}})$  that give a high value of  $\langle C \rangle_{\text{MF}}$  and  $(\beta^{p^{\text{opt}}}, \gamma^{p^{\text{opt}}})$  that give high  $p^{\text{opt}}$ :  $\gamma = t\gamma^{p^{\text{opt}}} + (1-t)\gamma^{(C)}_{\text{MF}}$  and  $\beta = t\beta^{p^{\text{opt}}} + (1-t)\beta^{(C)}_{\text{MF}}$ . We note that the parameters  $\beta^{(C)}_{\text{MF}}, \gamma^{(C)}_{\text{MF}}$  ( $t = 1$ ) give a very low value of  $p^{\text{opt}}$  and vice versa. This suggests that significant performance gains are possible if parameters are chosen differently with respect to the two figures of merit, rather than using one as a proxy for the other as is commonly done in QAOA research.<sup>20,21</sup> Similar observations have been made in Refs. 22 and 23, though the difference observed between the two figures of merit is more drastic in our case due to the hardness of the problem considered.

In the numerical experiments presented in the main test, we use time to solution as the target metric. For completeness, we include the results showing QAOA performance as an approximation algorithm. We observe that QAOA performs poorly on the LABS problem with respect to the expected merit factor  $\langle C \rangle_{\text{MF}}$ . Specifically, we observe that QAOA fails to outperform even simple classical techniques at high depth. Fig. S3 shows the expected merit factor of QAOA for fixed  $p = 100$  as a function of  $N$ , as well as examples of how the expected merit factor grows with  $p$  for two fixed values of  $N$ . We observe that as  $N$  grows, QAOA at  $p = 100$  achieves an expected merit factor  $\langle C \rangle_{\text{MF}}$  of less than 5. Note that explicit analytical sequences achieving merit factor  $> 6$  are known.<sup>7</sup> Moreover, we see that  $\langle C \rangle_{\text{MF}}$  grows increasingly slowly as  $N$  and  $p$  increase, suggesting that a prohibitively high value of  $p$  would be required to achieve a high expected merit factor.

We can understand this behavior by examining the values of the merit factor attainable by binary strings (in physics terms, we are examining the spectrum of the LABS Hamiltonian). For the  $N = 25$  problem, presented in Fig. S2, only the 4 lowest energy levels correspond to a merit factor better than 6. This means that QAOA must concentrate all the wave function mass on a superposition of a small number of computational basis states. Due to QAOA preserving the  $D_4$  symmetry of the problem,<sup>24</sup> this is a highly entangled state, which is hard to prepare.<sup>25</sup>

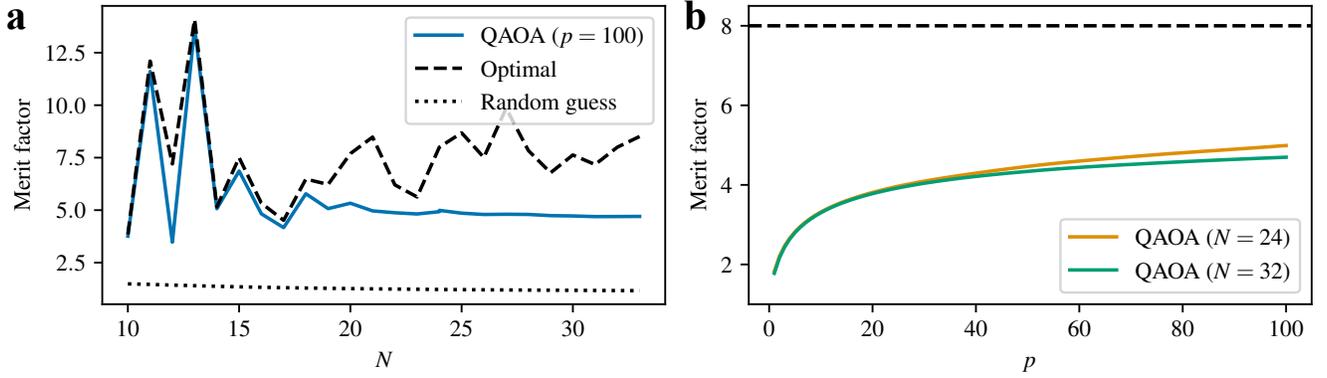


FIG. S3. **a**, Performance of QAOA as an approximation algorithm. Explicit constructions of skew-symmetric sequences exist that achieve  $\mathcal{F} \approx 6.34$  for large  $N$ .<sup>7</sup> Simulated annealing achieves  $\mathcal{F} \approx 5$  for large  $N$ .<sup>9</sup> For QAOA, the expected value of merit factor  $\langle C \rangle_{\text{MF}}$  is plotted. The expected merit factor of QAOA output is below the values easily attainable classically. **b**, For both  $N = 24$  and  $N = 32$ , the optimal merit factor is 8 (dashed line).

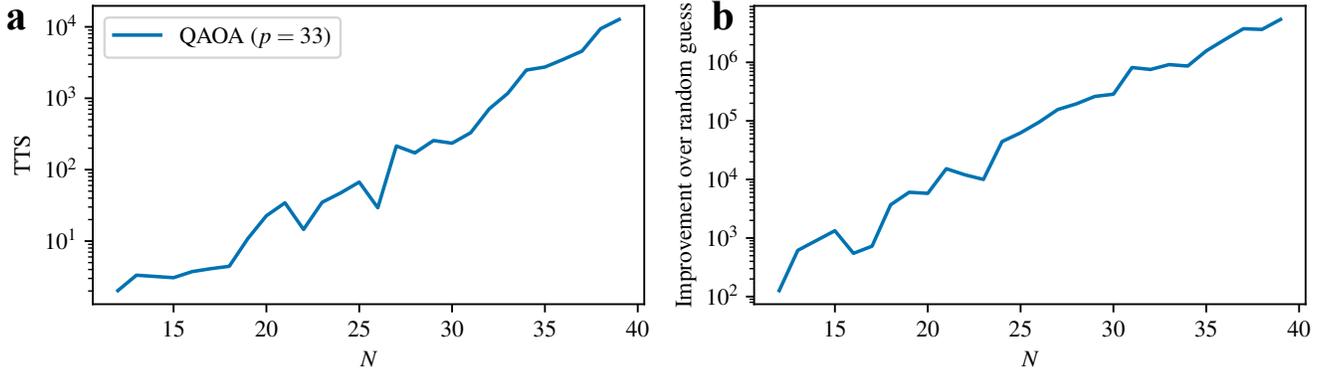


FIG. S4. **a**, Time-to-solution (TTS) of QAOA with parameters chosen with respect to  $p^{\text{opt}}$ . **b**, Improvement over random guess. For the largest case numerically considered, which is the  $N = 39$  problem at  $p = 33$ , the expected number of shots required to solve it is  $1.2 \times 10^4$ . This is a  $5.4 \times 10^6$  factor improvement over random guess.

As a result, QAOA requires a very large value of  $p$  to obtain a high expected merit factor. If, on the other hand, we choose the probability of sampling the exact optimal solution  $p^{\text{opt}}$  (overlap between QAOA state and the ground state of the LABS Hamiltonian) as the target metric, a much lower value of  $p$  is needed to obtain good success probability. Qualitatively, Fig. S2 shows how this state preparation succeeds by allowing a significant part of the QAOA state to “leak” to high energy levels. This can be observed by noticing how the population of energy levels  $> 50$  stays relatively high for QAOA with parameters optimized with respect to  $p^{\text{opt}}$ , but becomes negligible if QAOA parameters are chosen with respect to  $\langle C \rangle_{\text{MF}}$ .

The success of QAOA in preparing states with a large overlap with the ground state of the LABS Hamiltonian (i.e. states having a high probability of measuring exactly optimal solution) motivates the choice of time to solution as the target metric for QAOA evaluation. We show the time to solution at the largest  $p$  explored numerically ( $p = 33$ ) in Figure S4. We remark that the QAOA succeeds at achieving high overlap at  $N \leq 40$ .

Our results suggest a new way of viewing the potential of QAOA to provide algorithmic speedups and provide an important caveat to theoretical results bounding the approximations attainable by QAOA with constant depth.<sup>26</sup> Even in the regime where the expected solution quality of QAOA is bounded, it can be still useful as a tool to obtain a high probability of measuring the exact optimal solution.

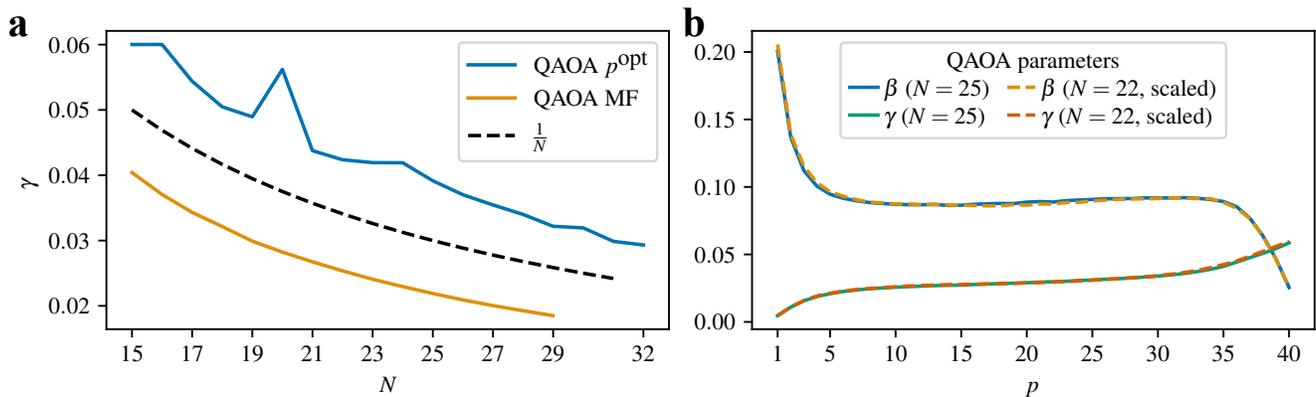


FIG. S5. **a**, Scaling of the optimized value of QAOA parameter  $\gamma$  with  $N$  for  $p = 1$  when optimized with respect to expected merit factor  $\langle C \rangle_{\text{MF}}$  (“QAOA MF”) and probability of obtaining optimal solution  $p^{\text{opt}}$  (“QAOA  $p^{\text{opt}}$ ”). Optimized  $\gamma^*$  decreases with  $N$  as  $\frac{1}{N}$ . **b**, QAOA parameters optimized with respect to  $\langle C \rangle_{\text{MF}}$  for  $N \in \{22, 25\}$ . For  $N = 22$ , the parameters  $\gamma$  are scaled by  $22/25$ . After rescaling, the parameters for  $N = 22$  and  $N = 25$  are visually indistinguishable.

### SIII. DETAILS OF NUMERICAL STUDIES

We now present in detail how the fixed parameters were chosen. Our procedure for doing so is as follows. First, we optimize the QAOA parameters using the FOURIER<sup>12</sup> reparameterization. We show evidence that the optimization of the reparameterized QAOA gives the same performance as the more extensive optimization of the standard parameterization. Second, we set our fixed parameters to be the arithmetic mean of the (appropriately rescaled) optimized parameters for smaller  $N$ . We provide evidence that for smaller  $N$  where directly optimized parameters are available, the fixed parameters lead to QAOA performance that is close to that with the optimized parameters. We note that better parameter setting schemes can only improve our performance.

#### A. Optimized QAOA parameters for LABS change with $N$

First, we observe that the optimized QAOA parameters are not invariant with  $N$ . Specifically, we observe that the optimized value of  $\gamma^*$  goes down with  $N$  as  $\frac{1}{N}$ . S5a plots this for  $p = 1$ . We note that  $\beta^*$  is roughly constant with  $N$ . We observe this scaling for all  $p$ . As an example, S5b plots optimized parameters for two different values of  $N$ . After rescaling  $\gamma^*$  by  $N$ , the two sets of parameters are visually indistinguishable. Below, we take advantage of this scaling of  $\gamma^*$  in two ways. First, we improve the convergence of local optimization runs by rescaling the initialization and the initial step size of the local optimizer. Second, we use it to correctly account for scale when executing QAOA with fixed parameters.

The scaling of  $\gamma^*$  with extremal properties of the objective function has been observed before for other problems. For example, the normalized value of the maximum cut on  $D$ -regular graphs grows with  $D$  (to the first order) and  $\gamma^*$  decreases as  $\frac{1}{\sqrt{D}}$ .<sup>27,28</sup> Analogous scaling has been observed for weighted problems.<sup>29</sup> In LABS, the energy  $\mathcal{E}_{\text{sidelobe}}(\mathbf{s})$  grows as  $N^2$ , so  $\gamma^*$  decreases as  $\frac{1}{N}$ . Formally establishing this connection is a promising direction for future research.

#### B. QAOA parameter optimization with FOURIER scheme

For a given figure of merit, we optimize the QAOA parameters as follows. In all cases below, we use the nlopt<sup>30</sup> implementation of BOBYQA<sup>31</sup> gradient-free local optimization algorithm. In all cases, we run BOBYQA until convergence, with convergence specified by relative tolerances on changes in parameters and in objective function value of  $10^{-8}$ . BOBYQA has been shown to outperform other local optimizers on the task of optimizing QAOA parameters.<sup>13</sup> We expect similar results with any other local hill-climbing algorithm, albeit at a potentially different cost in terms of the number of iterations.

For  $p = 1$ , we optimize the QAOA parameter exhaustively by running the local optimizer from 400 initial points. We set the initial step size (rhobeg) to  $0.01/N$ . The initial points  $\beta^{\text{init}}, \gamma^{\text{init}}$  are chosen uniformly at random from

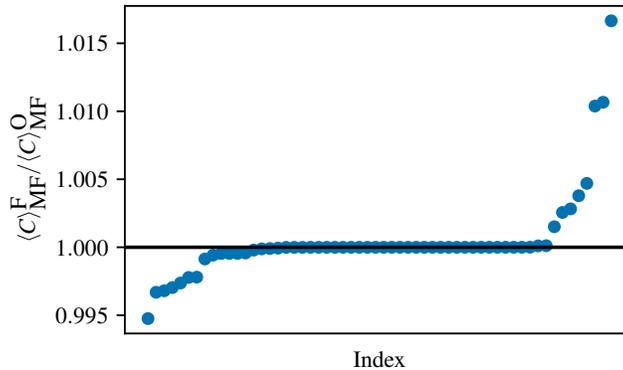


FIG. S6. The ratio between the expected merit factor of QAOA with parameters optimized by directly running local optimization from many initial points ( $\langle C \rangle_{\text{MF}}^{\text{O}}$ ) and with parameters optimized using the FOURIER $[\infty, 0]$  scheme ( $\langle C \rangle_{\text{MF}}^{\text{F}}$ ) for  $N > 12$ . We observe that the difference in the quality of the parameters obtained by the two optimization schemes is small.

$\beta^{\text{init}} \in [0.1, 0.2]$ ,  $\gamma^{\text{init}} \in [0, 0.85/N]$  when optimizing with respect to  $\langle C \rangle_{\text{MF}}$ , and  $\beta^{\text{init}} \in [0.15, 0.3]$ ,  $\gamma^{\text{init}} \in [0.6, 1.2/N]$  when optimizing with respect to  $p^{\text{opt}}$ . The regions for initializations are read off from grid search results for  $p = 1$ .

For  $p > 1$ , we follow the FOURIER $[\infty, 0]$  scheme of Ref. 12. Specifically, we change the QAOA parameterization to the frequency domain as follows:

$$\gamma_i = \sum_{k=1}^p u_k \sin \left[ \left( k - \frac{1}{2} \right) \left( i - \frac{1}{2} \right) \frac{\pi}{p} \right], \quad (\text{S1})$$

$$\beta_i = \sum_{k=1}^p v_k \cos \left[ \left( k - \frac{1}{2} \right) \left( i - \frac{1}{2} \right) \frac{\pi}{p} \right]. \quad (\text{S2})$$

Then we optimize over the new parameters  $\mathbf{u}, \mathbf{v}$ . We take the optimized parameters  $\mathbf{u}_{p-1}^*, \mathbf{v}_{p-1}^*$  for  $p - 1$  and run one local optimization from  $\mathbf{u}_p^* = (\mathbf{u}_{p-1}^*, 0)$ ,  $\mathbf{v}_p^* = (\mathbf{v}_{p-1}^*, 0)$ . The initial step size (`rhobeg`) for local optimizer is set to  $0.01/N$ . An initial step size that is small and decreasing with  $N$  is central to the robust convergence of a local optimizer, due to QAOA parameters having different scales for different  $N$ . We do not find it necessary for obtaining high-quality parameters to perform objective function rescaling of the type explored in Refs. 29 and 32, though we expect that it may reduce the number of iterations required by the local optimizer to converge.

### C. Evidence of the success of the FOURIER reparameterization heuristic

To evaluate the success of the FOURIER parameter optimization heuristic, we compare the quality of optimized parameters it finds with the quality of the parameters obtained by running a local optimizer with  $100p$  seeds from initial points sampled uniformly from  $\beta^{\text{init}} \in [0.1, 0.2]^p$ ,  $\gamma^{\text{init}} \in [0, 0.85/N]^p$ .

We find that the very expensive direct optimization performs very similarly to one local optimization run used in the FOURIER scheme, as shown in Figure S6. Specifically, the mean difference between the two schemes is  $< 0.05\%$ , and in the worst case of the ones considered, FOURIER gives parameters that are only  $< 0.5\%$  worse. Therefore below, we simply consider parameters optimized using the FOURIER $[\infty, 0]$  scheme.

### D. Procedure for obtaining the fixed parameters

The procedure we follow for obtaining the fixed parameters is described in the main text. We reiterate it here for completeness. We optimize QAOA parameters for smaller values of  $N$  for which the simulation is relatively inexpensive

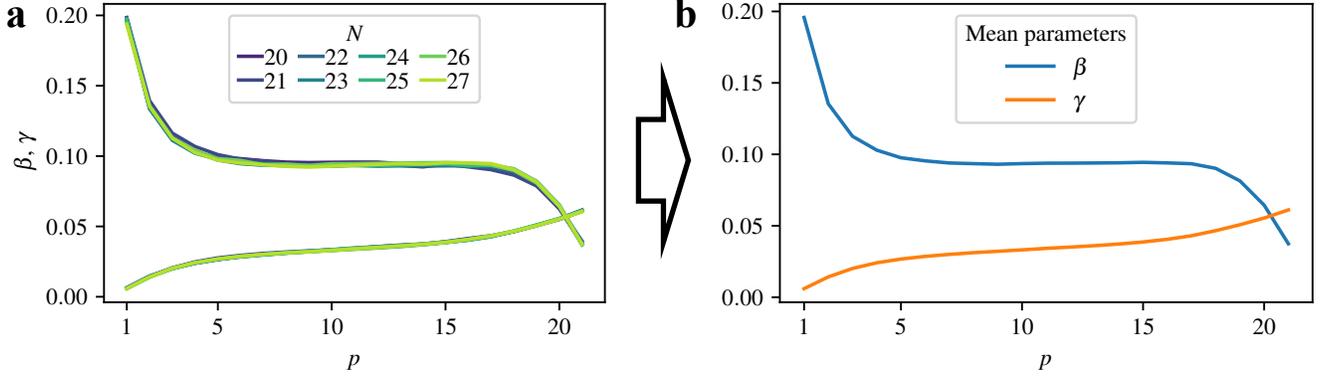


FIG. S7. Visualization of how the fixed parameters with respect to  $\langle C \rangle_{\text{MF}}$  are obtained. Visualization for  $p^{\text{opt}}$  is presented in the main text. **a**, Optimized QAOA parameters  $\beta$  (top line) and  $\gamma$  (bottom line) for  $p = 21$ .  $\gamma$  is scaled by  $N/24$ , with the constant factor of  $1/24$  added for figure readability. **b**, Fixed parameters obtained by taking the arithmetic mean over the optimized parameters.

and set the fixed parameters to be the mean over the optimized parameters:

$$\beta^{\text{Fixed}} = \frac{1}{M} \sum_{N \in \{N_1, \dots, N_M\}} \beta_N^*, \quad (\text{S3})$$

$$\gamma^{\text{Fixed}} = \frac{1}{M} \sum_{N \in \{N_1, \dots, N_M\}} N \gamma_N^*, \quad (\text{S4})$$

where  $\beta_N^*$ ,  $\gamma_N^*$  are the QAOA parameters optimized for  $N$ . The fixed parameters used in QAOA applied to a LABS instance of size  $N$  are then given by  $\beta^{\text{Fixed}}$ ,  $\frac{\gamma^{\text{Fixed}}}{N}$ . This process is visualized for parameters optimized with respect to  $\langle C \rangle_{\text{MF}}$  in Figure S7. We use optimized parameters for  $20 \leq N \leq 27$  when computing parameters for  $\langle C \rangle_{\text{MF}}$  and  $24 \leq N \leq 31$  for  $p^{\text{opt}}$ .

### E. Evidence of the success of the fixed parameter scheme

To evaluate the quality of fixed parameters, we compare the QAOA performance with fixed parameters and with directly optimized parameters. Figure S8 presents the comparison. We observe that the two are very close for small  $p$ , with the ratio between the two growing for higher  $p$ . Specifically, for parameters optimized with respect to the expected merit factor  $\langle C \rangle_{\text{MF}}$ , the median difference in  $\langle C \rangle_{\text{MF}}$  between QAOA evaluated with fixed and directly optimized parameters is less than 0.01% for  $p = 50$ , with the difference even lower for smaller  $p$  (S8a). For parameters optimized with respect to  $p^{\text{opt}}$ , the difference in  $p^{\text{opt}}$  is larger and growing with  $p$  (S8c). Note that due to the exponentially small value of  $p^{\text{opt}}$ , small absolute differences (including due to precision limitations) translate into large relative differences. Nonetheless, we observe good performance with fixed parameters at high  $N$ , as visualized for  $N = 32$  in Figure S8d. QAOA performance with fixed parameters with respect to both figures of merit monotonically improves with  $p$  for all values of  $p$  considered. As the performance gap between fixed and optimized parameters grows with  $p$ , further improvements to fixed parameters are likely to yield even better scaling of QAOA TTS.

### F. Scaling coefficient of QAOA TTS is not sensitive to the choice of $N_{\text{min}}$

In the Main text we motivate the choice of the cutoff  $N_{\text{min}}$  for the range of  $N$ s to be included in the fit by examining the stability of the quality of fit with varying  $p$ . We obtain  $N_{\text{min}} = 28$  as the minimum value required to maintain a stable fit, and estimate QAOA scaling at  $1.46^N$  at  $p = 12$ . We now provide evidence that this value is not sensitive to the choice of  $N_{\text{min}}$ .

Fig. S9 shows the scaling exponent for QAOA TTS for varying choices of the cutoff  $N_{\text{min}}$ . Taking the average over exponents obtained when performing a fit with  $20 \leq N_{\text{min}} \leq 35$  gives the estimated scaling of  $1.45^N$  (Fig. S9a), which is slightly better than the one reported in the main text. As shown in Fig. S9b, for sufficiently large  $p$  and small  $N_{\text{min}}$

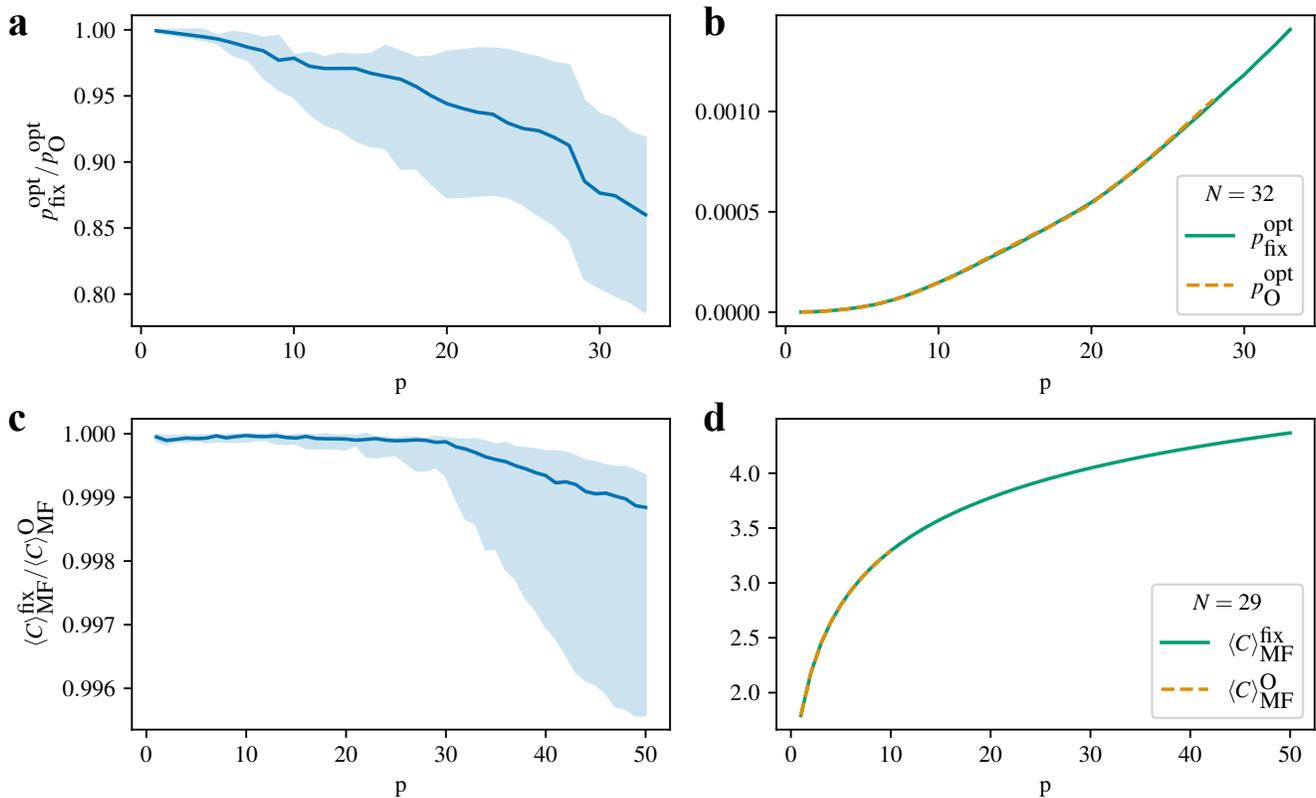


FIG. S8. Comparison of QAOA performance with fixed ( $\langle C \rangle_{MF}^{fix}, p_{fix}^{opt}$ ) and optimized ( $\langle C \rangle_{MF}^O, p_O^{opt}$ ) parameters. **a,c**, Shaded area shows 95% confidence interval. **b,d**, Despite the relative differences between performance with optimized and transferred parameters growing with  $p$ , we observe that for all cases considered, QAOA performance improves monotonically as expected. Since the absolute differences are small, the performance with fixed and directly optimized parameters is visually indistinguishable.

the exponent changes as  $N_{min}$  is increased, indicating that a larger regime of  $N$  must be explored to obtain a stable linear scaling.

### G. Scaling coefficient of QAOA TTS does not follow power law

One of the important findings of this work is that for LABS problem increasing QAOA depth  $p$  beyond a certain small constant does not lead to better scaling. This puts the findings of this work in contrast to those of Refs. 19, 33, and 34. We observe that unlike in Ref. 19, the scaling coefficient  $c$  in  $TTS = \Theta(2^{-cN})$  does not follow a power law. This is shown in Figure S10. For the coefficient  $c$  to follow a power law, it must depend on  $p$  as  $c_1 \times p^{c_2}$  for some constants  $c_1, c_2$ . When the cutoff is chosen to ensure good fit ( $N = 28$ ), we see clear deviation from a power law. We note that if we include smaller values of  $N$  in the fit and ignore the low quality of the fit, we observe a clear power law scaling of  $c \sim 0.77 \times p^{-0.18}$ .

### H. Comparison of performance between QAOA and amplitude amplification

In this work we propose a strategy for using QAOA as a building block for algorithmic speedups by combining it with amplitude amplification (AA) or, more specifically, generalized minimum-finding as described in Appendix C of Ref. 35. Specifically, we propose running the quantum minimum-finding algorithm with constant-depth QAOA as a subroutine. If the QAOA circuit prepares the state with  $p^{opt}$  overlap with the ground state of the LABS Hamiltonian, the minimum-finding algorithm would need apply the QAOA unitary and an oracle for computing the LABS cost function  $O(1/\sqrt{p^{opt}})$  times to obtain a constant probability of measuring a bitstring corresponding to the optimal

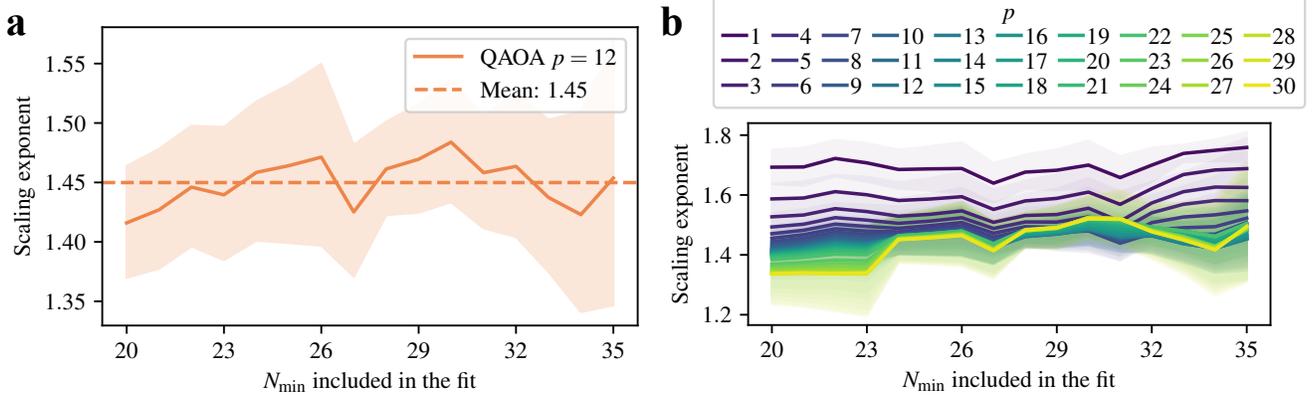


FIG. S9. The scaling exponent for QAOA TTS for varying choices of  $N_{\min}$  included in the range of values of  $N$  for fit at **a**,  $p = 12$  and **b**, for varying  $p$ . If smaller  $N_{\min}$  are included, the scaling exponent continues to improve with  $p$ , with the quality of fit decaying with  $p$  (see Main text and Fig. S10). For sufficiently high  $N_{\min}$ , the exponent does not improve beyond  $p \approx 10$ . At  $p = 12$ , the scaling exponent is not sensitive to the choice of  $N_{\min}$  (**a**).

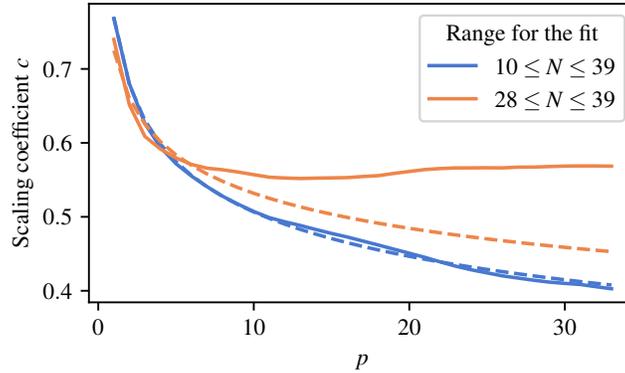


FIG. S10. The scaling coefficient  $c$  in  $TTS = \Theta(2^{-cN})$  as a function of  $p$ . The blue line is included for illustration as including  $N < 28$  leads to low quality of the fit. When the quality of the fit is high ( $28 \leq N \leq 38$ ), the scaling coefficient does not follow a power law.

solution.

We note that our numerical results suggest that increasing the QAOA depth is always beneficial as compared to doing a smaller number of QAOA steps and increasing the number of amplitude amplification steps. This is visualized in Figure S11. As this figure shows, for sufficiently high  $p$  the gain from a step of QAOA is very close to the gain from amplitude amplification, which provides an upper bound on the expected gain from one step of minimum-finding. At the same time one step of QAOA is much simpler to implement as it only requires the phase oracle and a series of one-qubit gates (mixer). We note that we do not provide exact gains from a step of the minimum-finding algorithm as they are non-trivial to estimate. However, asymptotically it is equivalent to amplitude amplification with an oracle that marks optimal solutions.

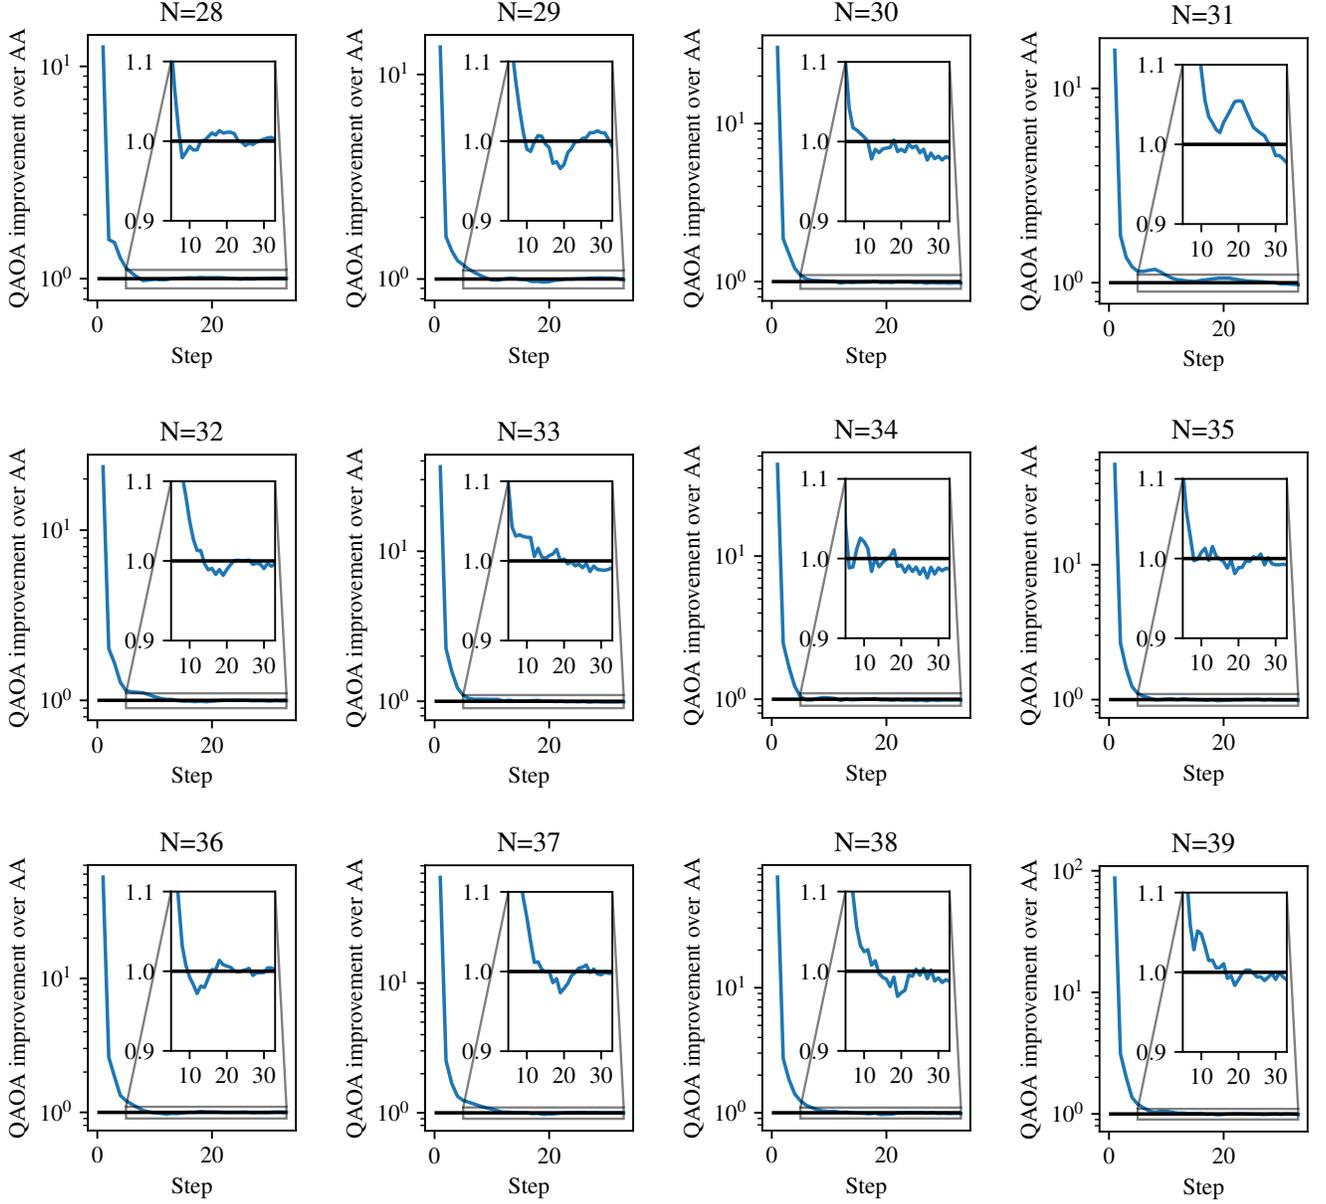


FIG. S11. The ratio between the gains in the probability of measuring the optimal solution from one step of QAOA and amplitude amplification (AA). QAOA provides orders of magnitude larger improvements for the first few steps, and then behaves approximately like amplitude amplification.

### I. Proof of Theorem 1

We now present the proof of Theorem 1 of the main text. Our technique is based on the generalized minimum-finding procedure outlined in Lemma 48 of Ref. 35. For reader's convenience and completeness, we begin by restating our algorithm and the Theorem.

**Lemma 1** (Exponential Quantum Search, Ref. 36). *Let  $|\psi\rangle = U|0\rangle^{\otimes N}$  be a quantum state in a  $2^N$ -dimensional Hilbert space with computational basis elements indexed by  $N$ -bit bitstrings, and  $m: \{0,1\}^N \rightarrow \{0,1\}$  be a marking function such that  $\sum_{\{x|m(x)=1\}} |\langle\psi|x\rangle|^2 \geq p$ . There exists a quantum algorithm **EQSearch**( $U, m, \delta$ ) that outputs an element  $x^*$  such that  $m(x^*) = 1$  with probability at least  $\delta$  using  $O\left(\frac{1}{\sqrt{p}} \log\left(\frac{1}{\delta}\right)\right)$  applications of  $U$  and  $m$ .*

---

**Algorithm 1** QAOA Enhanced with Quantum Minimum-Finding
 

---

**Require:** Unitary  $U_{\text{QAOA}}$  acting on  $\mathbb{C}^{2^N}$  such that  $|\langle x^* | U_{\text{QAOA}} | 0 \rangle^{\otimes N}| \geq 1/\sqrt{p_{\text{opt}}}$  for unknown  $p_{\text{opt}}$ ,  $V_{\text{LABS}}$  for computing  $\mathcal{E}_{\text{sidelobe}}$  into a register, and  $\delta \in (0, 1)$ , positive number  $M \leq 2^N$ ,  $C$  is the constant corresponding to the  $O(\cdot)$  in Lemma 1

**Ensure:** If  $M$  is greater than  $1/\sqrt{p_{\text{opt}}}$ , output  $x^*$  with  $\geq 1 - \delta$  probability using  $O(\log(1/\delta)M)$  calls to  $U_{\text{QAOA}}$  and  $V_{\text{LABS}}$  (and their inverses).  
 $x_{\text{res}}$  is set to an empty list.  
**for**  $i \leftarrow 1$  **to**  $\lceil \log(1/\delta) \rceil$  **do**  
    $t \leftarrow 0$ ;  $s_0 \leftarrow \infty$   
   **while** number of calls to  $U_{\text{QAOA}}$  &  $V_{\text{LABS}}$  is  $< 3CMN$  **do**  
      $t \leftarrow t + 1$   
     Define  $m_t: \{0, 1\}^N \rightarrow \{0, 1\}$  such that  $m_t(x) = 1$  if and only if  $\mathcal{E}_{\text{sidelobe}}(x) < s_{t-1}$ . Note that  $m_t$  can be coherently evaluated using one query to  $V_{\text{LABS}}$ .  
     Set  $s_t = \mathbf{EQSearch}(U_{\text{QAOA}}, m_t, 1/(6 \cdot 2^N))$ .  
   **end while**  
   Append  $s_t$  to  $x_{\text{res}}$ .  
**end for**  
 Output minimum of  $x_{\text{res}}$ .

---

**Theorem 1.** *Suppose a constant-depth QAOA circuit  $U_{\text{QAOA}}$  prepares a state  $|\psi\rangle = U_{\text{QAOA}}|0\rangle^{\otimes N}$  with  $N \geq 3$ , such that we have  $|\langle x^* | \psi \rangle| \geq 1/\sqrt{p_{\text{opt}}}$ , where  $|x^*\rangle$  encodes an optimal solution to the  $N$ -bit LABS problem in a computational basis state, and we assume that  $p_{\text{opt}} \geq 1/N$ . Then, running Algorithm 1 with parameters  $M \geq 1/\sqrt{p_{\text{opt}}}$  and failure probability  $\delta$ , runs with a gate complexity of  $O(\text{poly}(N) \log(1/\delta)M)$  and finds  $x^*$  with probability at least  $1 - \delta$ .*

*Proof.* We will first analyze the randomized Algorithm 1 as a *Las Vegas* algorithm, i.e. we assume that the internal **while** loop is infinite. In the upcoming analysis we will assume that every call to **EQSearch** in the internal loop behaves as intended. Note that we choose the failure probability of each such call to be  $1/(6 \cdot 2^N)$ . The total number of calls cannot be more than  $2^N$  and so by the union bound, every call succeeds except with probability at most  $1/6$ .

The algorithm generates a monotonically decreasing sequence of samples, uses **EQSearch** in each iteration to search for a sample that is strictly lesser than the previous one. Since the sequence of samples is strictly decreasing and there are  $N$  possible distinct samples, the algorithm eventually returns the minimum. Since correctness is eventually guaranteed, we can simply bound the expected number of iterations before the sequence finds the minimum. If this expected number is  $m$ , we can run the internal loop  $2m$  times to ensure that we find the minimum with probability at least  $1/2$ . Consequently,  $2m \log(\frac{1}{\delta})$  iterations suffice to ensure that we find the minimum with probability at least  $1 - \delta$ . It remains to show that the expected number of iterations before the minimum is found is at most  $O\left(\frac{1}{\sqrt{p_{\text{opt}}}}\right)$ .

For this argument, we define the following quantities.  $(x_1 = x^*, x_2, \dots, x_n)$  is the list of bit-strings sorted in ascending order of the value of  $\mathcal{E}_{\text{sidelobe}}$ , and define  $P(\xi(X))$  to be the probability of event  $\xi(X)$  when  $X$  is a bitstring sampled by measuring  $|\psi\rangle$  in the computational basis. Suppose in some iteration  $t$ , the sample returned by **EQSearch** is  $s_t$ . In the next iteration, **EQSearch** searches for an element with sidelobe energy less than  $s_t$ . The central observation is the following: for any  $x_k$  where  $k \in [N]$ , the probability that some  $s_t = x_k$  given that  $t$  is the first iteration where  $x_k$  appears in the list of obtained samples is given by  $P(X = x_k)/P(X \leq x_k)$ . To see this, we observe as in Ref. 35 that:

$$\begin{aligned}
\Pr(s_t = x_k | s_t \leq x_k \wedge s_{t-1} > x_k) &= \frac{\Pr(s_t = x_k)}{\Pr(s_t \leq x_k \wedge s_{t-1} > x_k)} \\
&= \sum_{x_l > x_k} \frac{\Pr(s_t = x_k \wedge s_{t-1} = x_{l-1})}{\Pr(s_t \leq x_k \wedge s_{t-1} > x_k)} \\
&= \sum_{x_l > x_k} \frac{\Pr(s_t = x_k \wedge s_{t-1} = x_{l-1})}{\Pr(s_t \leq x_k \wedge s_{t-1} > x_k)} \\
&= \sum_{x_l > x_k} \frac{\Pr(s_t = x_k \wedge s_{t-1} = x_{l-1}) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)}{\Pr(s_t \leq x_k \wedge s_{t-1} > x_k) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)} \\
&= \sum_{x_l > x_k} \frac{\Pr(s_t = x_k \wedge s_{t-1} = x_{l-1}) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)}{\Pr(s_t \leq x_k \wedge s_{t-1} > x_k) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)} \\
&= \sum_{x_l > x_k} \frac{\Pr(s_t = x_k | s_{t-1} = x_{l-1}) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)}{\Pr(s_t \leq x_k | s_{t-1} = x_k) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} > x_l)} \\
&= \sum_{x_l > x_k} \frac{P(X = x_k) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} = x_l)}{P(X \leq x_k) \Pr(s_{t-1} = x_l) \Pr(s_t \leq x_k \wedge s_{t-1} > x_l)} \\
&= \frac{P(X = x_k)}{P(X \leq x_k)} \tag{S5}
\end{aligned}$$

Notice that since a value can be sampled at most once, and the minimum is obtained within  $n$  steps, the probability that a given value  $x_k$  occurs in the list of observed samples is  $P(X = x_k)/P(X \leq x_k)$ .

To bound the expected number of queries before the minimum ( $x_1$ ) is found by the algorithm, we associate with each bitstring  $x_l \in \{x_n, x_{n-1}, \dots, x_2\}$  the probability that it is an obtained sample in some iteration, and the cost of performing the corresponding search for an element less than  $x_l$ . The cost of the search, for our chosen parameters is at most  $\frac{C \log(6 \cdot 2^N)}{\sqrt{P(X < x_l)}} \leq \frac{2CN}{\sqrt{P(X < x_l)}}$ , where the last inequality holds whenever  $N \geq 3$ .

The total expected number of queries before the minimum  $x_1$  is found, is therefore given by

$$\sum_{i=2}^N 2CN \cdot \frac{P(X = x_i)}{P(X \leq x_i)} \sqrt{\frac{1}{P(X < x_i)}} \leq \sum_{i=2}^N \frac{2CN}{P(X \leq x_i)} \sqrt{\frac{1}{P(X < x_i)}} \leq 2CN \int_{p_{\text{opt}}}^1 r^{-3/2} dr \leq \frac{CN}{\sqrt{p_{\text{opt}}}} \tag{S6}$$

If we run the inner loop of Algorithm 1 more than 3 times the expected number of queries required to find  $x^*$ , as prescribed if  $M \geq 1/\sqrt{p_{\text{opt}}}$ , we fail to find  $x^*$  with probability at most 1/3 by the Markov inequality as long as no query to **EQSearch** fails. Additionally, by the earlier discussion, a query in the internal loop fails with probability at most 1/6. Therefore, by a union bound, each application of the inner loop finds  $x^*$  with probability of at least 1/2. Repeating the inner loop  $\log(1/\delta)$  times ensures that  $x^*$  is found with probability at least  $1 - \delta$  (if not, the inner loop has to fail to find  $x^*$  in  $\log(1/\delta)$  independent trials).  $\square$

## J. Details of the classical solver scaling

The scaling of the commercial branch-and-bound solvers is presented in Figure S12. For each solver, we run it with 100 random seeds for  $N \leq 32$  and 10 random seeds for  $N > 32$  and report the mean. The minimum  $N$  to include in the fit was chosen to maximize the quality of fit. We set the Gurobi parameters as follows: `Cuts=0`, `Heuristics=0`. For the other parameters in Gurobi and CPLEX the defaults are used. We observe that the performance of the two commercial solvers considered is within a 95% confidence interval of each other, with the TTO scaling matching that reported in Ref. 37.

We report complete results for the Memetic Tabu solver in Figure S13. The scaling is obtained by extrapolating the number of cost function evaluations made by the Memetic Tabu algorithm at different lengths. This quantity is fixed over repeated seeds for a given seed and length, unlike the execution time that may fluctuate depending on the runtime environment. The fluctuation in running time is much larger for Memetic Tabu as compared to branch-and-bound due to lower absolute value of the runtime ( $< 1$  sec for  $N = 40$ ) The time to evaluate the cost function on a sequence of length  $N$  scales only as  $N^2$ , which does not produce a consistent effect on runtime scaling at small

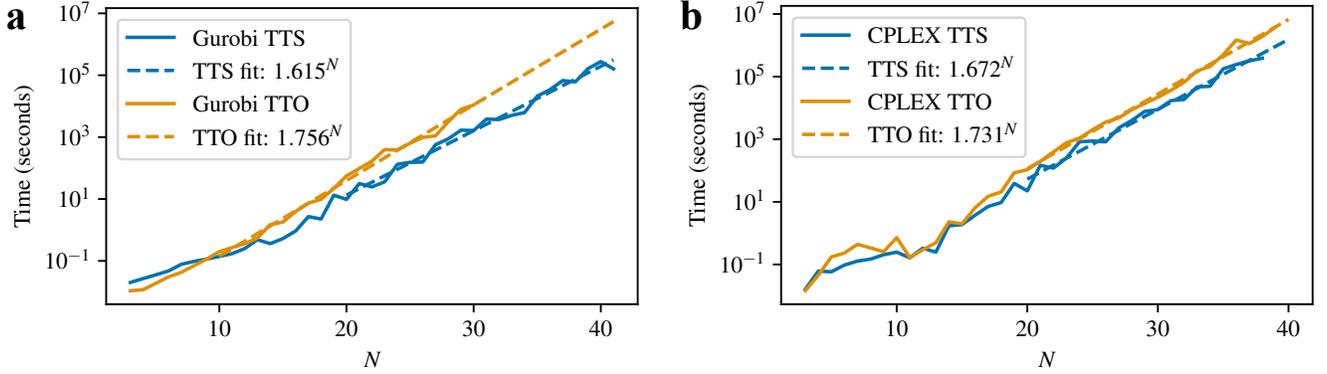


FIG. S12. Time-to-solution (TTS) and time to obtain a certificate of optimality (time-to-optimality or TTO) of **a**, Gurobi and **b**, CPLEX. For Gurobi, the 95% confidence interval (CI) for TTS is (1.571, 1.659) and for TTO is (1.721, 1.792). For CPLEX, the 95% CI for TTS is (1.609, 1.737) and for TTO is (1.693, 1.770).

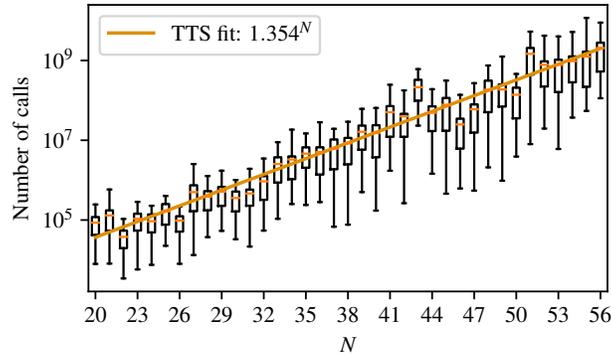


FIG. S13. Box plots showing the run time of Memetic Tabu. The scaling fit is reported with respect to the mean. The 95% confidence interval (CI) is (1.325, 1.383). The whiskers are minimum and maximum, box is showing quartiles and the horizontal line in the box is the mean.

lengths. The TTS scaling is therefore essentially the same as the scaling of the number of function evaluations,<sup>10</sup> and we choose to report the latter, more stable quantity. The seeds chosen for the runs are a contiguous block of 50 integers chosen from a random starting point. The Memetic Tabu solver is run in single-threaded mode for all our experiments, to avoid the overestimation of cost function evaluations arising from race conditions between exploration and termination checks.

## SIV. EXPERIMENTS ON TRAPPED-ION SYSTEMS

### A. Experimental system

The experiments in this work were performed on Quantinuum H1 and H2 platforms.<sup>38,39</sup> The system design is based on the QCCD architecture with multiple separate gate zones. Each gate zone is used to perform operations on an arbitrary pair of two qubits at a time, suppressing crosstalk and maintaining high fidelity. The hyperfine approximate clock states of  $^{171}\text{Yb}^+$  in the  $^2S_{1/2}$  state are used to encode qubit information. Namely,  $|0\rangle \equiv |F=0, m_f=0\rangle$  and  $|1\rangle \equiv |F=1, m_f=0\rangle$ . After loading, the qubits can be prepared in  $|0\rangle$  via optical pumping.<sup>38,39</sup>

The systems have all-to-all connectivity with two-qubit gates between pairs of qubits implemented by ion transport, which brings the pairs into the same gate zone. To implement two-qubit gates, a phase-sensitive Mølmer-Sørensen (MS) gate sandwiched between single-qubit wrapper pulses is used. It in turn gives the ZZ gate  $R_{ZZ}(\gamma) = \exp(-i\gamma ZZ/2)$ , where the rotation angle can be precisely controlled by changing the parameters in the MS gate.<sup>38,39</sup> Both the H1-1

and H2 systems used in this work have a typical average two-qubit infidelity of  $2 \times 10^{-3}$ , with single qubit infidelity two orders of magnitude smaller. The qubit state can be read out via state-dependent resonance fluorescence measurement. Note that mid-circuit measurement and reset can be implemented while causing a small crosstalk error due to the stray light from the measurement and reset laser beams.

### B. Circuit compilation and optimization

We now present the circuit compilation procedure for the experiments on the trapped-ion quantum processor. The H-series devices used in this work have at most five gate zones,<sup>39</sup> i.e. at most five two-qubit gates can be executed in parallel. This implies that optimizing the circuit for full parallelism may result in diminishing returns past five parallel gates. In any case, the highest error operation on the devices is the two-qubit gate, so the primary limiting factor for achieving high-fidelity results is the two-qubit gate count. Thus, we chose to first optimize the two-qubit gate count. In addition, the cost operator is the composition of diagonal gates, and thus we are free to apply them in any order. We optimize the order to maximize the number of gate cancellations. A similar approach has been used in Ref. 40 for devices with nearest-neighbor connectivity.

We start by decomposing the four-body terms  $R_{zzzz}(\gamma)$  into four CNOTs and a single  $R_{zz}(\gamma)$ , where  $R_{zzzz}(\gamma)$  and  $R_{zz}(\gamma)$  denote evolution under zzzz and zz coupling with angle  $\gamma$ , respectively. Note that the  $R_{zz}(\gamma) = e^{-i\frac{\gamma}{2}zz}$  is the native gate for the Quantinuum H-series trapped-ion processors.

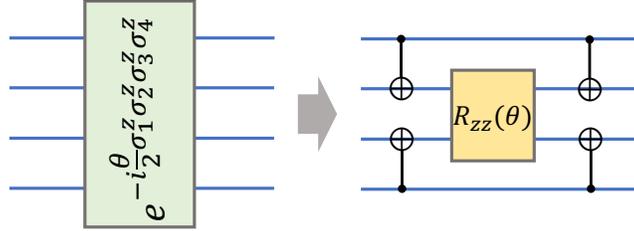


FIG. S14. Decomposition of four-body interaction terms into a two-body  $R_{zz}$  gate and four CNOT's. Figure reproduced from the main text.

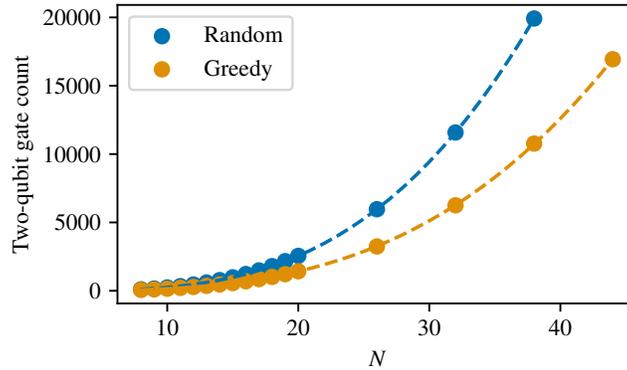


FIG. S15. Comparison of two-qubit gate count (number of  $R_{zz}(\gamma) + R_{zz}(\pi/2)$ ) of QAOA circuit at  $p = 1$  with random term ordering and with the greedy optimization. The “random” line is the average over 20 random orderings of the two- and four-body terms. Both random and greedy are further optimized by tket,<sup>41</sup> and the resulting gate count is plotted. Cubic fit line added to guide the eye.

The goal of the first step of the compilation procedure is to schedule the  $R_{zzzz}(\gamma)$  gates corresponding to four-body terms in a way that greedily cancels as many CNOTs as possible (see Figure S14 for the decomposition). The second step deals with the two-body terms and attempts to schedule each  $R_{zz}(\gamma)$  gate near a four-body term where one of the CNOT acts on the same qubits as the two-body term. This is to leverage two-qubit gate resynthesis implemented in tket.<sup>41</sup> These steps are described in detail in Algorithm 2. Then, the resulting circuit is passed to tket circuit

optimizer to transpile the circuit into the H-series native two-qubit gates:  $R_{zz}(\gamma)$  and  $R_{zz}(\pi/2)$ . The preliminary step of greedily optimizing the layout of the interactions reduced the two-qubit gate count by 1.7 times on average compared to tket alone. The improvement in two-qubit gate count from the greedy CNOT cancellation is shown in Fig. S15. Lastly, we compared a variety of circuit optimizers and found that tket resulted in the most significant gate-count reduction.

---

**Algorithm 2** Greedy cost-operator circuit optimization
 

---

**Require:**

List of four-tuples of  $(i, j, k, l)$ , with  $i < j < k < l$ , where  $(i, j)$  and  $(k, l)$ , respectively correspond to the qubits of top two  $\text{CNOT}_{ij}$  and bottom two  $\text{CNOT}_{lk}$  in decomposition of  $R_{zzzz}$  presented in Figure S14. Note the indices corresponding to the control and target, respectively, are reversed for  $(i, j)$  and  $(k, l)$

List of two-tuples  $(i, j)$ , with  $i < j$ , indicating which qubits to apply each  $R_{zz}$  rotation to.

**Ensure:** Output a single list that contains all of the terms (both four- and two-body) in the order in which they should be applied, according to the greedily-optimized circuit.

circuit  $\leftarrow$  empty list

**for** each collection of four body terms  $(i, j, k, l)$  grouped by locality  $d := j - i (= l - k$  for LABS) **do**

current  $\leftarrow$  uniformly randomly sample (and remove) a term  $(i, j, k, l)$  from the collection of terms of locality  $d$ .

add current to the circuit list

tops  $\leftarrow$  list initialized with tuple  $(i, j)$

bottoms  $\leftarrow$  list initialized with tuple  $(k, l)$

**while** there are still more terms in the collection **do**

**for** each term  $(r, s, t, v)$  in the collection **do**

**if**  $(r, s) \in$  tops or  $(t, v) \in$  bottoms **then**

Assign the term a score of +1

**else**

Assign the term a score of -1.

**end if**

**if**  $\exists m \mid (m, r) \in$  bottoms or  $\exists a \mid (a, t) \in$  tops **then**

Subtract 1 from terms score. // This implies that inserting this term to the circuit would mean that there is some  $\text{CNOT}_{mr}$  or  $\text{CNOT}_{ta}$  currently in the circuit that can never be cancelled.

**end if**

**end for**

current  $\leftarrow$  term  $(a, b, c, d)$  in collection with highest score

add current to the circuit list

Add  $(a, b)$  to tops (if not already in) and  $(c, d)$  to bottoms (if not already in)

**end while**

**end for**

**for** each two body term  $(i, j)$  **do**

**for** each four body term  $(r, s, t, v)$  in circuit **do**

insert  $(i, j)$  after  $(r, s, t, v)$  in circuit if  $(i, j) = (r, s)$  or  $(i, j) = (t, v)$  and break inner loop

**end for**

**if**  $(i, j)$  was not inserted into circuit **then**

Add  $(i, j)$  to end of circuit

**end if**

**end for**

output circuit

---

We note that there exists an alternative proposal<sup>42</sup> for reducing the cost of implementing the LABS cost operator. In this approach, the phase operator is replaced by the evolution under the Hamiltonian corresponding to

$$\mathcal{E}_{\text{sidelobe}}(\mathbf{s}) = \sum_{k=1}^{N-1} |\mathcal{A}_k(\mathbf{s})|, \quad (\text{S7})$$

where

$$\mathcal{A}_k(\mathbf{s}) = \sum_{i=1}^{N-k} s_i s_{i+k}. \quad (\text{S8})$$

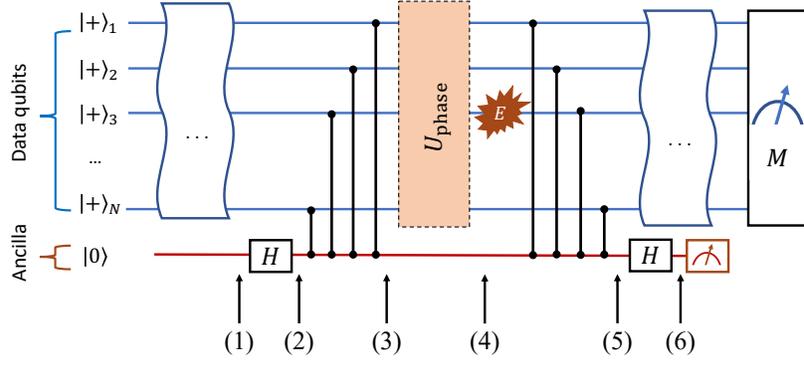


FIG. S16. Overview of one step of the error-detection scheme. A part of the circuit  $U_{\text{phase}}$  is “sandwiched” between two parity checks. Any error on data qubits that does not commute with the check is guaranteed to be detected.

This is in contrast to the approach mentioned in the main text:

$$\mathcal{E}_{\text{sidelobe}}(\mathbf{s}) = \sum_{k=1}^{N-1} \mathcal{A}_k^2(\mathbf{s}). \quad (\text{S9})$$

The Hamiltonian corresponding to the absolute value of the autocorrelations has the same ground state space and energy levels as the one in Equation S9. While this reduces the asymptotic complexity for computing the energy to  $\Theta(N^2)$ , it requires quantum arithmetic, putting it beyond the capability of current hardware. Thus, we focus on optimizing the cost operator corresponding to Equation S9. Further gate count reductions may be possible by fixing some of the variables and applying the techniques of Ref. 43, though doing so is outside of the scope of this work.

### C. Summary of the error-detection scheme

We now briefly summarize the error-detection scheme. The proof that our scheme is capable of detecting an arbitrary single-qubit error in the phase operator circuit under the assumption of noiselessly implemented parity checks is a special case of Ref. 44, Theorem 1. We include a brief discussion of the scheme here for completeness and refer interested readers to Ref. 44 for a detailed discussion. We first consider the case of when just one parity is checked (either x or z), and then show how the analysis generalizes to the case when both parities are checked simultaneously. Figure S16 presents the overview of the circuit. Let the state of the data qubits before the check be  $\rho_{\text{init}}$  and assume the ancilla is perfect and initialized to the pure state  $|0\rangle$ . The z parity check is given by the operator  $C_z = \frac{1}{2} (\mathbb{I} \otimes |0\rangle\langle 0| + Z^{\otimes N} \otimes |1\rangle\langle 1|)$ . Denote the part of the phase operator sandwiched between the checks as  $U_{\text{phase}}$ . Note that  $U_{\text{phase}}$  can encompass the full phase operator or only a part of it. At (1), the state is  $\rho_{(1)} = \rho_{\text{init}} \otimes |0\rangle\langle 0|$ . At (2), the state is  $\rho_{(2)} = \rho_{\text{init}} \otimes \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2}$ . After the first check is applied, at (3) the state becomes

$$\rho_{(3)} = \frac{1}{2} (\rho_{\text{init}} \otimes |0\rangle\langle 0| + Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} \otimes |1\rangle\langle 1|). \quad (\text{S10})$$

If no error during the application of  $U_{\text{phase}}$  occurs, then the state at point (4) is

$$\rho_{(4)}^{\text{no error}} = \frac{1}{2} \left( U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |0\rangle\langle 0| + U_{\text{phase}} Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} U_{\text{phase}}^\dagger \otimes |1\rangle\langle 1| \right), \quad (\text{S11})$$

and

$$\rho_{(5)}^{\text{no error}} = \frac{1}{2} \left( U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |0\rangle\langle 0| + Z^{\otimes N} U_{\text{phase}} Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} U_{\text{phase}}^\dagger \otimes |1\rangle\langle 1| \right) \quad (\text{S12})$$

$$= \frac{1}{2} \left( U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |0\rangle\langle 0| + U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |1\rangle\langle 1| \right) \quad (\text{S13})$$

$$= U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2}, \quad (\text{S14})$$

so the final state is

$$\rho_{(6)}^{\text{no error}} = U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |0\rangle \langle 0|. \quad (\text{S15})$$

Therefore if no error occurred, measuring the ancillary qubit will always give the measurement outcome 0.

If a single-qubit Pauli error  $E$  occurred during the execution of  $U_{\text{phase}}$ , the state at (4) becomes

$$\rho_{(4)} = \frac{1}{2} \left( EU_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger E^\dagger \otimes |0\rangle \langle 0| + EU_{\text{phase}} Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} U_{\text{phase}}^\dagger E^\dagger \otimes |1\rangle \langle 1| \right), \quad (\text{S16})$$

and

$$\rho_{(5)} = \frac{1}{2} \left( EU_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger E^\dagger \otimes |0\rangle \langle 0| + Z^{\otimes N} EU_{\text{phase}} Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} U_{\text{phase}}^\dagger E^\dagger Z^{\otimes N} \otimes |1\rangle \langle 1| \right). \quad (\text{S17})$$

If  $E$  is a Pauli, it can either commute or anti-commute with the check  $Z^{\otimes N}$ . If it anti-commutes, then

$$\rho_{(6)} = \frac{1}{2} \left( EU_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger E^\dagger \otimes |0\rangle \langle 0| - EZ^{\otimes N} U_{\text{phase}} Z^{\otimes N} \rho_{\text{init}} Z^{\otimes N} U_{\text{phase}}^\dagger E^\dagger Z^{\otimes N} \otimes |1\rangle \langle 1| \right) \quad (\text{S18})$$

$$= EU_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger E^\dagger \otimes \frac{|0\rangle \langle 0| - |1\rangle \langle 1|}{2}, \quad (\text{S19})$$

and the final state

$$\rho_{(6)} = U_{\text{phase}} \rho_{\text{init}} U_{\text{phase}}^\dagger \otimes |1\rangle \langle 1|. \quad (\text{S20})$$

Then measuring the ancillary qubit will return 1, and the error will be detected. A Pauli error that commutes with the check, however, will go undetected.

When both  $x^{\otimes N}$  and  $z^{\otimes N}$  parities are checked, no odd-weight Pauli commutes with both checks. Therefore, any odd-weight Pauli error will be detected by the implemented scheme. If the checks are noiseless, increasing the frequency of checks by reducing the size of the circuit  $U_{\text{phase}}$  between the checks can only improve the final fidelity. However, in practice the checks are noisy, introducing a trade-off between the increase in the errors detected and the errors introduced by the checks themselves.

#### D. Performance of the error-detection scheme

In this section, we provide details on the implementation of the proposed error-detection scheme. For a circuit with  $m$  checks, we separate the phase operator into  $m$  parts, where each part has approximately the same number of two-qubit gates. We observe that increasing the frequency of the parity checks up to  $m = 3$  improves the QAOA performance. Fig. S17a shows improvements in expected merit factor after post-selection. More measurements increases the ratio of samples with detected errors, increasing the overhead of the error-detection scheme in terms of the number of repetitions. Fig. S17b shows how this overhead increases with  $N$  by plotting the decay of the ratio of samples with no error detected to all samples (“post-selection ratio”). This ratio drops to below 10% at  $N \geq 15$  and  $m = 3$ . To trade off the number of repetitions required against the fidelity of the final result, we set  $m = 3$  in our experiments. We further note that in our experiments, the post-selection typically keeps the bitstring with the highest merit factor sampled from experiments, as shown in Fig. S18. Note that in a practical optimization setting, the best of all bitstrings corresponding to valid solutions would be chosen as the output.

An important benefit of our error-detection scheme is reduced time to a high-quality sample, i.e. a sample with no errors detected. The time improvement comes from the ability to stop the execution when a mid-circuit measurement detects an error. This capability is particularly relevant to trapped-ion systems with relatively low clock speeds and very long coherence times enabling such classical feedback. Although available hardware supports this feature, we do not implement the early stopping in our hardware experiments. The time savings provided below are estimates.

We denote the probability that no detectable error occurs during part  $i$  as  $p_i$ . For the case without any mid-circuit syndrome measurement, the average time needed to reach a measurement result with a high merit factor, i.e., no parity error detected for all the check measurements, is given by

$$\bar{t}_1 = t_0 / \left( \prod_i^m p_i \right), \quad (\text{S21})$$

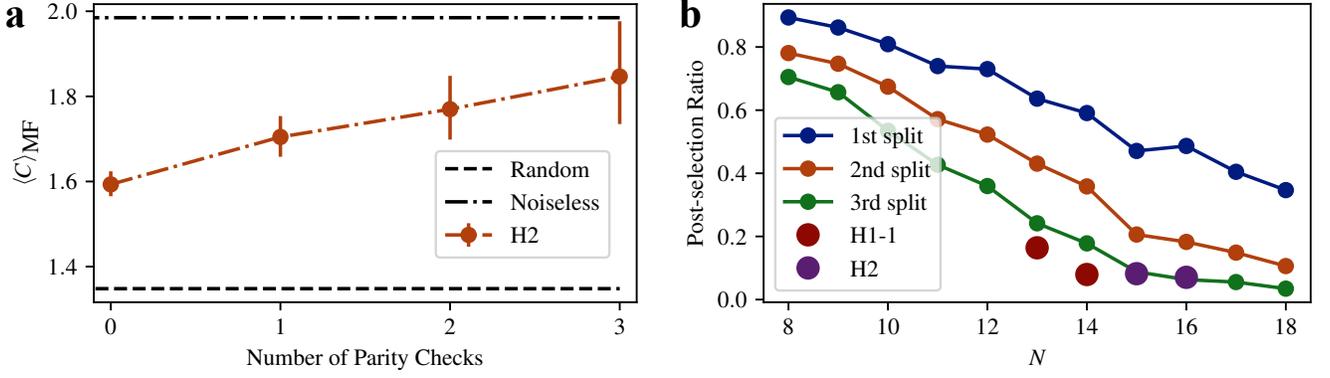


FIG. S17. **a**, Expected merit factor as a function of the number of error checks used when post-selecting the data qubit results for the  $N = 15$  experiment on the H2 hardware system with 5000 repetitions. Here  $m = 3$  is the number of z- and x-parity checks. The error bars become larger as we discard more shots due to more errors being detected with more parity checks. The first data point corresponds to the case when we keep all the results from measurement on the data qubits and do not do post-selection. **b**, Post-selection ratio when splitting the phase operator into three parts and performing parity syndrome measurement at the end of each part. The curves show the simulation results using realistic parameters of the Quantinuum H2 trapped-ion device. The measured results on both H1 and H2 hardware with three checks are in good agreement with simulations.

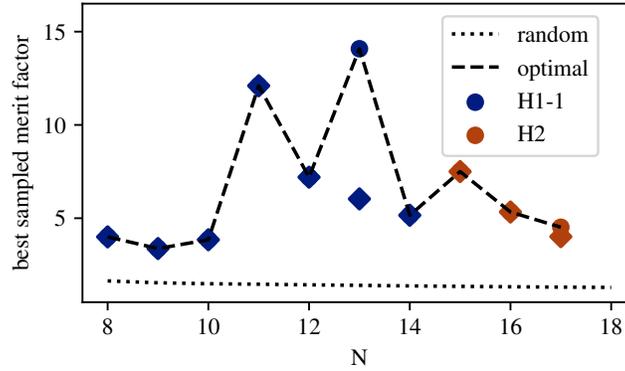


FIG. S18. Experimentally sampled bitstring that has the highest merit factor for different  $N$ . The number of shots taken is 2000 for  $N = 8, 9$  and 5000 for other  $N$ . Circles (diamonds) are the data without (with) post-selection. They have exact same values except for  $N = 13$  and  $N = 17$ . The best sampled bitstrings before post-selection have the same merit factor as true optimal bitstrings for all instances studied here.

where  $t_0$  is the total circuit time. With mid-circuit check and feed-forward discard of the remaining circuit conditioned on the check result, the average time to get a bitstring result for which the merit factor has a high value reads as

$$\bar{t}_2 = t_0 / \left( \prod_i p_i \right) \times \left( \sum_i \left( \prod_{j=0}^{i-1} p_j (1 - p_i) \frac{i}{m} \right) + \prod_k p_k \right), \quad (\text{S22})$$

with  $p_0 = 1$ . Here we neglect the gate time between data qubits and ancillary qubits. The comparison between the two is shown in Fig. S19, indicating that our error check method would reduce the time to get a bitstring with a high merit factor.

\* Corresponding author. Email: ruslan.shaydulin@jpmchase.com

<sup>1</sup> A. Boehmer, *IEEE Transactions on Information Theory* **13**, 156 (1967).

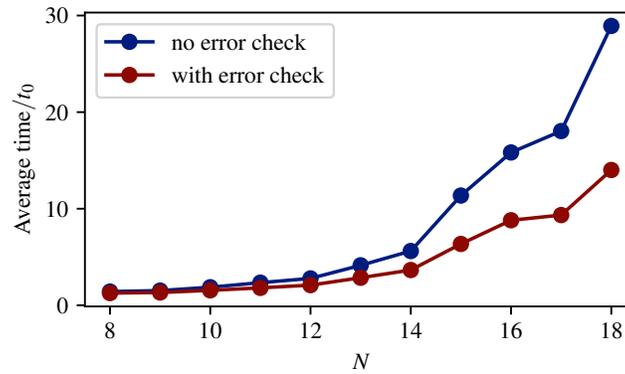


FIG. S19. Simulation of average time (normalized by  $t_0$ ) to a bitstring without detectable parity check errors. The number of splits  $m$  is set to 3. The simulations are performed using realistic parameters of the Quantinuum H2 trapped-ion device.

- <sup>2</sup> M. Schroeder, *IEEE Transactions on Information Theory* **16**, 85 (1970).
- <sup>3</sup> M. Golay, *IEEE Transactions on Information Theory* **18**, 449 (1972).
- <sup>4</sup> M. Golay, *IEEE Transactions on Information Theory* **23**, 43 (1977).
- <sup>5</sup> M. Golay and D. Harris, *IEEE Transactions on Information Theory* **36**, 1163 (1990).
- <sup>6</sup> G. Beenker, T. Claasen, and P. Hermens, *Philips Journal of Research* **40**, 289 (1985).
- <sup>7</sup> J. Jedwab, D. J. Katz, and K.-U. Schmidt, *Journal of Combinatorial Theory, Series A* **120**, 882 (2013).
- <sup>8</sup> M. Golay, *IEEE Transactions on Information Theory* **28**, 543 (1982).
- <sup>9</sup> J. Bernasconi, *Journal de Physique* **48**, 559 (1987).
- <sup>10</sup> B. Bošković, F. Brglez, and J. Brest, *Applied Soft Computing* **56**, 262 (2017).
- <sup>11</sup> E. Farhi, J. Goldstone, and S. Gutmann, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).
- <sup>12</sup> L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, *Physical Review X* **10**, 021067 (2020).
- <sup>13</sup> R. Shaydulin, I. Safro, and J. Larson, in *IEEE High Performance Extreme Computing Conference* (2019).
- <sup>14</sup> G. E. Crooks, [arXiv:1811.08419](https://arxiv.org/abs/1811.08419) (2018).
- <sup>15</sup> M. Streif and M. Leib, *Quantum Science and Technology* **5**, 034008 (2020).
- <sup>16</sup> X. Lee, Y. Saito, D. Cai, and N. Asai, in *International Conference on Quantum Computing and Engineering* (IEEE, 2021).
- <sup>17</sup> S. H. Sack and M. Serbyn, *Quantum* **5**, 491 (2021).
- <sup>18</sup> O. Amosy, T. Danzig, E. Porat, G. Chechik, and A. Makmal, [arXiv:2208.09888](https://arxiv.org/abs/2208.09888) (2022).
- <sup>19</sup> S. Boulebnane and A. Montanaro, [arXiv:2208.06909](https://arxiv.org/abs/2208.06909) (2022).
- <sup>20</sup> J. Larkin, M. Jonsson, D. Justice, and G. G. Guerreschi, *Quantum Science and Technology* **7**, 045014 (2022).
- <sup>21</sup> P. C. Lotshaw, T. S. Humble, R. Herrman, J. Ostrowski, and G. Siopsis, *Quantum Information Processing* **20** (2021).
- <sup>22</sup> L. Li, M. Fan, M. Coram, P. Riley, and S. Leichenauer, *Physical Review Research* **2** (2020).
- <sup>23</sup> P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, *Quantum* **4**, 256 (2020).
- <sup>24</sup> R. Shaydulin, S. Hadfield, T. Hogg, and I. Safro, *Quantum Information Processing* **20** (2021).
- <sup>25</sup> S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, *Physical Review Letters* **125** (2020).
- <sup>26</sup> J. Basso, D. Gamarnik, S. Mei, and L. Zhou, in *IEEE Symposium on Foundations of Computer Science* (2022).
- <sup>27</sup> Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, *Physical Review A* **97** (2018).
- <sup>28</sup> J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, *Proceedings of the Conference on the Theory of Quantum Computation, Communication and Cryptography* **7**, 1 (2022).
- <sup>29</sup> S. H. Sureshbabu, D. Herman, R. Shaydulin, J. Basso, S. Chakrabarti, Y. Sun, and M. Pistoia, [arXiv:2305.15201](https://arxiv.org/abs/2305.15201) (2023).
- <sup>30</sup> S. G. Johnson, “The NLOpt nonlinear-optimization package,” (2022), [http://github.com/stevengj/nlopt](https://github.com/stevengj/nlopt).
- <sup>31</sup> M. J. Powell, *Cambridge NA Report NA2009/06*, University of Cambridge, Cambridge **26** (2009).
- <sup>32</sup> R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, *ACM Transactions on Quantum Computing* **4**, 1 (2023).
- <sup>33</sup> P. C. Lotshaw, G. Siopsis, J. Ostrowski, R. Herrman, R. Alam, S. Powers, and T. S. Humble, “Approximate Boltzmann distributions in quantum approximate optimization,” (2022).
- <sup>34</sup> V. Akshay, H. Philathong, E. Campos, D. Rabinovich, I. Zacharov, X.-M. Zhang, and J. D. Biamonte, *Physical Review A* **106** (2022).
- <sup>35</sup> J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, *Quantum* **4**, 230 (2020).
- <sup>36</sup> G. Brassard, P. Høyer, M. Mosca, and A. Tapp, *Contemporary Mathematics*, 53 (2002).
- <sup>37</sup> T. Packebusch and S. Mertens, *Journal of Physics A: Mathematical and Theoretical* **49**, 165001 (2016).
- <sup>38</sup> 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, *Nature* **592**, 209 (2021).
- <sup>39</sup> S. A. Moses *et al.*, [arXiv:2305.03828](https://arxiv.org/abs/2305.03828) (2023).

- <sup>40</sup> J. Unger, A. Messinger, B. E. Niehoff, M. Fellner, and W. Lechner, [arXiv:2211.11287](#) (2022).
- <sup>41</sup> S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, *Quantum Science and Technology* **6**, 014003 (2020).
- <sup>42</sup> Y. R. Sanders, D. W. Berry, P. C. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, *PRX Quantum* **1** (2020).
- <sup>43</sup> R. Ayanzadeh, N. Alavisamani, P. Das, and M. Qureshi, in *International Conference on Architectural Support for Programming Languages and Operating Systems* (ACM, 2023).
- <sup>44</sup> A. Gonzales, R. Shaydulin, Z. H. Saleem, and M. Suchara, *Scientific Reports* **13** (2023).