

Lawrence Berkeley National Laboratory

LBL Publications

Title

Calibrating the Classical Hardness of the Quantum Approximate Optimization Algorithm

Permalink

<https://escholarship.org/uc/item/6sv7q730>

Authors

Dupont, Maxime

Didier, Nicolas

Hodson, Mark J

et al.

Publication Date

2022-06-13

Calibrating the classical hardness of the quantum approximate optimization algorithm

Maxime Dupont,^{1,2,3,*} Nicolas Didier,³ Mark J. Hodson,³ Joel E. Moore,^{1,2} and Matthew J. Reagor³

¹*Department of Physics, University of California, Berkeley, California 94720, USA*

²*Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

³*Rigetti Computing, 775 Heinz Ave, Berkeley CA 94710, USA*

Trading fidelity for scale enables approximate classical simulators such as matrix product states (MPS) to run quantum circuits beyond exact methods. A control parameter, the so-called bond dimension χ for MPS, governs the allocated computational resources and the output fidelity. Here, we characterize the fidelity for the quantum approximate optimization algorithm by the expectation value of the cost function it seeks to minimize and find that it follows a scaling law $\mathcal{F}(\ln \chi/N)$ with N the number of qubits. With $\ln \chi$ amounting to the entanglement that an MPS can encode, we show that the relevant variable for investigating the fidelity is the entanglement per qubit. Importantly, our results calibrate the classical computational power required to achieve the desired fidelity and benchmark the performance of quantum hardware in a realistic setup. For instance, we quantify the hardness of performing better classically than a noisy superconducting quantum processor by readily matching its output to the scaling function. Moreover, we relate the global fidelity to that of individual operations and establish its relationship with χ and N . We pave the way for noisy quantum computers to outperform classical techniques at running a quantum optimization algorithm in speed, size, and fidelity.

I. INTRODUCTION

Quantum advantage is achieved when a quantum computer outpaces a classical one at doing a specific task in speed, size, or quality [1–6]. Whether a quantum processor succeeds is mainly a binary question. It can be hard to appreciate for noisy intermediate-scale quantum (NISQ) devices due to their small number of qubits and the presence of inherent noise, limiting the success rate of quantum algorithms [7]. For instance, Google’s Sycamore experiment sampling from a random circuit was believed to be out of reach for classical hardware [4]. However, further development in tensor network techniques showed that the corresponding quantum circuits could be executed classically with better performances in a reasonable time [8, 9].

In general, the ability to classically execute a given circuit is different than claiming quantum advantage (it is the same for sampling from a circuit) but can otherwise provide a broader flavor of the original definition. Classically running quantum circuits requires a simulator. Exact tools include, e.g., state vectors, Feynman paths [10], and tensor networks [11] with some superiorly handling qubits over circuit depth and vice versa. Approximate simulators make running large-scale circuits feasible beyond exact ones. It comes at a price of a lower fidelity in the execution, governed nontrivially by a control parameter χ . These methods are mainly based on tensor networks such as matrix product states [12], projected entangled-pair states [13, 14], tree tensor networks [15], or the multi-scale entanglement renormalization ansatz [16].

Matrix product states (MPS) [12] are perhaps best established due to their simplicity and extensive literature, see, e.g., Refs. 17 and 18 for a review. Furthermore, we have a good understanding of how χ , known as the bond dimension, affects the properties of a quantum state represented as an MPS. It amounts to the quantum entanglement that can be

encoded in the system and relates to the entanglement entropy $S \sim \ln \chi$. For N qubits, the memory usage of MPS scales as $O(N\chi^2)$ and the time complexity of individual computational operations scale at most as $O(\chi^3)$. Depending on the classical resources at hand, one can hope to perform calculations up to $\chi \approx 10^2 - 10^4$ (or effectively more when incorporating the symmetries of a circuit). For instance, in the con-




FIG. 1. This summarizes the main results of this work. We perform MPS simulations with bond dimension χ of QAOA for the Max-Cut problem on N qubits. The best output cost that the algorithm seeks to minimize is used for defining the “Fidelity” as the ratio $C(N, \chi)/C(N, \chi_{\text{exact}})$ with $\chi_{\text{exact}} = 2^{N/2}$ the bond dimension required for exact calculations. Each data point corresponds to the statistical average of a hundred QAOA instances with one layer on random 3-regular graphs. Plotting the data as a function of $X = \ln \chi/N$ makes them collapse onto a single curve $Y = \mathcal{F}(X)$, thus revealing a scaling relation. For instance, by matching the output of noisy quantum hardware $Y = 37(3)\%$ on $N = 8$ qubits to the scaling function, one obtains a measure of the classical hardness $\chi = \exp[\mathcal{F}^{-1}(Y)N] \approx 2$ of performing better than the quantum processor.

* Corresponding author: dupont@berkeley.edu

text of quantum computing, MPS have been used to simulate Google’s Sycamore circuits [19], boson sampling [20], Shor’s algorithm [21–23], and the quantum approximate optimization algorithm (QAOA) [24].

QAOA [25–27] is a quantum optimization algorithm seeking to solve quadratic unconstrained binary optimization problems [28] arising in finance, logistics, machine learning, and basic science, to cite but a few. These problems are often NP-hard, with no efficient classical algorithm to address them. For that reason, QAOA is subject to intense research [24–27, 29–72].

Our main contribution is to readily evaluate the classical computational power required to obtain the desired QAOA output and thus provide a practical benchmark for NISQ devices to outperform approximate simulators at running circuits in both size and fidelity. Precisely, we investigate the performance of MPS simulations of QAOA circuits versus N and χ for the paradigmatic Max-Cut problem. We reveal that analyzing the average output cost as a function of a single variable $\ln \chi/N$, which can be interpreted as the entanglement per degree of freedom, leads to a data collapse onto a single curve, unveiling a scaling relation, see Fig. 1. This central result calibrates the classical hardness of obtaining the desired output fidelity, including for quantum hardware. For instance, we benchmark the superconducting quantum chip Rigetti Aspen-M-1 on $N = 8$ qubits by estimating its corresponding χ . Besides, by relating the global fidelity to that of individual operations f , we determine how engineering efforts in increasing N and improving f translate into classical hardness. In addition to the average cost function, we show that a scaling relation involving χ and N also emerges for the probability of finding the bitstring solving the Max-Cut problem, straightforwardly extending our results regarding the cost to this quantity. Finally, we find that deliberately computationally low-cost MPS QAOA simulations may speed up the optimization of finding good circuit parameters, adding to other strategies [30, 47, 49].

II. DEFINITIONS AND METHODS

A. Quantum Approximate Optimization Algorithm

A QAOA circuit for N qubits with p layers reads [25–27],

$$|\beta, \gamma\rangle = \left(\prod_{\ell=1}^p U_{\beta_\ell} U_{\gamma_\ell} \right) H^{\otimes N} |0\rangle^{\otimes N}, \quad (1)$$

with H the Hadamard gate applied on the individual qubits. The parametrized unitaries read $U_{\beta_\ell} = \prod_j \exp(-i\beta_\ell X_j/2)$ and $U_{\gamma_\ell} = \exp(-i\gamma_\ell C/2)$. The cost function operator related to the Max-Cut problem $C = \sum_{\{i,j\} \in E} w_{ij} Z_i Z_j$ [73], is defined for a graph $G = (V, E)$ with edges $\{i, j\} \in E$ carrying a weight $w_{ij} > 0$ (a non-unit weight is often referred to as a weighted Max-Cut problem). X_i and Z_i are Pauli operators on qubit i . The goal is to classically optimize the parameters for the circuit output to minimize the cost function C . We carry out the minimization using the Broyden-Fletcher-Goldfarb-Shanno algorithm [74–77]. For each problem considered, we

repeat the optimization procedure for $\approx 10^3 - 10^4$ random initializations of the parameters and only keep the best result. The initial $2p$ parameters are drawn from the uniform distributions $\gamma \in [0, 2\pi]$ and $\beta \in [0, \pi]$. Strategies have been proposed and benchmarked on the best practices for speeding up the optimization [30, 47, 49].

B. Matrix Product States

An MPS represents a quantum state in a local form where each degree of freedom is associated with a tensor A . The tensors are arranged linearly [12, 17, 18],

$$|\beta, \gamma\rangle = \sum_{\{s\}} A_{a_1}^{s_1} A_{a_1 a_2}^{s_2} A_{a_2 a_3}^{s_3} \cdots A_{a_{N-1}}^{s_N} |s\rangle, \quad (2)$$

where Einstein summation notation is used. The physical index s_i has a dimension two (with value 0 or 1 for a qubit), and a_i is the bond index of dimension χ . One-qubit gates—Hadamard and U_{β_ℓ} in Eq. (1)—are straightforward to apply since the operation only involves a contraction between the gate and the corresponding tensor A of the qubit. We represent the unitary U_{γ_ℓ} as a matrix product operator and apply it onto the MPS [78]. We keep the MPS bond dimension to the desired χ value through singular value decompositions after each QAOA layer. Exact simulations require $\chi \equiv \chi_{\text{exact}} = 2^{N/2}$. The state $|\beta, \gamma\rangle$ is normalized.

III. RESULTS

We consider 3-regular graphs with unit weight $w_{ij} = 1$ together with QAOA depth $p = 1$ (case of Fig. 1) and $p = 2$, as well as complete graphs with uniform random weights $w_{ij} \in [0, 1]$ with QAOA depth $p = 4$. For each of the three cases, we average the expectation value of the cost function operator related to the Max-Cut problem over a hundred randomly generated graphs. We repeat the protocol versus χ for different sizes $N \leq 20$ and note the result $C(N, \chi)$. Due to the nature of the graphs considered, there is no obvious mapping to the native linear topology of an MPS, and we randomly assign graph vertices to tensors. Data plotted in Figs. 2(a) and (b) show that the cost decreases as the bond dimension increases. It suggests that, when rescaled by $C(N, \chi_{\text{exact}})$, the cost can serve as a proxy to the fidelity, i.e., a normalized number assessing the quality of the simulation, one being perfect.

The exact cost $C(N, \chi_{\text{exact}})$ converges asymptotically with the circuit depth p to the absolute minimum cost C_{min} . The absolute solution to the Max-Cut problem, for which there are N_s bitstrings s of cost C_{min} , can be written as $\sum_{\{s\}} e^{i\phi_s} |s\rangle / \sqrt{N_s}$ with ϕ_s a phase. If N_s is not extensive, the state has an exact MPS representation with finite χ . For example, the complete graphs with randomized weights have two solutions ($N_s = 2$): One bitstring and its reverse obtained by the transformation $0 \leftrightarrow 1$. This cat state, similar to a generalized GHZ state for N qubits, has an exact MPS representation with $\chi = 2$. The initial state $H^{\otimes N} |0\rangle^{\otimes N}$ in the QAOA algorithm also has an exact MPS representation with $\chi = 1$ (product state). Nonetheless,




FIG. 2. Left column: 3-regular graphs with unit weight, QAOA depth $p = 2$. Right column: Complete graphs with uniform random weight $w_{ij} \in [0, 1]$, QAOA depth $p = 4$. Top row: Cost versus the bond dimension χ for different sizes N . Bottom row: Same data as the top row but with the Y axis divided by the cost at $\chi_{\text{exact}} = 2^{N/2}$ and the X axis changed to $\ln \chi/N$. The data collapse onto a single curve $Y = \mathcal{F}(X)$, revealing a scaling law. Note the deviation for the smallest size $N = 8$ from the scaling.

simulating QAOA in the restricted space of MPS with bond dimension $\chi = 2$ does not lead to the correct solution. It can be understood by studying the entanglement production and entanglement spreading in QAOA circuits which need to accommodate volume law entanglement $S \sim N$, requiring $\chi \sim \exp(N)$ in intermediate steps [79].

A. Quantifying Classical Hardness

1. Scaling Relation

When plotting this effective fidelity as a function of $\ln \chi/N$, all the data points collapse onto a single curve $Y = \mathcal{F}(X)$, see Fig. 1 and Figs. 2(c) and (d). It reveals a scaling relation of the form,

$$C(N, \chi) / C(N, \chi_{\text{exact}}) = \mathcal{F}(\ln \chi/N), \quad (3)$$

valid for all values of χ and N considered, and supposedly any value. The relation holds for different graphs and different QAOA depths, suggesting universality—though the function \mathcal{F} is different, as discussed in the following. Scaling laws involving χ typically arise in the context of MPS representing quantum critical states in 1 + 1 dimensions [80, 81]. In that case, a finite bond dimension induces a finite length scale

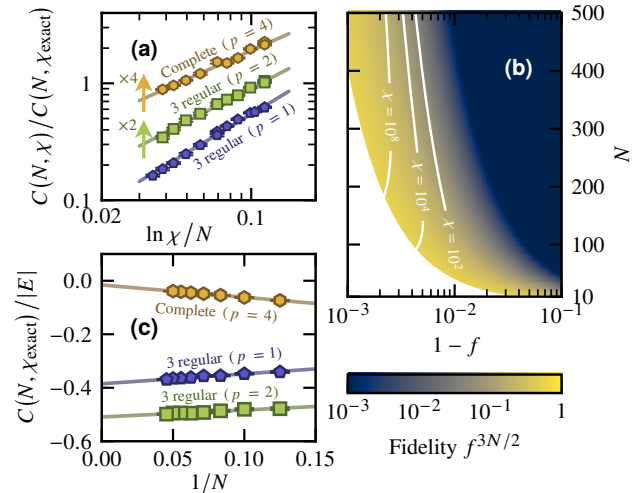


FIG. 3. (a) Scaling function $Y = \mathcal{F}(X \rightarrow 0)$ for unit weight 3-regular graphs with $p = 1$ and $p = 2$ (data shifted vertically by a factor 2 for readability) and complete graphs with uniform random weights $w_{ij} \in [0, 1]$ and $p = 4$ (data shifted vertically by a factor 4 for readability). $N = 8$ data was discarded due to its deviation to the scaling. Bold lines are fit following Eq. (4) with parameters $A = 6.9(6)$, $4.0(3)$, $3.1(2)$, and $\alpha = 1.10(3)$, $0.95(3)$, and $0.82(3)$, respectively. (b) Overall fidelity f^K with f the average fidelity of individual operations and $K = 3N/2$ the number of operations. Using the parameters $A \approx 6.9$ and $\alpha \approx 1.10$ extracted from (a) for the case of unit weight 3-regular graphs with $p = 1$, we plot iso- χ lines according to Eq. (7). (c) Extrapolation of the exact cost for arbitrary N for the three cases considered. Straight lines are fits according to Eq. (5) with parameters $c_0 = -0.385(2)$ and $c_1 = 0.37(3)$, $c_0 = -0.509(3)$ and $c_1 = 0.26(4)$, and $c_0 = -0.0157(4)$ and $c_1 = -0.462(6)$ for unit weight 3-regular graphs with $p = 1$ and $p = 2$ and complete graphs with uniform random weights $w_{ij} \in [0, 1]$ and $p = 4$, respectively.

$\xi \sim \chi$ whose effect is analytically understood by conformal field theories. However, this is not applicable here, leaving the derivation of Eq. (3) an open question. Nevertheless, the rescaled variable $\ln \chi/N$ can be interpreted as the entanglement per degree of freedom and is the relevant quantity to characterize the fidelity of QAOA circuits in MPS simulations, see also App. C [79].

Turning our attention to the scaling function $Y = \mathcal{F}(X)$, we find that in the limit of small X , it follows the functional form,

$$Y = \mathcal{F}(\ln \chi/N \rightarrow 0) \simeq A(\ln \chi/N)^\alpha, \quad (4)$$

valid for $\ln \chi/N \lesssim 0.1$. The algebraic dependence is visible in Fig. 3(a) for the three cases considered in this work. A and α are non-universal parameters depending on the QAOA depth and graphs considered, and which can be extracted by least-square fitting.

2. Exact Cost Normalization

Characterizing the fidelity of a QAOA circuit output by the expectation value of the cost it seeks to minimize requires the

exact cost $C(N, \chi_{\text{exact}})$ for normalization, see Eq. (3). When considering a typical graph within an ensemble, the cost can be extrapolated from data at small N —not necessarily obtained with MPS simulations but simply any exact method. Indeed, we find that the average cost of the ensemble follows,

$$C(N, \chi_{\text{exact}})/|E| \simeq c_0 + c_1/N + O(N^{-2}), \quad (5)$$

with $|E|$ the number of edges, and c_0 and c_1 fitting parameters. Such an extrapolation technique is custom to estimate the thermodynamic limit ($N \rightarrow +\infty$) ground state energy density of quantum many-body Hamiltonians based on system sizes N that can be simulated. We show in Fig. 3(c) that Eq. (5) works well for the three cases considered. Note that extrapolating the exact cost $C(N, \chi_{\text{exact}})$ does not provide information regarding the quantum state $|\beta, \gamma\rangle$ of Eq. (1) for larger N and thus does not solve the original Max-Cut problem of interest.

In perfect simulations, the quality of an output is measured by the approximation ratio $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$ with $C_{\min}(N)$ the absolute minimum cost. In the limit $p \rightarrow +\infty$, the definition of the fidelity of Eq. (3) is the same as r . At finite p , making the substitution $C(N, \chi_{\text{exact}}) \rightarrow C_{\min}(N)$ in Eq. (3) supposes that the two are related by a N -independent factor for the scaling relation to remain valid. It was found not to be the case [35, 47], which we also confirm in App. B, although the small sizes accessible may bias the observation with respect to the large size limit.

3. Extracting Classical Hardness

Fig. 1 and Figs. 2(c) and (d) are calibration curves quantifying the classical computational power required for MPS to achieve the desired fidelity Y for a typical QAOA circuit. By inverting the scaling function $X = \mathcal{F}^{-1}(Y)$, one finds that the corresponding bond dimension for N qubits is,

$$\chi(Y) = \exp[\mathcal{F}^{-1}(Y)N], \quad (6)$$

which provides a direct measure of the classical hardness of running the circuit. In addition, the relationship shows that the difficulty increases exponentially with N . With quantum hardware heading towards $N \sim 10^3$ qubits in the coming years, a fidelity as low as 10% for QAOA on 3-regular graphs with $p = 1$ would require $\chi \sim 2 \times 10^9$ (see Fig. 1) several order or magnitudes beyond the reach of classical computers. For comparison, as of today, the largest reported exact simulation of a QAOA circuit with $p = 1$ on a 3-regular graph involved 210 qubits [44]. Besides, it was argued that a quantum computer should contain at least 420 flawless qubits for QAOA to show a quantum advantage over classical algorithms solving the same class of problems [45], see also Ref. 43.

B. Benchmarking Quantum Hardware

One can use Eq. (6) to benchmark a quantum device by extracting its corresponding χ on a calibration curve $Y = \mathcal{F}(X)$. We perform the experiment on Rigetti Aspen-M-1




FIG. 4. Cost landscape for the $N = 8$ unit weight 3-regular graph of Fig. 5(b) versus the two parameters γ_1 and β_1 of a one-layer QAOA circuit. (a), (b), and (c) Matrix product state simulations with $\chi \equiv \chi_{\text{exact}} = 16$, $\chi = 4$, and $\chi = 2$, respectively. (d) Quantum processor output (Rigetti Aspen-M-1).

superconducting quantum chip. We consider the eight vertices 3-regular graph of Fig. 5(b) with unit weight together with QAOA depth $p = 1$ (case of Fig. 1). There are two parameters γ_1 and β_1 which we discretize on a two-dimensional grid. For each pair of parameters, we collect 1024 output bitstrings from which we compute the expectation value of the cost function operator of the Max-Cut problem. Additional experimental details are available in App. A.

We plot the experimental result in Fig. 4(d). As a comparison, we perform MPS simulations with $\chi \equiv \chi_{\text{exact}} = 16$, $\chi = 4$, and $\chi = 2$, see Figs. 4(a), (b), and (c), respectively. Hardware noise or small bond dimensions reduce the contrast without moving so much the minimum's position, which has also been reported elsewhere [36, 50, 51, 53, 54, 63]. However, low-fidelity QAOA does not solve the original problem. Indeed, a lessened contrast means a higher average cost, indicating that the output distribution of bitstrings also follows this trend and that the bitstrings which minimize the cost function are less likely to appear. We investigate in App. D the probability of finding the bitstrings with absolute minimum cost versus the bond dimension χ and the number of qubits N .

The minimum experimental cost in Fig. 4(d) is $C_{\text{exp}} = -1.9(1)$, associated to a fidelity of 46(3)% with respect to the average exact cost $C(N = 8, \chi_{\text{exact}})$. We repeat the experiment for all possible other $N = 8$ vertices unit-weight 3-regular graphs in App. E (there are five of them in total). Considering all the graphs, we find an average fidelity of 37(3)%, which we report on Fig. 1. We readily extract the corresponding average bond dimension $\chi \approx \exp(0.07 \times 8) \approx 2$, quantifying the classical hardness of doing at least as good with a MPS simulation. On a qualitative level, the MPS simulation with

$\chi = 2$ looks similar to the experimental data with a lower contrast compared to the exact simulation, see Figs. 4(c) and (d), as well as Fig. 9 in App. E.

C. Dissecting the Fidelity

It has been found that a good approximation of the overall fidelity of executing quantum circuit follows the product of each individual operation’s fidelity (assuming no or very little correlations between errors) [4, 19], including the expectation value of the QAOA cost function (see App. E in Ref. 63). Hence, for a circuit comprised of K operations with an average fidelity f , one finds a global fidelity decaying exponentially with K , i.e., f^K . For instance, a total of 16 two-qubit gates were performed experimentally to generate Fig. 4(d) with an overall fidelity of 46(3)%. We estimate the effective fidelity as $f \approx 0.46^{1/16} \approx 95\%$, which is in line with the hardware specification displaying an average fidelity of 96% for two-qubit gates (see App. A), assuming that the two-qubit gates fidelity is the main source of errors.

When put in connection with the scaling relation of Eq. (3), one can relate the number of qubits, the bond dimension, and f together through $f^K \simeq \mathcal{F}(\ln \chi/N)$. Using the functional form of Eq. (4) for the scaling function, one gets,

$$\chi \simeq \exp\left(A^{-1/\alpha} N f^{K/\alpha}\right) \quad \text{for } \ln \chi/N \lesssim 0.1. \quad (7)$$

For the case of unit weight 3-regular graphs with $p = 1$, corresponding to Fig. 1, we get $A = 6.9(6)$ and $\alpha = 1.10(3)$ extracted from Fig. 3(a). We use $K = 3N/2$, corresponding to the number of edges in the graphs, and thus minimum number of two-qubit gates required to be applied for the QAOA circuit, assuming all-to-all connectivity in the topology (this is the case in ion-trapped quantum computers) and that all other operations are perfectly executed. We plot in Fig. 3(b) the overall fidelity as well as iso- χ lines versus f and N . It shows that, for a few hundred of qubits, even small improvements in f can lead to a huge increase in the required bond dimension χ for MPS simulations to beat the hardware. For instance, for $N = 300$ qubits in the regime $f \approx 99.4\%$, improving the error rate $1 - f$ by a factor 1.5, i.e., $f \approx 99.6\%$, makes the bond dimension increase by a factor 100. Fig. 3(b) provides physicists and engineers with a quantitative idea of how incremental experimental efforts translate into classical hardness. We emphasize again that Eq. (7) is valid in the regime $\ln \chi/N \lesssim 0.1$, which is the overall low-fidelity regime. Therefore, hundreds of qubits running a QAOA circuit with a global fidelity as low as $\approx 10^{-2} - 10^{-1}$ are out of reach for approximate MPS simulations. Per the central limit theorem, resolving experimentally such a low fidelity requires a number of shots scaling as its inverse square, which fits in the scope of current processor capabilities.

IV. CONCLUSION AND OUTLOOK

In summary, we address whether noisy qubits in the hundreds and thousand, which is where quantum devices are head-

ing in the coming years, can have a computational advantage over a classical technique such as MPS for the quantum optimization algorithm QAOA. The scaling relation shown in Fig. 1 calibrates through χ how increasing the number of qubits and improving the fidelity of individual operations in a quantum computer makes it harder for classical simulators to beat it—and eventually fail to do so.

All sources of noise were put on the same footing through an average fidelity parameter f . It would be interesting to perform noisy simulations to disentangle how specific errors (relaxation, dephasing, readout, etc.) modify the global fidelity of QAOA circuits, thus bringing the noise relation with χ one level closer to quantum hardware specifications.

We observed that for one-layer QAOA circuits, the main effect of reduced bond dimension is to reduce the contrast of the cost landscape without moving the position of the minimum so much—it has also been reported elsewhere [36, 50, 51, 53, 54, 63]. If this holds for deeper QAOA circuits, running low-fidelity (i.e., low-cost) MPS simulations may reliably estimate the best angles $\{\gamma, \beta\}$, which could then be used to speed up the optimization of larger- χ simulations. It would add to the other strategies that have been proposed [30, 47, 49].

A question left open is why a scaling relation emerges when considering the variable $\ln \chi/N$ —or equivalently $\ln \chi/\ln \chi_{\text{exact}}$ or S/N , see App. C—which to the best of our knowledge, has not been reported elsewhere. It would be interesting to see whether such a scaling law appears in other contexts, such as variational [82] or quantum machine learning algorithms [83, 84], which ultimately seek to minimize a cost function. As reported in this work, it can be used to establish when and how noisy quantum computers will outperform classical techniques at running these circuits, both in size and fidelity. In addition, we reported in App. D that there exists a scaling relation involving χ for the probability of finding the correct bitstrings solving Max-Cut with QAOA. This suggests that the calibration strategy may apply to quantum algorithms with a unique bitstring for an answer, such as Shor’s and Grover’s algorithms [21, 85]—although it is computationally more challenging to evaluate the probability of finding a bitstring than evaluating the average of a cost function.

The calibration of classical hardness through the bond dimension is specific to MPS, and it may be that other approximate classical simulators can perform better. However, the topology of the graphs considered in this work does not clearly make one choice of tensor network-based method better suited than another. For very specific graphs, it may be that a careful ordering of the qubits onto the linear MPS topology leads to substantial improvements compared to a random assignment. Recently, shallow neural networks based on restricted Boltzmann machines [86–88] were used to simulate QAOA circuits [55] with promising results. Unlike our work based on MPS, there is no easy way to estimate the classical difficulty of simulating the circuit for the desired fidelity and no understanding of what would control this—the entanglement per qubit with MPS.

Note that outperforming classical techniques at running circuits is different than a quantum advantage. Indeed, there may exist classical algorithms solving—or attempting to—the




FIG. 5. (a) Native topology on Rigetti Aspen-M-1 for the eight qubits labeled 30, 31, 32, 33, 34, 35, 36, and 37 that we use for the experiment. (b) (c) (d) (e) (f) There exist five different unit-weight 3-regular graphs with $N = 8$ vertices. Graph (b) is considered in Fig. 4 of the main text. Graphs (c) (d) (e) (f) are considered in Fig. 9.

problem of interest differently than by running a quantum circuit. For instance, even if solving the Max-Cut problem is NP-hard, there exist efficient polynomial-time algorithms that find an approximate solution $r\%$ as good as the exact one with high probability. The Goemans-Williamson algorithm [89] guarantees $r = 87.8\%$ for general graphs, a value that can be improved in certain cases [90].

ACKNOWLEDGMENTS

We acknowledge discussions with B. Evert, A.D. Hill, S. Jeffrey, S.L. Tomarken, and J.A. Valery. We also acknowledge discussions with S. Cohen, J-S. Kim, and S. Stanwyck at Nvidia during the early stages of this work. J.E.M. was supported by the Quantum Science Center (QSC), a National Quantum Information Science Research Center of the U.S. Department of Energy (DOE), and a Simons Investigatorship. This research used the Lawrence computational cluster resource provided by the IT Division at the Lawrence Berkeley National Laboratory (supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Award No. DE-AC02-05CH11231). This research also used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC Award No. DDR-ERCAP0022242. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. The experimental results presented here are based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under agreement No. HR00112090058.

Appendix A: Running on Rigetti Aspen-M-1

1. Device Specifications

We use the eight qubits labeled 30, 31, 32, 33, 34, 35, 36, and 37 on Rigetti Aspen-M-1, which are connected in a ring topology, see Fig. 5(a). These qubits have an average relaxation time $T_1 \simeq 29 \mu\text{s}$ and dephasing time $T_2 \simeq 38 \mu\text{s}$. One-qubit $RX(\mathbb{Z}) = \exp(-i\mathbb{Z}\pi X/2)$ gates have an average fidelity of 99.6% under single-qubit randomized benchmarking [91] and $RZ(\theta) = \exp(-i\theta Z/2)$ gates are virtual and thus error-free. The average readout fidelity is 96%. The native two-qubit gates are $\text{CPHASE}(\theta) = \text{diag}(1, 1, 1, e^{i\theta})$ and $\text{XY}(\theta) = \exp[-i\theta(XX + YY)/2]$, which both have an average fidelity of 96%, estimated by single two-qubit gate randomized benchmarking at $\theta = \pi$.

2. QAOA Gates with Native Gates

Implementing QAOA with depth $p = 1$ for 3-regular graphs on the native ring topology of Fig. 5(a) requires the use of SWAP gates. Up to one-qubit gates, it can be implemented as,

$$\text{SWAP} \sim \text{XY}(\pi)\text{CPHASE}(\pi). \quad (\text{A1})$$

The two-qubit gate $\exp(-i\gamma w_{ij} Z_i Z_j / 2)$ with γ the QAOA angle and w_{ij} the weight of the edge between vertices i and j can be implemented using a single two-qubit gate (up to one-qubit gates),

$$\exp(-i\gamma w_{ij} Z_i Z_j / 2) \sim \text{CPHASE}(-2\gamma w_{ij}). \quad (\text{A2})$$

Appendix B: Exact Versus Minimum Cost

We investigate the relationship between the exact cost $C(N, \chi_{\text{exact}})$ for a fixed QAOA circuit and the absolute minimum $C_{\min}(N)$ of the cost function that QAOA seeks to minimize. The ratio $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$ asymptotically converges to one with the QAOA depth p .

We find that r shows an N -dependence (especially visible for larger p values), see Fig. 6(a). It was also observed in other works [35, 47]. Because exact calculations can only be performed on small systems, one cannot exclude that this dependence will disappear as $N \rightarrow +\infty$. With $C(N, \chi_{\text{exact}})$ and $C_{\min}(N)$ related by an N -dependent constant, the scaling relation Eq. (3) does not hold when substituting $C(N, \chi_{\text{exact}})$ by $C_{\min}(N)$. We attempt the data collapse anyway in Figs. 6(b), (c), and (d), which works reasonably well for the unit weight 3-regular graphs where $p = 1$, as the corresponding $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$ curve is mostly flat versus N . If the approximation ratio r was to flatten versus N in a generic case for a fixed QAOA circuit as $N \rightarrow +\infty$, one could probably get good data collapse even for small system sizes by including corrections to the scaling [92, 93].

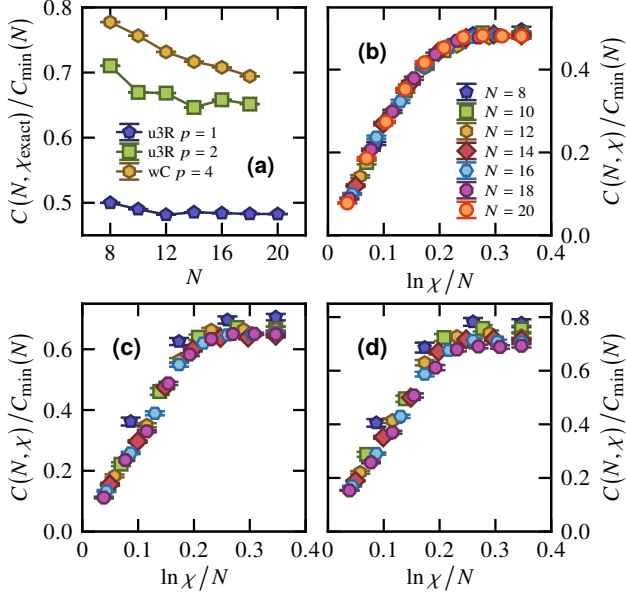


FIG. 6. (a) Approximation ratio $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$ versus the system size N for unit weight 3-regular graphs with $p = 1$ and $p = 2$ and complete graphs with uniform random weights $w_{ij} \in [0, 1]$ and $p = 4$. (b) (c) (d) Data collapse attempts according to Eq. (3) substituting $C(N, \chi_{\text{exact}})$ by $C_{\min}(N)$.

Appendix C: Relation Between the Bond Dimension and the Entanglement

The entanglement entropy puts a number on the degree of quantum entanglement between two subsets of qubits A and B of a system defined over $A \cup B$. The reduced density matrix of the subsystem A , $\rho_A = \text{tr}_B |\beta, \gamma\rangle\langle\beta, \gamma|$ of the pure state $|\beta, \gamma\rangle$ is used to compute the bipartite Von Neumann entanglement entropy between A and B ,

$$S = -\text{tr}(\rho_A \ln \rho_A) = -\sum_k \lambda_k \ln \lambda_k, \quad (\text{C1})$$

with λ_k the eigenvalues of ρ_A . Here, we compute S by cutting the system in half with respect to the left- and right-hand sides of the linear MPS topology according to Eq. (2). The setting is the same as in the main text and we consider the same three cases: Unit weight 3-regular graphs, QAOA depth $p = 1, 2$ as well as complete graphs with uniform random weights $w_{ij} \in [0, 1]$, QAOA depth $p = 4$.

We plot in Figs. 7 (a), (c), and (e) the bipartite Von Neumann entanglement entropy and S versus the bond dimension χ . We verify that $S \sim \ln \chi$ until the entropy saturates for large χ to a volume-law value, i.e., $S \propto N$ [79]. In Figs. 7 (b), (d), and (f), we display the entanglement per qubit versus the rescaled variable $\ln \chi/N$, which is the relevant variable of the scaling relation of Eq. (3). We find a good data collapse, except for small system sizes N , especially in the limit of small bond dimensions χ . Discarding these data points, the collapse suggests the relation $S/N \approx \mathcal{G}(\ln \chi/N)$ with \mathcal{G} a non-universal function with a linear behavior at small $\ln \chi/N$ (from

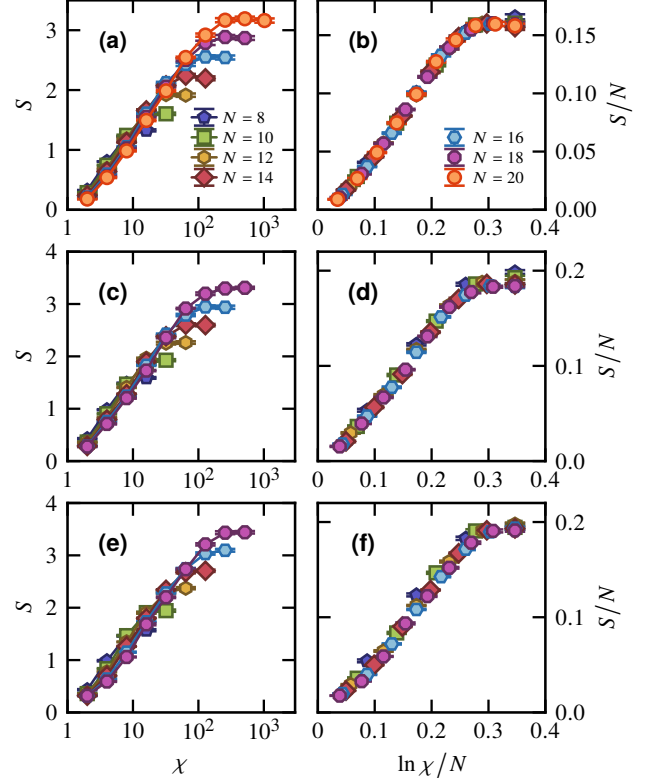


FIG. 7. Top row: 3-regular graphs with unit weight, QAOA depth $p = 1$. Middle row: 3-regular graphs with unit weight, QAOA depth $p = 2$. Bottom row: Complete graphs with uniform random weight $w_{ij} \in [0, 1]$, QAOA depth $p = 4$. Left column: Average bipartite Von Neumann entanglement entropy S from a cut in the middle of the MPS versus the bond dimension χ . Right column: Average bipartite Von Neumann entanglement entropy S from a cut in the middle of the MPS rescaled by the system size N versus the rescaled variable $\ln \chi/N$.

which one recovers $S \sim \ln \chi$) before it saturates. It means that the variable $\ln \chi/N$ in Eq. (3) can indeed be interpreted as the entanglement per qubit S/N .

Appendix D: Probability of Finding the Absolute Solution to the Max-Cut Problem

Through a cost function minimization, QAOA seeks to output bitstrings $\{s\}$ with absolute minimum cost. These bitstrings solve the Max-Cut problem of interest. The probability of finding these bitstrings is,

$$p_{\min} = \sum_{\{s\}} |\langle s | \beta, \gamma \rangle|^2, \quad (\text{D1})$$

with $\langle \beta, \gamma | \beta, \gamma \rangle = 1$. For small enough systems, we can enumerate all bitstrings to find the set $\{s\}$. We plot in Figs. 8(a), (c), and (e) the probability versus the bond dimension χ for different system sizes, for the same three cases considered throughout this work: Unit weight 3-regular graphs, QAOA

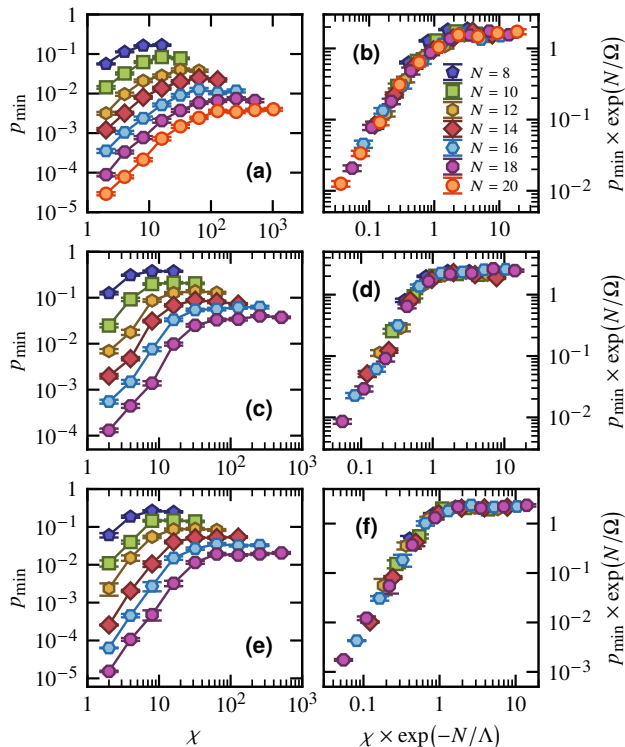


FIG. 8. Top row: 3-regular graphs with unit weight, QAOA depth $p = 1$. Middle row: 3-regular graphs with unit weight, QAOA depth $p = 2$. Bottom row: Complete graphs with uniform random weight $w_{ij} \in [0, 1]$, QAOA depth $p = 4$. Left column: Probability p_{\min} of finding a bitstring minimizing the cost function of the Max-Cut problem versus the bond dimension χ for various system sizes N . Right column: Same data as the left column except that the x axis has been rescaled by $\exp(-N/\Lambda)$ and the y axis by $\exp(N/\Omega)$, with Λ and Ω two parameters such that the data collapse. We find $\Lambda \approx 5$ in the three cases and $\Omega \approx 3.3, 4.3,$ and 3.8 in (b), (d), and (f), respectively.

depth $p = 1, 2$ as well as complete graphs with uniform random weights $w_{ij} \in [0, 1]$, QAOA depth $p = 4$. We observe that the probability decreases with the system size and increases with the bond dimension.

We find that the data points collapse onto a single curve with the following scaling relation,

$$p_{\min} = e^{-N/\Omega} \mathcal{H}\left(\chi e^{-N/\Lambda}\right), \quad (\text{D2})$$

with Ω and Λ parameters, and \mathcal{H} a scaling function, which are a priori non-universal. For $X \equiv \chi e^{-N/\Lambda} \lesssim 1$, we observe an algebraic dependence for the scaling function $\mathcal{H}(X) \simeq DX^\delta$ with parameters D and δ that can be extracted by least-square fitting. For instance, for the case of unit weight 3-regular graphs with QAOA depth $p = 1$, we find $D \approx 1.9$ and $\delta \approx 1.6$. This scaling shows that in its regime of validity, at fixed bond dimension χ , the probability decreases exponentially with N . Similarly, to maintain a fixed probability, the bond

dimension has to scale exponentially with the system size. By randomly sampling bitstrings, the probability of finding the solution $\propto 2^{-N}$ is also exponentially small with the system size. However, QAOA makes the factor in the exponential more favorable. Although it is beyond the scope of this work, one may expect that additional QAOA layers in the circuit will also make it more favorable.

Appendix E: Additional Experiments

In addition to the 3-regular graph of Fig. 5(b) discussed in the main text, we consider all the other unit weight 3-regular graphs with 8 vertices, which are displayed in Figs. 5(c), (d), (e), and (f). For a single layer QAOA, the two parameters γ_1 and β_1 are discretized on a two-dimensional grid. For each pair of parameters, we collect 1024 output bitstrings from which we compute the expectation value of the cost function operator of the Max-Cut problem.

We plot the experimental results in Fig. 9 and compare with MPS simulations using $\chi \equiv \chi_{\text{exact}} = 16$, $\chi = 4$, and $\chi = 2$. As observed previously, hardware noise reduces the contrast without changing the landscape too much otherwise. We attribute the discontinuous behavior of Figs. 9(c) and (o) to the hard cutoff set by the bond dimension $\chi = 2$. A more gentle truncation should adapt χ by considering the numerical value of the singular values being thrown away when performing the singular value decompositions in the MPS method. The relevant numerical values related to the experiments are reported in Table I. The fidelity associated with the graphs of Figs. 5(d), (e), and (f) is lower than for the graphs of Figs. 5(b) and (c) due to the higher number of SWAP gates required to map the graphs to the native topology.

Graph	Exp. Min. Cost	Exact Min. Cost	Fidelity	# SWAP	# Layers 2Q Gates	f
Fig. 5(b)	-1.9(1)	-3.484	46(3)%	2	5	95.3(4)%
Fig. 5(c)	-2.4(2)	-4.612	57(3)%	4	5	97.2(3)%
Fig. 5(d)	-1.0(1)	-4.005	24(3)%	5	11	93.7(6)%
Fig. 5(e)	-1.3(1)	-4.285	32(3)%	4	7	94.5(5)%
Fig. 5(f)	-1.1(1)	-4.612	26(3)%	6	7	94.5(5)%
Average	-1.5(1)	-4.200	37(3)%	4.2	7	95.0(5)%

TABLE I. Relevant numerical data for the five unit weight 3-regular graphs with $N = 8$ vertices that exist. The fidelity is computed with respect to the average exact minimum cost. The individual fidelity f is computed from the global fidelity assuming a number of operations equal to the number of two-qubit gates applied in the circuit. This number equals 12 (number of edges in a 3-regular graph with $N = 8$ vertices) plus twice the number of SWAP gates according to Eqs. (A2) and (A1).

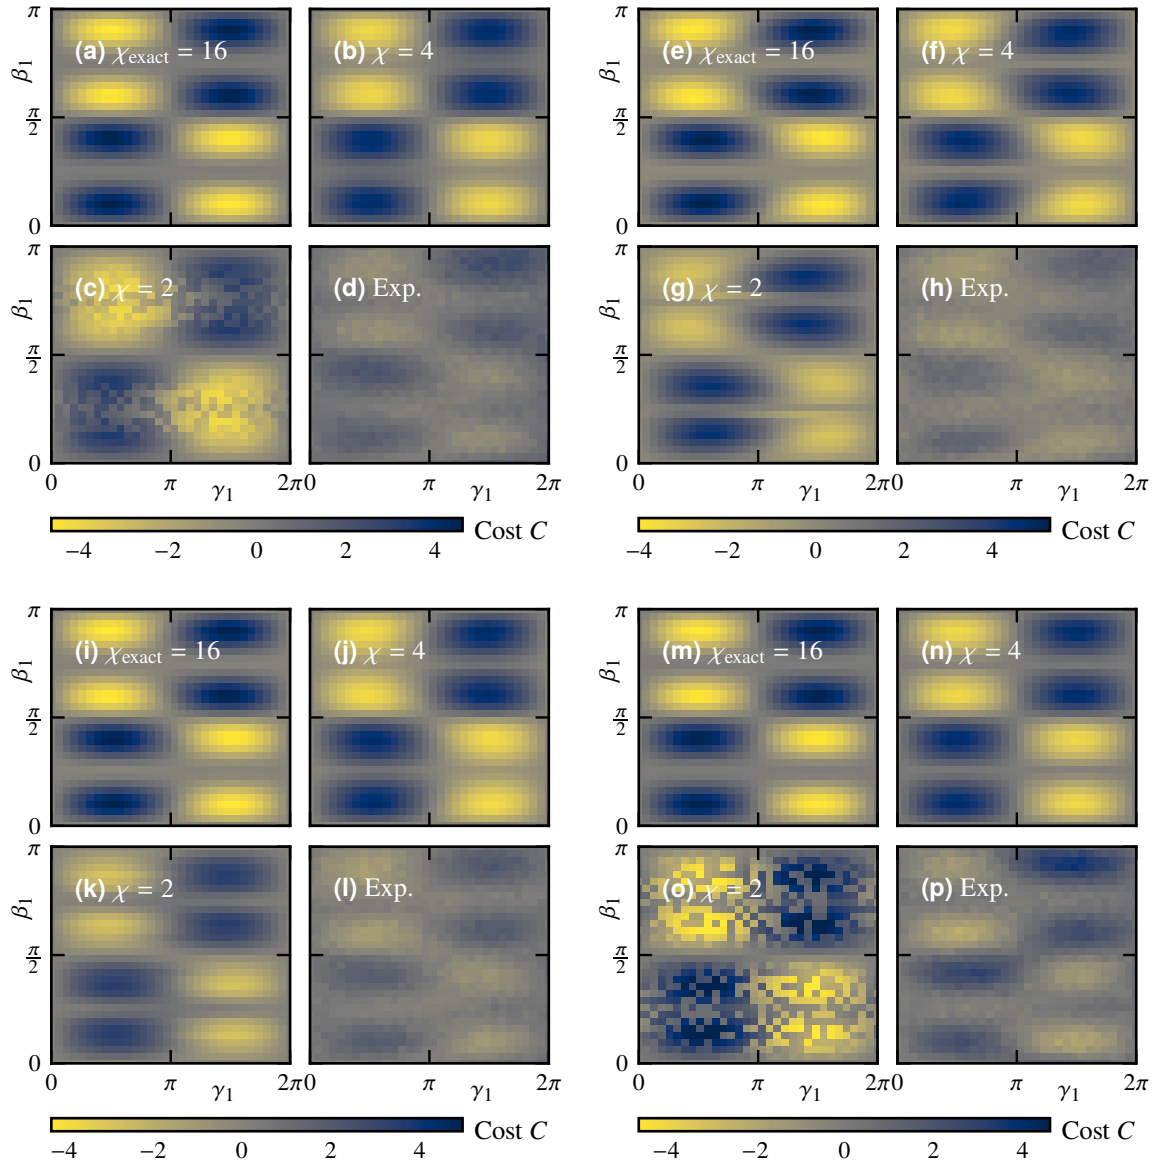


FIG. 9. Cost landscape for $N = 8$ unit weight 3-regular graphs versus the two parameters γ_1 and β_1 of a one-layer QAOA circuit. Comparing matrix product state simulations with $\chi \equiv \chi_{\text{exact}} = 16$, $\chi = 4$, $\chi = 2$ together with the quantum processor output (Rigetti Aspen-M-1). (a) (b) (c) (d) 3-regular graph of Fig. 5(f). (e) (f) (g) (h) 3-regular graph of Fig. 5(d). (i) (j) (k) (l) 3-regular graph of Fig. 5(e). (m) (n) (o) (p) 3-regular graph of Fig. 5(c).

-
- [1] J. Preskill, Quantum computing and the entanglement frontier, [arXiv:1203.5813](https://arxiv.org/abs/1203.5813) (2012).
- [2] A. W. Harrow and A. Montanaro, Quantum computational supremacy, *Nature* **549**, 203 (2017).
- [3] S. Boixo *et al.*, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [4] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [5] H.-S. Zhong *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
- [6] L. S. Madsen *et al.*, Quantum computational advantage with a programmable photonic processor, *Nature* **606**, 75 (2022).
- [7] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [8] H. Cupjin *et al.*, Quantum computing and the entanglement frontier, [arXiv:2005.06787](https://arxiv.org/abs/2005.06787) (2020).
- [9] F. Pan and P. Zhang, Simulation of quantum circuits using the big-batch tensor network method, *Phys. Rev. Lett.* **128**, 030501 (2022).
- [10] E. Bernstein and U. Vazirani, Quantum complexity theory,

- SIAM J. Comput. **26**, 1411 (1997).
- [11] J. C. Bridgeman and C. T. Chubb, Hand-waving and interpretive dance: an introductory course on tensor networks, *J. Phys. A Math.* **50**, 223001 (2017).
- [12] G. Vidal, Efficient simulation of one-dimensional quantum many-body systems, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [13] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, *arXiv:cond-mat/0407066* (2004).
- [14] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the area law, and the computational power of projected entangled pair states, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [15] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [16] G. Vidal, Class of quantum many-body states that can be efficiently simulated, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [17] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011), january 2011 Special Issue.
- [18] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117 (2014).
- [19] Y. Zhou, E. M. Stoudenmire, and X. Waintal, What limits the simulation of quantum computers?, *Phys. Rev. X* **10**, 041038 (2020).
- [20] C. Oh, K. Noh, B. Fefferman, and L. Jiang, Classical simulation of lossy boson sampling using matrix product operators, *Phys. Rev. A* **104**, 022407 (2021).
- [21] P. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124–134.
- [22] D. S. Wang, C. D. Hill, and L. C. L. Hollenberg, Simulations of shor’s algorithm using matrix product states, *Quantum Inf. Process.* **16**, 176 (2017).
- [23] A. Dang, C. D. Hill, and L. C. L. Hollenberg, Optimising Matrix Product State Simulations of Shor’s Algorithm, *Quantum* **3**, 116 (2019).
- [24] T. L. Patti, J. Kossaifi, A. Anandkumara, and S. F. Yelin, Variational quantum optimization with multi-basis encodings, *arXiv:2106.13304* (2021).
- [25] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028* (2014).
- [26] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, *arXiv:1412.6062* (2014).
- [27] E. F. andd Aram W Harrow, Quantum supremacy through the quantum approximate optimization algorithm, *arXiv:1602.07674* (2016).
- [28] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, The unconstrained binary quadratic programming problem: a survey, *J. Comb. Optim.* **28**, 58 (2014).
- [29] D. Wecker, M. B. Hastings, and M. Troyer, Training a quantum optimizer, *Phys. Rev. A* **94**, 022309 (2016).
- [30] G. G. Guerreschi and M. Smelyanskiy, Practical optimization for hybrid quantum-classical algorithms, *arXiv:1701.01450* (2017).
- [31] J. S. Otterbach *et al.*, Unsupervised machine learning on a hybrid quantum computer, *arXiv:1712.05771* (2017).
- [32] Z. Jiang, E. G. Rieffel, and Z. Wang, Near-optimal quantum circuit for grover’s unstructured search using a transverse field, *Phys. Rev. A* **95**, 062317 (2017).
- [33] G. Verdon, M. Broughton, and J. Biamonte, A quantum algorithm to train neural networks using low-depth circuits, *arXiv:1712.05304* (2017).
- [34] S. Lloyd, Quantum approximate optimization is computationally universal, *arXiv:1812.11075* (2018).
- [35] G. E. Crooks, Performance of the quantum approximate optimization algorithm on the maximum cut problem, *arXiv:1811.08419* (2018).
- [36] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O’Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph, J. B. Wang, J. L. O’Brien, M. G. Thompson, and J. C. F. Matthews, Large-scale silicon quantum photonics implementing arbitrary two-qubit processing, *Nat. Photonics* **12**, 534 (2018).
- [37] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for maxcut: A fermionic view, *Phys. Rev. A* **97**, 022304 (2018).
- [38] W. Lechner, Quantum approximate optimization with parallelizable gates, *arXiv:1802.01157* (2018).
- [39] E. R. Anschuetz, J. P. Olson, A. Aspuru-Guzik, and Y. Cao, Variational quantum factoring, *arXiv:1808.08927* (2018).
- [40] B. Sundar, R. Paredes, D. T. Damanik, L. D. nas Osorio, and K. R. A. Hazzard, A quantum algorithm to train neural networks using low-depth circuits, *arXiv:1908.01745* (2019).
- [41] M. B. Hastings, Classical and quantum bounded depth approximation algorithms, *arXiv:1905.07047* (2019).
- [42] S. Hadfield, Z. Wang, B. O’Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 10.3390/a12020034 (2019).
- [43] G. G. Guerreschi and A. Y. Matsuura, Qaoa for max-cut requires hundreds of qubits for quantum speed-up, *Sci. Rep.* **9**, 6903 (2019).
- [44] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, Tensor network quantum simulator with step-dependent parallelization, *arXiv:2012.02430* (2020).
- [45] A. M. Dalzell, A. W. Harrow, D. E. Koh, and R. L. La Placa, How many qubits are needed for quantum computational supremacy?, *Quantum* **4**, 264 (2020).
- [46] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to variational quantum optimization from symmetry protection, *Phys. Rev. Lett.* **125**, 260505 (2020).
- [47] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [48] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [49] D. Wierichs, C. Gogolin, and M. Kastoryano, Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer, *Phys. Rev. Research* **2**, 043246 (2020).
- [50] G. Pagano *et al.*, Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 25396 (2020).
- [51] A. Bengtsson *et al.*, Improved success probability with greater circuit depth for the quantum approximate optimization algorithm, *Phys. Rev. Applied* **14**, 034010 (2020).
- [52] Z. Wang, N. C. Rubin, J. M. Dominy, and E. G. Rieffel, *xy* mixers: Analytical and numerical results for the quantum alternating operator ansatz, *Phys. Rev. A* **101**, 012320 (2020).
- [53] M. Willsch, D. Willsch, F. Jin, H. De Raedt, and K. Michielsen, Benchmarking the quantum approximate optimization algorithm, *Quantum Inf. Process.* **19**, 197 (2020).
- [54] D. M. Abrams, N. Didier, B. R. Johnson, M. P. d. Silva, and C. A. Ryan, Implementation of *xy* entangling gates with a single

- calibrated pulse, *Nat. Electron.* **3**, 744 (2020).
- [55] M. Medvidović and G. Carleo, Classical variational simulation of the quantum approximate optimization algorithm, *npj Quantum Inf.* **7**, 101 (2021).
- [56] V. Kremenetski, T. Hogg, S. Hadfield, S. J. Cotton, and N. M. Tubman, Quantum alternating operator ansatz (qaoa) phase diagrams and applications for quantum chemistry, [arXiv:2108.13056](https://arxiv.org/abs/2108.13056) (2021).
- [57] B. Barak and K. Marwaha, Classical algorithms and quantum limitations for maximum cut on high-girth graphs, [arXiv:2106.05900](https://arxiv.org/abs/2106.05900) (2021).
- [58] D. Fitzek, T. Ghandriz, L. Laine, M. Granath, and A. F. Kockum, Applying quantum approximate optimization to the heterogeneous vehicle routing problem, [arXiv:2110.06799](https://arxiv.org/abs/2110.06799) (2021).
- [59] J. Lee, A. B. Magann, H. A. Rabitz, and C. Arenz, Progress toward favorable landscapes in quantum combinatorial optimization, *Phys. Rev. A* **104**, 032401 (2021).
- [60] E. F. Dumitrescu *et al.*, Benchmarking treewidth as a practical component of tensor network simulations, *PloS one* **13**, e0207827 (2018).
- [61] R. Herrman *et al.*, Globally optimizing qaoa circuit depth for constrained optimization problems, [arXiv:2108.03281](https://arxiv.org/abs/2108.03281) (2021).
- [62] V. Akshay, H. Philathong, I. Zacharov, and J. Biamonte, Reachability Deficits in Quantum Approximate Optimization of Graph Problems, *Quantum* **5**, 532 (2021).
- [63] M. P. Harrigan *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [64] P. Díez-Valle, D. Porras, and J. J. García-Ripoll, Qaoa pseudo-boltzmann states, [arXiv:2201.03358](https://arxiv.org/abs/2201.03358) (2022).
- [65] Y. Zhu *et al.*, Multi-round qaoa and advanced mixers on a trapped-ion quantum computer, [arXiv:2201.12335](https://arxiv.org/abs/2201.12335) (2022).
- [66] M. S. Alam *et al.*, Practical verification of quantum properties in quantum-approximate-optimization runs, *Phys. Rev. Applied* **17**, 024026 (2022).
- [67] S. Ebadi *et al.*, Quantum optimization of maximum independent set using rydberg atom arrays, [arXiv:2202.09372](https://arxiv.org/abs/2202.09372) (2022).
- [68] T. Stollenwerk and S. Hadfield, Diagrammatic analysis for parameterized quantum circuits, [arXiv:2204.01307](https://arxiv.org/abs/2204.01307) (2022).
- [69] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, Filtering variational quantum algorithms for combinatorial optimization, *Quantum Sci. Technol.* **7**, 015021 (2022).
- [70] P. C. Lotshaw *et al.*, Scaling quantum approximate optimization on near-term hardware, [arXiv:2201.02247](https://arxiv.org/abs/2201.02247) (2022).
- [71] G. González-García, R. Trivedi, and J. I. Cirac, Error propagation in nisq devices for solving classical optimization problems, [arXiv:2203.15632](https://arxiv.org/abs/2203.15632) (2022).
- [72] J. Weidenfeller *et al.*, Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware, [arXiv:2202.03459](https://arxiv.org/abs/2202.03459) (2022).
- [73] Another custom definition of the Max-Cut cost function is $\tilde{C} = \sum_{\{i,j\} \in E} w_{ij}(1 - Z_i Z_j)/2$, which directly counts the number of cuts that one seeks to maximize. Up to the constant $\sum_{\{i,j\} \in E} w_{ij}/2$ and a factor one-half, maximizing \tilde{C} is equivalent to minimizing $C = \sum_{\{i,j\} \in E} w_{ij} Z_i Z_j$.
- [74] C. G. Broyden, The convergence of a class of double-rank minimization algorithms 1. general considerations, *IMA J. Appl. Math.* **6**, 76 (1970).
- [75] R. Fletcher, A new approach to variable metric algorithms, *Comput. J.* **13**, 317 (1970).
- [76] D. Goldfarb, A family of variable-metric methods derived by variational means, *Math. Comput.* **24**, 23 (1970).
- [77] D. F. Shanno, Conditioning of quasi-Newton methods for function minimization, *Math. Comput.* **24**, 647 (1970).
- [78] Another strategy, which can be scaled to larger systems, would be to decompose the unitary U_{γ_ℓ} as a product of $|E|$ two-qubit gates, individually applied on the corresponding MPS tensors with the use of internal SWAP gates to accommodate for the native linear topology of MPS. However, scaling to larger systems would not allow us to perform exact simulations for benchmarking. Nevertheless, the two methods have similar effects, see Ref. 19 which used the gate-based one for investigating Google's Sycamore experiment.
- [79] M. Dupont *et al.*, in preparation.
- [80] L. Tagliacozzo, T. R. de Oliveira, S. Iblisdir, and J. I. Latorre, Scaling of entanglement support for matrix product states, *Phys. Rev. B* **78**, 024410 (2008).
- [81] F. Pollmann, S. Mukerjee, A. M. Turner, and J. E. Moore, Theory of finite-entanglement scaling at one-dimensional quantum critical points, *Phys. Rev. Lett.* **102**, 255701 (2009).
- [82] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [83] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
- [84] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [85] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, NY, USA, 1996) p. 212–219.
- [86] G. E. Hinton, Training products of experts by minimizing contrastive divergence, *Neural Comput.* **14**, 1771–1800 (2002).
- [87] G. E. Hinton and R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* **313**, 504 (2006).
- [88] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [89] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [90] E. Halperin, D. Livnat, and U. Zwick, MAX CUT in cubic graphs, *J. Algorithms* **53**, 169 (2004).
- [91] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Phys. Rev. A* **77**, 012307 (2008).
- [92] K. S. D. Beach, L. Wang, and A. W. Sandvik, Data collapse in the critical region using finite-size scaling with subleading corrections, [arXiv:cond-mat/0505194](https://arxiv.org/abs/cond-mat/0505194) (2005).
- [93] L. Wang, K. S. D. Beach, and A. W. Sandvik, High-precision finite-size scaling analysis of the quantum-critical point of $s = 1/2$ heisenberg antiferromagnetic bilayers, *Phys. Rev. B* **73**, 014431 (2006).