

UC Berkeley

UC Berkeley Previously Published Works

Title

Calibrating the Classical Hardness of the Quantum Approximate Optimization Algorithm

Permalink

<https://escholarship.org/uc/item/35d642pr>

Journal

PRX Quantum, 3(4)

ISSN

2691-3399

Authors

Dupont, Maxime

Didier, Nicolas

Hodson, Mark J

et al.

Publication Date

2022-12-01

DOI

10.1103/prxquantum.3.040339

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed


# Calibrating the Classical Hardness of the Quantum Approximate Optimization Algorithm

Maxime Dupont<sup>1,2,3,\*</sup>, Nicolas Didier,<sup>3</sup> Mark J. Hodson,<sup>3</sup> Joel E. Moore,<sup>1,2</sup> and Matthew J. Reagor<sup>3</sup>

<sup>1</sup>*Department of Physics, University of California, Berkeley, California 94720, USA*

<sup>2</sup>*Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

<sup>3</sup>*Rigetti Computing, 775 Heinz Avenue, Berkeley, California 94710, USA*

 (Received 17 June 2022; revised 2 November 2022; accepted 28 November 2022; published 30 December 2022)

The trading of fidelity for scale enables approximate classical simulators such as matrix product states (MPSs) to run quantum circuits beyond exact methods. A control parameter, the so-called bond dimension  $\chi$  for MPSs, 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 that it seeks to minimize and find that it follows a scaling law  $\mathcal{F}(\ln \chi / N)$ , where  $N$  is the number of qubits. With  $\ln \chi$  amounting to the entanglement that a 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  $\chi$  and  $N$ . We sharpen the requirements for noisy quantum computers to outperform classical techniques at running a quantum optimization algorithm in speed, size, and fidelity.

DOI: [10.1103/PRXQuantum.3.040339](https://doi.org/10.1103/PRXQuantum.3.040339)

## 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, which limit the success rate of quantum algorithms [7]. For instance, Google’s “Sycamore” experiment sampling from a random circuit has been believed to be out of reach for classical hardware [4]. However, further development in tensor-network techniques has shown that the corresponding quantum circuits could be executed classically with better performance in a reasonable time [8,9].

In order to run quantum circuits classically, a simulator is required. Exact tools include, e.g., state vectors, Feynman paths [10], and tensor networks [11] with some

handling qubits in a superior manner over circuit depth and vice versa. Approximate simulators make the running of large-scale circuits feasible beyond exact ones. This comes at a price of lower fidelity in the execution, governed nontrivially by a control parameter  $\chi$ . 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 multiscale entanglement-renormalization ansatz [16].

Matrix product states (MPSs) [12] are perhaps best established, due to their simplicity and the extensive literature (for a review, see, e.g., Refs. [17,18]). Furthermore, we have a good understanding of how  $\chi$ , known as the bond dimension, affects the properties of a quantum state represented as a MPS. It amounts to the quantum entanglement that can be encoded in the system and it relates to the entanglement entropy  $S \sim \ln \chi$ . For  $N$  qubits, the memory usage of MPSs scales as  $O(N\chi^2)$  and the time complexity of individual computational operations scales 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 context of quantum computing, MPSs have been used to simulate Google’s “Sycamore” circuits [19], boson sampling [20], Shor’s algorithm [21–23], and

\*[mdupont@rigetti.com](mailto:mdupont@rigetti.com)

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

the quantum approximate optimization algorithm (QAOA) [24].

The 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 nondeterministic polynomial time (NP) hard and there is no efficient classical algorithm to address them. For that reason, the QAOA is subject to intensive research [24–27,29–74].

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  $\chi$  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 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 that for quantum hardware. For instance, we benchmark the Rigetti Aspen-M-1 superconducting quantum chip on  $N = 8$  qubits by estimating its corresponding  $\chi$ . 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  $\chi$  and  $N$  also emerges for the probability of finding the bit string that solves 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],

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

where  $H$  is 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,  $C = \sum_{\{i,j\} \in E} w_{ij} Z_i Z_j$  [75], is defined for a graph  $G = (V, E)$  with edges  $\{i, j\} \in E$  carrying a weight  $w_{ij} > 0$  (a nonunit 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




FIG. 1. A summary of the main results of this work. We perform MPS simulations with bond dimension  $\chi$  of the 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  $C(N, \chi)/C(N, \chi_{\text{exact}})$ , where  $\chi_{\text{exact}} = 2^{N/2}$  is the bond dimension required for exact calculations. Each data point corresponds to the statistical average of 100 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.

output to minimize the cost function  $C$ . We carry out the minimization using the Broyden-Fletcher-Goldfarb-Shannon algorithm [76–79]. For each problem considered, we repeat the optimization procedure for approximately  $10^3$ – $10^4$  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  $\beta \in [0, \pi]$ . Although we do not make use of them, we note that strategies have been proposed and benchmarked for speeding up the optimization of the QAOA [30,47,49].

### B. Matrix product states

A 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]:

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

where Einstein summation notation is used. The physical index  $s_i$  has a dimension of two (with value 0 or 1 for a qubit) and  $a_i$  is the bond index of dimension  $\chi$ . One-qubit gates—Hadamard and  $U_{\beta_\ell}$  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. Unless specified otherwise, we represent the unitary  $U_{\gamma_\ell}$  as a matrix product operator and apply it onto the MPS [80]. We keep the MPS bond dimension to the desired  $\chi$  value through singular-value decompositions after each QAOA layer by keeping only the largest  $\chi$  singular values and dropping the rest (the smallest ones). For reference, exact simulations require  $\chi \equiv \chi_{\text{exact}} = 2^{N/2}$ . The state  $|\beta, \gamma\rangle$  is normalized at the end once the optimal angles are found. This strategy (i) helps to keep the cost landscape a smooth function of the angles [81] and (ii) matches experimental observations (in a sense that will become clear in the following).

### III. RESULTS

We consider 3-regular graphs with unit weight  $w_{ij} = 1$  together with QAOA depth  $p = 1$  (the 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 100 randomly generated graphs. We repeat the protocol versus  $\chi$  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 a MPS and we randomly assign graph vertices to tensors. The data plotted in Figs. 2(a) and 2(b) show that the cost decreases as the bond dimension increases. This suggests that, when rescaled by  $C(N, \chi_{\text{exact}})$ , the cost can serve as a proxy for the fidelity, i.e., a normalized number assessing the quality of the simulation, 1 being perfect.

The exact cost  $C(N, \chi_{\text{exact}})$  converges asymptotically with the circuit depth  $p$  to the absolute minimum cost  $C_{\text{min}}$ . The absolute solution to the Max-Cut problem, for which there are  $N_s$  bit strings  $\mathbf{s}$  of cost  $C_{\text{min}}$ , can be written as  $\sum_{\{\mathbf{s}\}} e^{i\phi_{\mathbf{s}}} |\mathbf{s}\rangle / \sqrt{N_s}$ , where  $\phi_{\mathbf{s}}$  is a phase. If  $N_s \sim O(1)$ , the state has an exact MPS representation with finite  $\chi$ . For example, the complete graphs with randomized weights have two solutions ( $N_s = 2$ ): one bit string and its reverse obtained by the transformation  $0 \leftrightarrow 1$ . This cat state, similar to a generalized Greenberger-Horne-Zeilinger (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  $\chi = 1$  (a product state). Nonetheless, simulating the QAOA in the restricted space of MPSs with bond dimension  $\chi = 2$  does not lead to the correct solution. This 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 [82].




FIG. 2. Left column: 3-regular graphs with unit weight and QAOA depth  $p = 2$ . Right column: complete graphs with uniform random weight  $w_{ij} \in [0, 1]$  and QAOA depth  $p = 4$ . Top row: the cost versus the bond dimension  $\chi$  for different sizes  $N$ . Bottom row: the 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.

#### 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 2(d)]. This reveals a scaling relation of the form

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

which is valid for all values of  $\chi$  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. 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 Appendix C) [82].

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

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

which is valid for  $\ln \chi / N \lesssim 0.1$ . The algebraic dependence is visible in Fig. 3(a) for the three cases considered

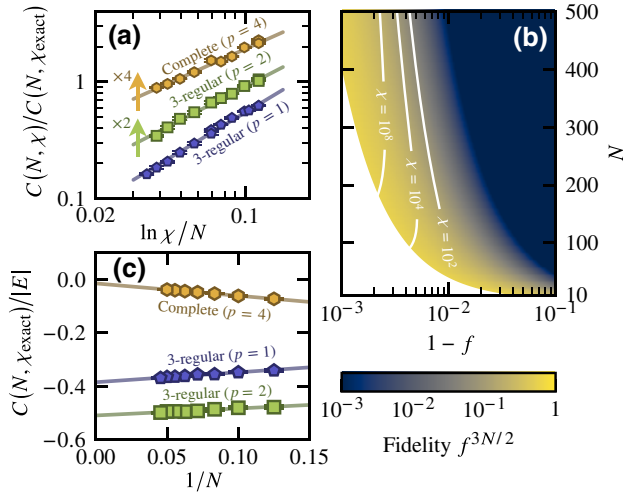


FIG. 3. (a) The scaling function  $Y = \mathcal{F}(X \rightarrow 0)$  for unit-weight 3-regular graphs with  $p = 1$  and  $p = 2$  (data shifted vertically by a factor of 2 for readability) and complete graphs with uniform random weights  $w_{ij} \in [0, 1]$  and  $p = 4$  (data shifted vertically by a factor of 4 for readability). The  $N = 8$  data are discarded due to their deviation to the scaling. The bold lines are a fit following Eq. (4) with parameters  $A = 6.9(6)$ ,  $4.0(3)$ , and  $3.1(2)$ , and  $\alpha = 1.10(3)$ ,  $0.95(3)$ , and  $0.82(3)$ , respectively. (b) The overall fidelity  $f^K$ , where  $f$  is the average fidelity of individual operations and  $K = 3N/2$  is 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- $\chi$  lines according to Eq. (7). (c) Extrapolation of the exact cost for arbitrary  $N$  for the three cases considered. The 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.

in this work.  $A \sim O(1)$  and  $\alpha \approx 1$  are parameters depending on the QAOA depth and graphs considered and which can be extracted by least-squares fitting.

Scaling laws involving  $\chi$  typically arise in the context of MPSs representing quantum critical states in  $1 + 1$  dimensions [83,84]. In that case, a finite bond dimension induces a finite length scale  $\xi \sim \text{poly}(\chi)$ , the effect of which is analytically understood by conformal field theories. However, this is directly not applicable here, leaving the derivation of Eq. (3) as an open question. We believe that one might get further insights into the scaling relation by considering the evolution from the paramagnetic to the ferromagnetic ground state of the one-dimensional quantum Ising model. The corresponding circuit is analogous to the QAOA and achieves a similar purpose. In that case, we know that the critical point between the paramagnetic and ferromagnetic phases governs the properties of the final state with respect to the speed of the evolution process (known as Kibble-Zurek physics; for a review, see, e.g., Ref. [85]). As stated

previously, plugging in a finite bond dimension in a static scenario modifies the critical properties by introducing a finite correlation length  $\xi \sim \text{poly}(\chi)$  [83,84]. Therefore, we believe that investigating Kibble-Zurek physics in the context of finite-bond-dimension simulations would perhaps shed light on some of our findings.

Nevertheless, one can establish a rationale behind Eqs. (3) and (4). As observed in Fig. 2(a) and Appendix D, for  $\chi \ll \chi_{\text{exact}}$ , the cost  $C(N, \chi)$ —non-normalized by data at  $\chi_{\text{exact}}$ —is independent of  $N$  and shows a behavior consistent with  $C(N, \chi) \sim \ln \chi$ . In this regime, one can interpret the algorithm as having too little information to figure out how many qubits there are in the graph. It is analogous to the effect of a finite length scale  $\chi$  in a critical system. As  $\chi \rightarrow \chi_{\text{exact}} \sim \exp(N)$ , the cost saturates to an  $N$ -dependent value, as discussed in the Sec. III A 2. Thus, when the exact cost is extensive  $\sim N$ , as is the case for 3-regular graphs, the previous observation leads directly to  $C(N, \chi) \sim N\mathcal{F}(\ln \chi/N)$  with  $\mathcal{F}(X) \sim X$  at small  $X$ . This is consistent with the exponent  $\alpha \approx 1$  found in Eq. (4) and the additional data of Appendix D, from which we anticipate  $\alpha = 1$ . In particular, we note that the fitted value  $\alpha = 1.10(3) > 1$  in Fig. 2(a) is unlikely to hold asymptotically, as it would imply that  $C(N \rightarrow +\infty, \chi \ll \chi_{\text{exact}}) \rightarrow 0$ , while Appendix D suggests otherwise, with the cost being constant at finite  $\chi \ll \chi_{\text{exact}}$ , independently of  $N$ . For complete graphs, whatever qubits are locally next to each other in the MPS are also connected in the underlying graphs. Hence, in the above picture, at fixed  $\chi \ll \chi_{\text{exact}}$  but increasing  $N$ , the number of edges contributing to the total cost also increases, which is corroborated by the fact that  $\alpha \lesssim 1$  in Eq. (4) and Fig. 3(a) and by the data of Fig. 2(b).

## 2. Exact cost normalization

The characterization of the fidelity of a QAOA circuit output by the expectation value of the cost that 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

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

where  $|E|$  is the total number of edges in the graph and  $c_0$  and  $c_1$  are fitting parameters. Such an extrapolation technique is customary 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 extrapolation of 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  $r = C(N, \chi_{\text{exact}})/C_{\text{min}}(N)$ , where  $C_{\text{min}}(N)$  is 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  $C(N, \chi_{\text{exact}}) \rightarrow C_{\text{min}}(N)$  in Eq. (3) supposes that the two are related by a  $N$ -independent factor for the scaling relation to remain valid. This is found not to be the case [35,47], which we also confirm in Appendix B, although the small sizes accessible may bias the observation with respect to the large size limit.

### 3. Extracting classical hardness

Figure 1 and Figs. 2(c) and 2(d) are calibration curves quantifying the classical computational power required for MPSs 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 toward  $N \sim 10^3$  qubits in the coming years, a fidelity as low as 10% for the QAOA on 3-regular graphs with  $p = 1$  would require  $\chi \sim 2 \times 10^9$  (see Fig. 1), several orders of magnitude 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 has involved 210 qubits [44]. Besides, it has been argued that a quantum computer should contain at least 420 flawless qubits for the 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  $\chi$  on a calibration curve  $Y = \mathcal{F}(X)$ . We perform the experiment on a 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  $p = 1$  (the case of Fig. 1). There are two parameters,  $\gamma_1$  and  $\beta_1$ , that we discretize on a two-dimensional grid. For each pair of parameters, we collect 1024 output bit strings, from which we compute the expectation value of the cost-function operator of the Max-Cut problem. Additional experimental details are available in Appendix A.

We plot the experimental result in Fig. 4(d). As a comparison, we perform MPS simulations with  $\chi \equiv \chi_{\text{exact}} =$

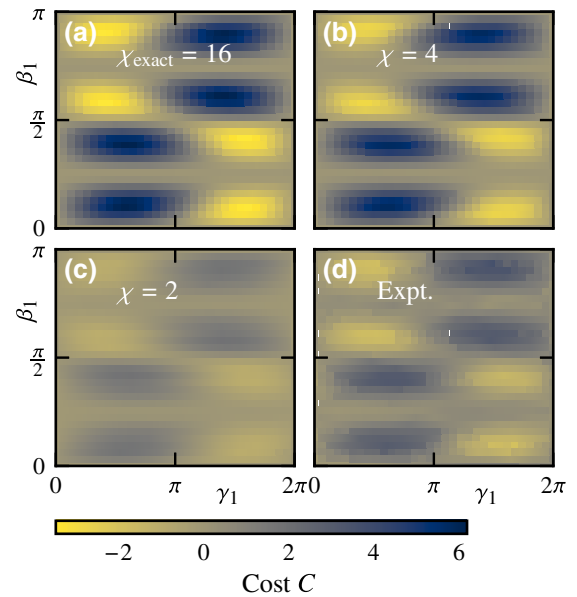


FIG. 4. The cost landscape for the  $N = 8$  unit-weight 3-regular graph of Fig. 5(b) versus the two parameters  $\gamma_1$  and  $\beta_1$  of a one-layer QAOA circuit. (a)–(c) Matrix-product-state simulations with  $\chi \equiv \chi_{\text{exact}} = 16$ ,  $\chi = 4$ , and  $\chi = 2$ , respectively. (d) The quantum processor output (Rigetti Aspen-M-1).

16,  $\chi = 4$ , and  $\chi = 2$  [see Figs. 4(a), 4(b), and 4(c), respectively]. Hardware noise or small bond dimensions reduce the contrast without moving the position of the minimum very much, which has also been reported elsewhere [36,50,51,53,54,63]. However, a low-fidelity QAOA does not solve the original problem. Indeed, a lessened contrast means a higher average cost, indicating that the output distribution of the bit strings also follows this trend and that bit strings that minimize the cost function are less likely to appear. In Appendix E, we investigate the probability of




FIG. 5. (a) The native topology on the 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)–(f) There exist five different unit-weight 3-regular graphs with  $N = 8$  vertices: (b) is considered in Fig. 4 of the main text, while (c)–(f) are considered in Fig. 10.

finding bit strings with absolute minimum cost versus the bond dimension  $\chi$  and the number of qubits  $N$ .

The minimum experimental cost in Fig. 4(d) is  $C_{\text{exp}} = -1.9(1)$ , associated with a fidelity of 46(3)% with respect to the average exact cost  $C(N = 8, \chi_{\text{exact}})$ . In Appendix F, we repeat the experiment for all other possible unit-weight 3-regular graphs having  $N = 8$  vertices (there are five of them in total). Considering all the graphs, we find an average fidelity of 37(3)%, which we report in 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 well using 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 4(d), as well as Fig. 10 in Appendix F].

### C. Dissecting the fidelity

It has been found that a good approximation of the overall fidelity of executing a quantum circuit follows the product of the fidelities of each individual operation (assuming no or very little correlations between errors) [4,19], including the expectation value of the QAOA cost function (see Appendix 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 are 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 the two-qubit gates (see Appendix A), assuming that the fidelity of the two-qubit gates 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, we obtain

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

For the case of unit-weight 3-regular graphs with  $p = 1$ , corresponding to Fig. 1, we obtain  $A = 6.9(6)$  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 the 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. In Fig. 3(b), we plot the overall fidelity as well as iso- $\chi$  lines versus  $f$  and  $N$ . This shows that, for a few hundred qubits, even small improvements in  $f$  can lead to a huge increase in the required bond dimension  $\chi$  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 of 100. Figure 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 approximately  $10^{-2} - 10^{-1}$  are out of reach for approximate MPS simulations. Per the central limit theorem, resolving such a low fidelity experimentally requires a number of shots scaling as its inverse square, which fits within the scope of current processor capabilities.

## IV. CONCLUSIONS AND OUTLOOK

In summary, we address whether noisy qubits in the hundreds and thousands, which is where quantum devices are heading in the coming years, can have a computational advantage over a classical technique such as the MPS for the QAOA. The scaling relation shown in Fig. 1 calibrates through  $\chi$  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 to eventually fail to do so.

All sources of noise are 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  $\chi$  one level closer to quantum hardware specifications.

We observe that for one-layer QAOA circuits, the main effect of a reduced bond dimension is to reduce the contrast of the cost landscape without moving the position of the minimum very much—this 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- $\chi$  simulations. This 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 Appendix 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 [86] or quantum machine-learning algorithms [87,88], which ultimately seek to minimize a cost function. As reported in this work, this 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 report in Appendix E that there exists a scaling relation involving  $\chi$  for the probability of finding the correct bit

strings solving the Max-Cut problem with the QAOA. This suggests that the calibration strategy may apply to quantum algorithms with a unique bit string for an answer, such as Shor’s and Grover’s algorithms [21,89]—although it is computationally more challenging to evaluate the probability of finding a bit string than to evaluate the average of a cost function.

The calibration of classical hardness through the bond dimension is specific to MPSs and it may be that other approximate classical simulators can perform better. For instance, while the topology of the graphs considered in this work does not clearly make one choice of tensor-network-based method better suited than another, it would be interesting to see how other techniques perform. Also, in the context of MPSs and for very specific graphs, a careful qubit ordering onto the linear MPS topology should lead to substantial improvements compared to a random assignment. Recently, shallow neural networks based on restricted Boltzmann machines [90–92] have been used to simulate QAOA circuits [55], with promising results. Unlike our work based on MPSs, 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 MPSs.

Note that outperforming classical techniques at running circuits is different than a quantum advantage. Indeed, there may exist classical algorithms solving—or attempting to solve—the 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  $r\%$  that is as good as the exact one with high probability. The Goemans-Williamson algorithm [93] guarantees  $r = 87.8\%$  for general graphs, a value that can be improved in certain cases [94].

**ACKNOWLEDGMENTS**

We are indebted to one of the anonymous referees, whose suggestions and insights led to the discussion at the end of Sec. III A 1 and Appendix D. 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.

*Note added.*—Recently, we became aware of Refs. [81] and [95], in which QAOA circuits with MPSs are also investigated.

**APPENDIX A: RUNNING ON THE RIGETTI ASPEN-M-1**

**1. Device specifications**

We use the eight qubits labeled 30, 31, 32, 33, 34, 35, 36, and 37 on the Rigetti Aspen-M-1, which are connected in a ring topology [see Fig. 5(a)]. These qubits have an average relaxation time  $T_1 \simeq 29 \mu s$  and dephasing time  $T_2 \simeq 38 \mu s$ . One-qubit rotation about the  $x$  axis  $R_x(\mathbb{Z}) = \exp(-i\mathbb{Z}\pi X/2)$  gates have an average fidelity of 99.6% under single-qubit randomized benchmarking [96] and one-qubit rotation about the  $z$  axis  $R_z(\theta) = \exp(-i\theta Z/2)$  gates are virtual and thus error free. The

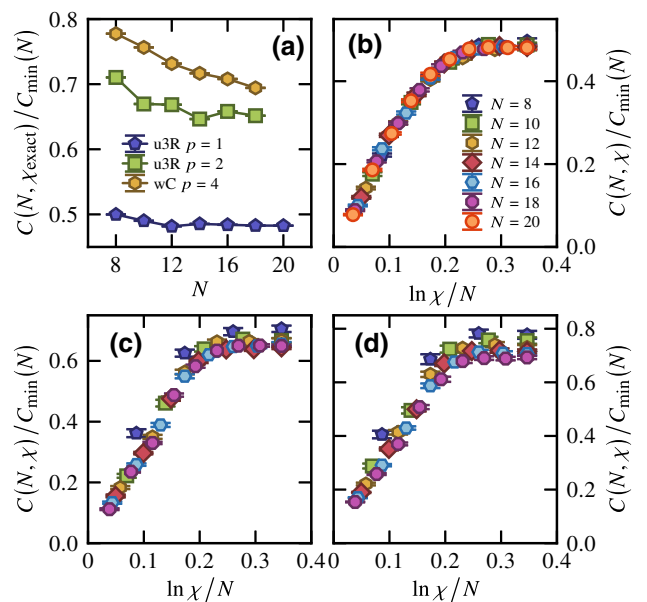


FIG. 6. (a) The approximation ratio  $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$  versus the system size  $N$  for unit-weight 3-regular graphs (u3R) with  $p = 1$  and  $p = 2$  and complete graphs (wC) with uniform random weights  $w_{ij} \in [0, 1]$  and  $p = 4$ . (b)–(d) Data-collapse attempts according to Eq. (3), substituting  $C(N, \chi_{\text{exact}})$  by  $C_{\min}(N)$ .



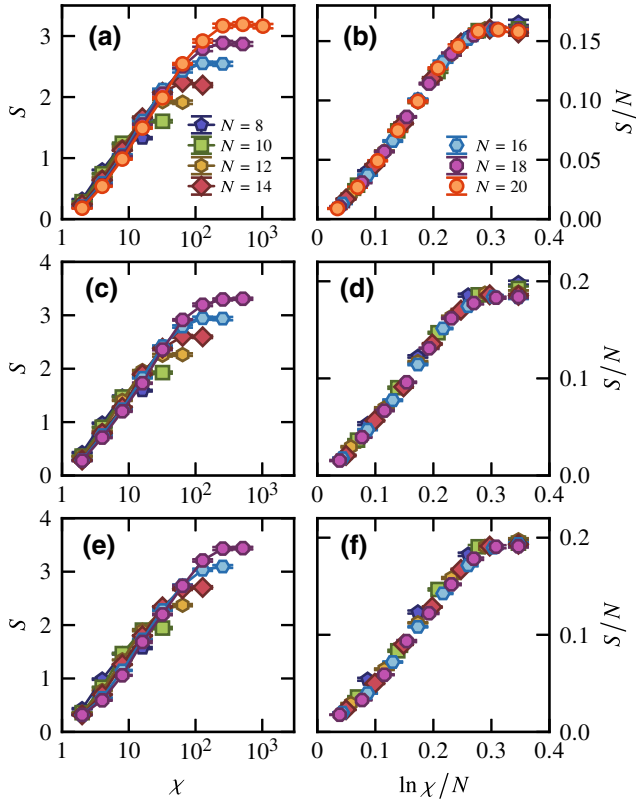


FIG. 7. Top row: 3-regular graphs with unit weight and QAOA depth  $p = 1$ . Middle row: 3-regular graphs with unit weight and QAOA depth  $p = 2$ . Bottom row: complete graphs with uniform random weight  $w_{ij} \in [0, 1]$  and QAOA depth  $p = 4$ . Left column: the average bipartite von Neumann entanglement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

average readout fidelity is 96%. The native two-qubit gates are the controlled-phase  $\text{CPHASE}(\theta) = \text{diag}(1, 1, 1, e^{i\theta})$  and  $\text{XY}(\theta) = \exp[-i\theta(\text{XX} + \text{YY})/2]$ , both of which have an average fidelity of 96%, estimated by single two-qubit gate randomized benchmarking at  $\theta = \pi$ .

## 2. QAOA gates with native gates

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

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

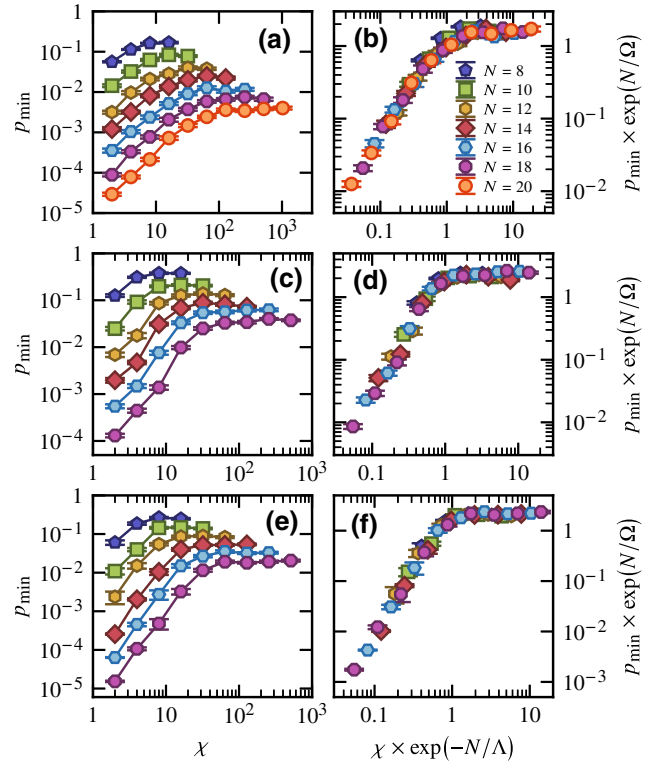


FIG. 8. Top row: 3-regular graphs with unit weight and QAOA depth  $p = 1$ . Middle row: 3-regular graphs with unit weight and QAOA depth  $p = 2$ . Bottom row: complete graphs with uniform random weight  $w_{ij} \in [0, 1]$  and QAOA depth  $p = 4$ . Left column: the probability  $p_{\min}$  of finding a bit string minimizing the cost function of the Max-Cut problem versus the bond dimension  $\chi$  for various system sizes  $N$ . Right column: the same data as the left column, except that the  $x$  axis is rescaled by  $\exp(-N/\Lambda)$  and the  $y$  axis by  $\exp(N/\Omega)$ , where  $\Lambda$  and  $\Omega$  are 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.

## 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 the QAOA seeks to minimize. The ratio  $r = C(N, \chi_{\text{exact}})/C_{\min}(N)$  asymptotically converges to 1 with QAOA depth  $p$ .

We find that  $r$  shows an  $N$  dependence [which is especially visible for larger  $p$  values; see Fig. 6(a)]. This has also been observed in other works [35,47]. Because exact calculations can only be performed on small systems, one cannot exclude the possibility 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)–6(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  $N$ .




FIG. 9. The cost versus the size  $N$  of 3-regular graphs with a unit weight (QAOA depth  $p = 1$ ) for various values of the bond dimension  $\chi$ . Each data point corresponds to the statistical average of 100 independent instances. The straight solid lines are a guide to the eye, highlighting that at fixed  $\chi \ll \chi_{\text{exact}}$ , the cost is roughly constant versus the size  $N$ . Note that unlike other plots in the paper, we do not use a matrix-product-operator representation of the quantum circuit but we apply one-qubit and two-qubit gates individually on the matrix product state representing the quantum state.

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 obtain good data collapse even for small system sizes by including corrections to the scaling [97,98].

#### APPENDIX C: RELATION BETWEEN THE BOND DIMENSION AND THE ENTANGLEMENT

In pure states, 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  $B$ :

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

where the  $\lambda_k$  are the eigenvalues of  $\rho_A$ . 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  $p = 1$  and  $p = 2$ , as well as complete graphs with uniform random weights  $w_{ij} \in [0, 1]$  and QAOA depth  $p = 4$ .

In Figs. 7(a), 7(c), and 7(e), we plot the bipartite von Neumann entanglement entropy and  $S$  versus the bond dimension  $\chi$ . We verify that  $S \sim \ln \chi$  until the entropy

saturates for large  $\chi$  to a volume-law value, i.e.,  $S \propto N$  [82]. In Figs. 7(b), 7(d), and 7(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 good data collapse, except for small system sizes  $N$ , especially in the limit of small bond dimensions  $\chi$ . Discarding these data points, the collapse suggests the relation  $S/N \simeq \mathcal{G}(\ln \chi/N)$ , where  $\mathcal{G}$  is a nonuniversal function with linear behavior at small  $\ln \chi/N$  (from which one recovers  $S \sim \ln \chi$ ) before it saturates. This means that the variable  $\ln \chi/N$  in Eq. (3) can indeed be interpreted as the entanglement per qubit  $S/N$ .

#### APPENDIX D: COST VERSUS SIZE AND BOND DIMENSION

When considering the cost  $C(N, \chi)$  without the need for rescaling by the cost of exact simulations, we can push to larger sizes with  $\chi \ll \chi_{\text{exact}}$ . Note that unlike in other sections of the paper, we do not use a matrix-product-operator representation of the quantum circuit but we apply one-qubit and two-qubit gates individually on the matrix product state representing the quantum state.

We focus on 3-regular graphs with a unit weight and a QAOA depth  $p = 1$ . We show the data in Fig. 9 as a function of  $N$  for various values of the bond dimension  $\chi$ . For  $\chi \ll \chi_{\text{exact}}$ , we find that the cost is independent of  $N$  and saturates to a finite value that decreases as  $\chi$  increases. At fixed  $N$ , the cost behaves as  $C(N, \chi) \sim \ln \chi$ . This behavior is also observed in Figs. 2(a) and 2(b) in different settings for smaller sizes. These data are valuable in Sec. III A 1 when establishing a rationale behind the emergence of the scaling function of Eq. (3).

#### APPENDIX E: PROBABILITY OF FINDING THE ABSOLUTE SOLUTION TO THE MAX-CUT PROBLEM

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

$$p_{\min} = \sum_{\{\mathbf{s}\}} |\langle \mathbf{s} | \beta, \gamma \rangle|^2, \quad (\text{E1})$$

where  $\langle \beta, \gamma | \beta, \gamma \rangle = 1$ . For small enough systems, we can enumerate all bit strings to find the set  $\{\mathbf{s}\}$ . In Figs. 8(a), 8(c), and 8(e), we plot the probability versus the bond dimension  $\chi$  for different system sizes, for the same three cases considered throughout this work: unit-weight 3-regular graphs, QAOA depth  $p = 1$  and  $p = 2$ , 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.




FIG. 10. The cost landscape for  $N = 8$  unit-weight 3-regular graphs versus the two parameters  $\gamma_1$  and  $\beta_1$  of a one-layer QAOA circuit—a comparison of the 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)–(d) the 3-regular graph of Fig. 5(f); (e)–(h) the 3-regular graph of Fig. 5(d); (i)–(l) the 3-regular graph of Fig. 5(e); (m)–(p) the 3-regular graph of Fig. 5(c).

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

$$p_{\min} = e^{-N/\Omega} \mathcal{H}(\chi e^{-N/\Lambda}), \quad (\text{E2})$$

where  $\Omega$  and  $\Lambda$  are parameters and  $\mathcal{H}$  is a scaling function, all of which are *a priori* nonuniversal. 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  $\delta$  that can be extracted by least-squares 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  $\chi$ , 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 bit strings, the probability of finding the solution  $\propto 2^{-N}$  is also exponentially small with the system size. However, the 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.

TABLE I. The 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 (the 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).

Graph	Experimental minimum cost	Exact minimum cost	Fidelity (%)	Number of SWAP gates	Number of layers of 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)

APPENDIX F: 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 eight vertices, which are displayed in Figs. 5(c), 5(d), 5(e), and 5(f). For a single-layer QAOA, the two parameters  $\gamma_1$  and  $\beta_1$  are discretized on a two-dimensional grid. For each pair of parameters, we collect 1024 output bit strings, from which we compute the expectation value of the cost-function operator of the Max-Cut problem.

We plot the experimental results in Fig. 10 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. The discontinuous behavior in the landscape, observed in Figs. 10(c) and 10(o), is due to the choice of mapping from the 3-regular graph to the one-dimensional topology of the MPS. This is verified in Fig. 11, which considers different mappings of the same problem. The relevant numerical values related to the experiments are reported in Table I. The fidelity associated with the graphs of Figs. 5(d), 5(e), and 5(f) is lower than for the graphs of Figs. 5(b) and 5(c) due to the higher number of SWAP gates required to map the graphs to the native topology.




FIG. 11. The cost landscape for the  $N = 8$  unit-weight 3-regular graph of Fig. 5(f) versus the two parameters  $\gamma_1$  and  $\beta_1$  of a one-layer QAOA circuit. (a) An exact simulation with  $\chi = 16$ . (b)–(d) An approximate simulation with  $\chi = 2$ . The mapping from the graph to the one-dimensional topology of the MPS is different in (b)–(d). One observes that the truncation to keep  $\chi = 2$  affects the landscape differently depending on the mapping.

- [1] J. Preskill, Quantum computing and the entanglement frontier, [ArXiv:1203.5813](#) (2012).
- [2] A. W. Harrow and A. Montanaro, Quantum computational supremacy, *Nature* **549**, 203 (2017).
- [3] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [4] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, and B. Burkett, 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, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie, 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] C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, and M. Szegedy, Classical simulation of quantum supremacy circuits, [ArXiv: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: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, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994), p. 124.
- [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. Kossaiifi, 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. Aram and 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, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, and P. Karalekas, 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. 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**, 34 (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. Res. 2, 043246 \(2020\)](#).
- [50] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator, [Proc. Natl. Acad. Sci. USA 117, 25396 \(2020\)](#).
- [51] A. Bengtsson, P. Vikstal, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Križan, D. Shiri, I-M. Svensson, G. Tancredi, G. Johansson, P. Delsing, G. Ferrini, and J. Bylander, Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm, [Phys. Rev. Appl. 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 \(2021\)](#).
- [57] B. Barak and K. Marwaha, in *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 215, edited by M. Braverman (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022), p. 14:1.
- [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 \(2021\)](#).
- [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, A. L. Fisher, T. D. Goodrich, T. S. Humble, B. D. Sullivan, A. L. Wright, and E. G. Dalla Torre, Benchmarking treewidth as a practical component of tensor network simulations, [PLoS ONE 13, e0207827 \(2018\)](#).
- [61] R. Herrman, L. Treffert, J. Ostrowski, P. C. Lotshaw, T. S. Humble, and G. Siopsis, Globally optimizing QAOA circuit depth for constrained optimization problems, [ArXiv: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. Diez-Valle, D. Porras, and J. J. García-Ripoll, QAOA pseudo-Boltzmann states, [ArXiv:2201.03358 \(2022\)](#).
- [65] Y. Zhu, Z. Zhang, B. Sundar, A. M. Green, C. H. Alderete, N. H. Nguyen, K. R. Hazzard, and N. M. Linke, Multi-round QAOA and advanced mixers on a trapped-ion quantum computer, [ArXiv:2201.12335 \(2022\)](#).
- [66] M. S. Alam, F. A. Wudarski, M. J. Reagor, J. Sud, S. Grabbe, Z. Wang, M. Hodson, P. A. Lott, E. G. Rieffel, and D. Venturelli, Practical Verification of Quantum Properties in Quantum-Approximate-Optimization Runs, [Phys. Rev. Appl. 17, 024026 \(2022\)](#).
- [67] S. Ebadi, *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, [Science 376, 1209 \(2022\)](#).
- [68] T. Stollenwerk and S. Hadfield, Diagrammatic analysis for parameterized quantum circuits, [ArXiv: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, T. Nguyen, A. Santana, A. McCaskey, R. Herrman, J. Ostrowski, G. Siopsis, and T. S. Humble, Scaling quantum approximate optimization on near-term hardware, [ArXiv: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 \(2022\)](#).
- [72] J. Weidenfeller, L. C. Valor, J. Gacon, C. Tornow, L. Bello, S. Woerner, and D. J. Egger, Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware, [ArXiv:2202.03459 \(2022\)](#).
- [73] Y. Chen, L. Zhu, C. Liu, N. J. Mayhall, E. Barnes, and S. E. Economou, How much entanglement do quantum optimization algorithms require? [ArXiv:2205.12283 \(2022\)](#).
- [74] G. C. Santra, F. Jendrzejewski, P. Hauke, and D. J. Egger, Squeezing and quantum approximate optimization, [ArXiv:2205.10383 \(2022\)](#).
- [75] 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 of  $1/2$ , maximizing  $\tilde{C}$  is equivalent to minimizing  $C = \sum_{\{i,j\} \in E} w_{ij} Z_i Z_j$ .
- [76] C. G. Broyden, The convergence of a class of double-rank minimization algorithms I. General considerations, [IMA J. Appl. Math. 6, 76 \(1970\)](#).
- [77] R. Fletcher, A new approach to variable metric algorithms, [Comput. J. 13, 317 \(1970\)](#).
- [78] D. Goldfarb, A family of variable-metric methods derived by variational means, [Math. Comput. 24, 23 \(1970\)](#).
- [79] D. F. Shanno, Conditioning of quasi-Newton methods for function minimization, [Math. Comput. 24, 647 \(1970\)](#).
- [80] Another strategy, which can be scaled to larger systems, would be to decompose the unitary  $U_{\gamma_\ell}$  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 MPSs. This other strategy is used in Appendix D and leads to similar results..

- [81] R. Sreedhar, P. Vikstål, M. Svensson, A. Ask, G. Johansson, and L. García-Álvarez, The quantum approximate optimization algorithm performance with low entanglement and high circuit depth, [ArXiv:2207.03404](#) (2022).
- [82] M. Dupont, N. Didier, M. J. Hodson, J. E. Moore, and M. J. Reagor, An entanglement perspective on the quantum approximate optimization algorithm, [ArXiv:2206.07024](#) (2022).
- [83] 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).
- [84] 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).
- [85] A. del Campo and W. H. Zurek, Universality of phase transition dynamics: Topological defects from symmetry breaking, *Int. J. Mod. Phys. A* **29**, 1430018 (2014).
- [86] 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).
- [87] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
- [88] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [89] L. K. Grover, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, 1996), p. 212.
- [90] G. E. Hinton, Training products of experts by minimizing contrastive divergence, *Neural Comput.* **14**, 1771 (2002).
- [91] G. E. Hinton and R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* **313**, 504 (2006).
- [92] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [93] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [94] E. Halperin, D. Livnat, and U. Zwick, MAX CUT in cubic graphs, *J. Algorithms* **53**, 169 (2004).
- [95] T. Ayril, T. Louvet, Y. Zhou, C. Lambert, E. M. Stoudenmire, and X. Waintal, A density-matrix renormalization group algorithm for simulating quantum circuits with a finite fidelity, [ArXiv:2207.05612](#) (2022).
- [96] 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).
- [97] 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](#) (2005).
- [98] 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).