

Large-scale quantum approximate optimization on nonplanar graphs with machine learning noise mitigation

Stefan H. Sack^{1,*} and Daniel J. Egger^{2,†}

¹*Institute of Science and Technology Austria (ISTA), Am Campus 1, 3400 Klosterneuburg, Austria*

²*IBM Quantum, IBM Research Europe – Zurich, Säumerstrasse 4, CH-8803 Rüschlikon, Switzerland*



(Received 18 August 2023; accepted 28 January 2024; published 1 March 2024)

Quantum computers are increasing in size and quality but are still very noisy. Error mitigation extends the size of the quantum circuits that noisy devices can meaningfully execute. However, state-of-the-art error mitigation methods are hard to implement and the limited qubit connectivity in superconducting qubit devices restricts most applications to the hardware's native topology. Here we show a quantum approximate optimization algorithm (QAOA) on nonplanar random regular graphs with up to 40 nodes enabled by a machine learning-based error mitigation. We use a swap network with careful decision-variable-to-qubit mapping and a feed-forward neural network to optimize a depth-two QAOA on up to 40 qubits. We observe a meaningful parameter optimization for the largest graph which requires running quantum circuits with 958 two-qubit gates. Our paper emphasizes the need to mitigate samples, and not only expectation values, in quantum approximate optimization. These results are a step towards executing quantum approximate optimization at a scale that is not classically simulable. Reaching such system sizes is key to properly understanding the true potential of heuristic algorithms like QAOA.

DOI: [10.1103/PhysRevResearch.6.013223](https://doi.org/10.1103/PhysRevResearch.6.013223)

I. INTRODUCTION

Quantum information processing holds the promise of drastically speeding up a number of interesting computational tasks ranging from optimization [1], chemistry [2], finance [3], machine learning [4,5], and high-energy physics [6,7]. The size and quality of noisy quantum computers is progressing rapidly. In particular, error mitigation tools [8,9], e.g., zero-noise extrapolation (ZNE) [8,10] and probabilistic error cancellation (PEC) [8], extend the reach of noisy hardware [11].

In ZNE, multiple logically equivalent copies of a circuit are run under different noise amplification factors c . Based on the noisy results, extrapolation to the zero-noise limit produces a biased estimation of the noiseless expectation value. ZNE can be performed with pulses, which results in small stretch factors c close to one [11]. However, pulse-based ZNE is almost impossible for users of a cloud-based quantum computer to implement due to the onerous calibration. As an alternative, digital ZNE folds gates such as CNOTs which produce large stretch factors $c = 2m + 1$, where m is the number of times a gate is folded [12]. If the original circuit is deep compared to the noise, then the first fold $m = 1$ results in noise

rendering the extrapolation useless. Partial folding prevents this by folding a subset of the gates in a circuit [13].

PEC learns a sparse model of the noise [14]. The non-physical inverse of the noise channel is applied through a quasiprobability distribution to recover an unbiased expectation value. However, the large shot overhead of PEC can be prohibitive. This is why large experiments resort to [robust] probabilistic error amplification (PEA), a form of ZNE in which the learnt error channels are amplified [15]. PEA avoids pulse calibration but requires an onerous noise learning like PEC.

Pulse-based approaches can also error mitigate variational algorithms enabled by, e.g., open pulse [16]. Scaled cross-resonance gates [17,18] can implement ZNE [19] and reduce the schedule duration without calibrating pulses [20]. Other approaches, inspired by optimal control [21], leave it up to the classical optimizer to shape the pulses resulting in shorter schedules [22,23].

Supervised learning benefits a wide range of scientific fields, including quantum physics [24]. In particular, it can mitigate hardware noise in quantum computations. Kim *et al.* [25] adjusted the probabilities estimated from measurements of quantum circuits with neural networks. They show an effective reduction in errors with a method that scales exponentially with system size. Czarnik *et al.* [26] proposed a scalable method to error mitigate observables, rather than the full state vector, with linear regression. They efficiently generated training data by computing expectation values of Clifford circuits on noiseless simulators and noisy quantum hardware. Similarly, Strikis *et al.* [27] presented a method that learns noise mitigation from Clifford data. They error mitigated a quantum circuit by simulating multiple versions of it in which non-Clifford gates are replaced with gates that are efficient to simulate classically. These methods successfully

*stefan.sack@ista.ac.at

†deg@zurich.ibm.com

mitigate noise on both real quantum hardware and simulations of imperfect quantum computers.

Combinatorial problems are regularly encountered in practical settings such as finance and vehicle routing. The quantum approximate optimization algorithm (QAOA) [28] may help solve such problems by mapping the cost function to a spin Hamiltonian and finding its ground state with a suitable variational ansatz [29]. Crucially, many problems of practical interest are nonplanar [30], but common superconducting qubit architectures have a grid [31] or heavy-hexagonal [32] coupling map. Recently, QAOA experiments with a connectivity matching the hardware coupling map have been reported for 27 [20] and 127 [33] qubits with up to QAOA depth-two. By contrast to industry relevant problems, these instances are very sparse. Moreover, classical solvers perform well, especially on sparse problems [34]. While brute-force classical simulation methods of quantum circuits can handle up to around 50 qubits [35,36] tensor-product-based methods are capable of simulating much larger QAOA circuits. For example, Lykov *et al.* [37] reported simulating a single depth-one QAOA amplitude with up to 210 qubits and 1785 gates on a supercomputer. There is, therefore, a dire need to implement denser and larger problems than those in current demonstrations on hardware.

In this paper, we make two contributions. Inspired by Refs. [26,27], we present an error mitigation strategy based on a neural network that uses measurements of noisy observables and compares them to their ideal values. Second, we go one step beyond the hardware-native topology by implementing in hardware random regular-three (RR3) graphs with up to 40 nodes. We achieve this by combining swap networks [20,31], and the SAT-based initial mapping of Matsuo *et al.* [38], which was so far only numerically studied.

This paper is structured as follows. In Sec. II, we introduce the QAOA and discuss its implementation on hardware. Section III discusses machine-learning assisted quantum error mitigation. In Sec. IV, we combine the QAOA implementation advances of Sec. II and the error mitigation approach of Sec. III to train depth-two QAOA circuits on hardware. We discuss our results and conclude in Sec. V.

II. QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

The QAOA was initially developed to solve the maximum cut (MAXCUT) problem [28], but it also applies to any quadratic unconstrained binary optimization (QUBO) as exemplified by Refs. [39–41]. MAXCUT requires cutting the set of nodes V of a given undirected graph $G = (V, E)$ into two groups to maximize the number of edges in E traversed by the cut. This problem, as many others, is equivalent to finding the ground state of an Ising Hamiltonian for an n -qubit system, where $n = |V|$ is the number of decision variables [29].

A depth- p QAOA for an unweighted MAXCUT minimizes the expectation value of the cost function Hamiltonian $H_C = \sum_{(i,j) \in E} \sigma_i^z \sigma_j^z$ under the variational state:

$$|\boldsymbol{\beta}, \boldsymbol{\gamma}\rangle = \prod_{i=1}^p e^{-i\beta_i H_B} e^{-i\gamma_i H_C} |+\rangle^{\otimes n}. \quad (1)$$

The initial product state $|+\rangle^{\otimes n}$ is an equal superposition of all possible solutions. It is also the ground state of the mixer Hamiltonian $H_B = -\sum_i \sigma_i^x$ [28]. The circuit depth, controlled by p , determines the number of applications of the Hamiltonians. A classical optimizer varies the angles $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ and $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)$ to minimize the energy expectation value $E(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \langle H_C \rangle$ in a closed-loop with the quantum computer until the parameters $\boldsymbol{\beta}, \boldsymbol{\gamma}$ converge. We denote the optimized parameters by $\boldsymbol{\theta}^* = (\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) = \arg \min_{\boldsymbol{\beta}, \boldsymbol{\gamma}} E(\boldsymbol{\beta}, \boldsymbol{\gamma})$.

A. Implementation on superconducting hardware

In hardware, $e^{-i\beta_i H_B}$ is trivially implemented by single-qubit R_X rotations applied to all qubits. The cost-operator, however, creates a network of R_{ZZ} gates that matches the graph connectivity. Noisy quantum hardware can run graphs with many nodes if their topology matches the connectivity of the qubits [33]. However, SWAP gates must be inserted in the circuit when the structure of G does not match the native coupling map between the qubits. This severely limits the number of nodes that can be considered [20,42].

Transpiler passes are responsible for routing quantum circuits, i.e., inserting SWAP gates. Transpilers that do not account for gate commutativity in $e^{-i\gamma_i H_C}$ are suboptimal [20]. Commutation-aware transpiler passes have thus been developed [43,44]. Predetermined networks of SWAP gates quickly transpile blocks of commuting two-qubit circuits and produce low-depth circuits compared to other methods [20,31]. However, for problems that are not fully connected, such as MAXCUT on RR3 graphs, predetermined swap networks produce even shallower quantum circuits if the initial mapping from the decision variables to the physical qubits is optimized to minimize the number of swap layers [38].

In this paper, we map the quantum circuits to the best line of qubits on the hardware using alternating layers of SWAP gates [20,31]. The line of qubits is chosen according to the fidelity of the CNOT gates as reported by the back end, see Appendix A. Furthermore, since RR3 graphs are sparse, we reorder the decision variables of the problem to minimize the number of SWAP layers. This is done by a SAT description of the initial mapping problem [38]. Details on the graph generation, transpilation, and SAT mapping are in Appendix A. We first consider two RR3 graphs with 30 and 40 nodes that can be mapped to the hardware with a total of six and seven swap layers, respectively, once the SAT initial mapping is solved. The resulting circuits are transpiled to the hardware native gate set $\{X, \sqrt{X}, R_Z(\theta), \text{ECR}\}$. Here X and \sqrt{X} are the Pauli X gate and its square root. $R_Z(\theta)$ is a rotational Z gate with angle θ and ECR is the echoed cross-resonance gate [45,46]. The ECR gate is equivalent to the standard two-qubit entangling CNOT gate up to single-qubit rotations.

RR3 graphs with 30 and 40 nodes result in large and dense quantum circuits. For example, a depth-one QAOA creates circuits with 305 and 479 ECR gates for $|V| = 30$ and $|V| = 40$, respectively, see Fig. 1(e) which also shows that these circuits leave little space for error suppression methods such as dynamical decoupling [47]. We run the circuits on *ibm_brisbane* and scan the values (γ_1, β_1) from $\pi/2$ to π in 25 steps to investigate if there is a signal without error

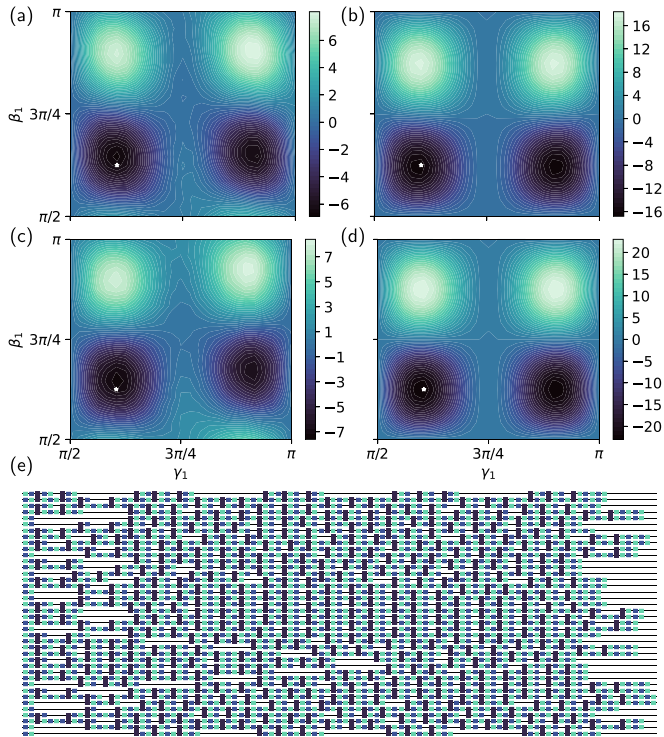


FIG. 1. (a), (b) Depth-one energy landscapes, measured on *ibm_brisbane* and an ideal simulator, respectively, of a 30-node RR3 graph. The circuit (not shown) has 305 ECR gates, 639 X and \sqrt{X} gates, and 804 virtual R_Z gates. (c), (d) The depth-one energy landscapes, measured on *ibm_brisbane* and an ideal simulator, respectively, of a 40 node RR3 graph. The white stars indicate a minima of the noiseless simulations. They reveal a small shift in the corresponding minimum of the hardware-measured data. (e) The quantum circuit of the depth-one QAOA of a RR3 40 node graph transpiled to a line of qubits with seven layers of SWAP gates and a SAT-based initial mapping. The circuit has 479 ECR gates (dark blue), 1021 X and \sqrt{X} gates (blue), and 1275 virtual R_Z gates (light green).

mitigation. We compare the hardware results to an efficient simulation of depth-one QAOA as described in Appendix F of Ref. [48]. The structure of the measured landscape matches the simulations, compare Figs. 1(a) and 1(c) to 1(b) and 1(d), respectively. For the 30 and 40 node graphs, the contrast, i.e., maximum less minimum, of the hardware-measured landscape is 43.0% and 33.8% of the contrast of the simulations, respectively, see the color scales in Fig. 1. The location of the hardware and simulation minima are identical in γ and shifted in β by one grid point, i.e., 65 mrad. Crucially, these results indicate that, despite the large gate count, the quantum computer produces a signal that we can further error mitigate to optimize the parameters of QAOA circuits with $p > 1$.

III. MACHINE LEARNING ASSISTED ERROR MITIGATION

Inspired by Refs. [25–27], we mitigate errors in the energy expectation value with supervised machine learning. We explore a machine-learning approach based on a neural network

to error mitigate QAOA circuits with $p > 1$ during the optimization of γ and β .

A. Supervised machine learning

A supervised machine-learning model requires input data $\mathcal{X} = \{X_i\}_{i=1}^M$ and target data $\mathcal{Y} = \{Y_i\}_{i=1}^M$ to learn the relation between \mathcal{X} and \mathcal{Y} and make predictions on unseen data. Here, M is the data size. We build \mathcal{X} from noisy local expectation values and \mathcal{Y} from the corresponding exact, noise-free, expectation values. The machine-learning model learns the relation from noisy data to the noise-free data. Our proposed method has three steps. First, we generate noisy input data \mathcal{X} on a quantum computer. Second, we simulate the quantum circuits classically to obtain noise-free target data \mathcal{Y} . Finally, we train a machine-learning model to learn the mapping from noisy to noise-free data. The trained model then error mitigates new, i.e., unseen, data.

B. Feed-forward neural network

There is a large number of sophisticated supervised machine learning models. Here, we use a standard fully connected feed-forward neural network (FFNN) [49] due to its simplicity and ease of use. A FFNN is a series of layers. Each layer has multiple neurons that are fully connected to all the neurons in the subsequent layer. This architecture allows the FFNN to model complex nonlinear relationships between the input and output data. We construct our FFNN with an input layer, a single hidden layer, and an output layer. Variational algorithms typically minimize the expectation value of a Hamiltonian built from a linear combination of Pauli expectation values $\sum_i \alpha_i P_i$ with α_i a coefficient and P_i a Pauli operator. To error mitigate a variational algorithm with a FFNN, the output layer must yield quantities that can be optimized. We therefore choose as output layer the correlators that build up the cost function to minimize. The input is a set of noisy observables measured on the quantum computer. The FFNN thus maps noisy observables $\langle P_i \rangle_{\mathcal{N}}$, measured on hardware, to error mitigated observables $\langle P_i \rangle_{\mathcal{M}}$. The subscripts \mathcal{N} and \mathcal{M} indicate noisy and error mitigated observables, respectively.

In the following, we apply the general ideas outlined above to QAOA on a graph $G = (V, E)$ with $|V| = n$ nodes. We choose an input layer with $n(n+1)/2$ neurons. n of these neurons correspond to n noisy local Pauli-Z observables $\langle \sigma_i^z \rangle_{\mathcal{N}}$. The other $n(n-1)/2$ neurons correspond to all possible $\langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{N}}$ correlators, where $i, j = 1, 2, \dots, n$. The output layer is made of $|E|$ neurons; one for each correlator $\langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{M}}$ corresponding to an edge $(i, j) \in E$. Therefore, a RR3 graph uses a FFNN with $3n/2$ output neurons. The number of neurons in the hidden layer is the average of the input and output number of neurons. This construction is illustrated in Fig. 2. A trained FFNN helps us run the QAOA on a quantum computer. Noisy observables are fed into the FFNN for error mitigation. The value of the output neurons is summed to produce an error-mitigated estimation of the energy expectation value $E(\beta, \gamma)_{\mathcal{M}} = \sum_{i,j \in E} \langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{M}}$. This helps optimize γ and β .




FIG. 2. Schematic of the error mitigation with a trained FFNN. The QAOA circuits of graph $G = (V, E)$ are run and the noisy expectation values $\langle \sigma_i^z \rangle_{\mathcal{N}} \forall i \in V$ and $\langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{N}} \forall i, j \in V$ are computed from the sampled counts. These expectation values are the input to a FFNN which outputs the noise mitigated correlators $\langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{M}} \forall (i, j) \in E$. The noise mitigated cost function to optimize is $\sum_{(i,j) \in E} \langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{M}}$.

C. Efficient training data generation

Training the FFNN requires input data \mathcal{X} and target data \mathcal{Y} . We generate \mathcal{X} and \mathcal{Y} by transforming the circuits to error mitigate into classically efficiently simulable circuits. This can be done in multiple ways. According to Ref. [26], it is advantageous to bias the training data towards the state of interest. The QAOA seeks the ground state of H_C , typically a classical product state. It applies the unitaries $e^{-i\beta_k H_B}$ and $e^{-i\gamma_k H_C}$ to drive the initial equal superposition $|+\rangle^{\otimes n}$ towards the ground state of H_C . To generate the training data, we could restrict the angles β_k and γ_k to reduce $e^{-i\beta_k H_B}$ and $e^{-i\gamma_k H_C}$ to Clifford circuits. This would, however, result in a small training data set and may not be possible if the edges in E have noninteger weights. Alternatively, we could randomly replace each R_Z rotation in the transpiled circuit of $e^{-i\gamma_k H_C}$ by a Clifford gate such as I , S , and Z . However, this alters the graph G by giving the edges in E random weights. This may be undesirable as the structure of the QUBO of interest is changed.

These considerations motivate us to train the FFNN on data obtained by sampling over random product states that have undergone a noise process qualitatively similar to the QAOA without altering G . First, we change the initial state from an equal superposition to a random partition of V by randomly applying X gates to the qubits. This initial state is followed by circuit instructions that generate noise similar to the noise in the QAOA. The cost operator (up to SWAP gates which we omit in the following for simplicity) is

$$e^{-i\gamma_k H_C} = \prod_{(i,j) \in E} e^{-i\gamma_k \sigma_i^z \sigma_j^z} = \prod_{(i,j) \in E} CX_{i,j} R_Z^j(2\gamma_k) CX_{i,j}, \quad (2)$$

where $CX_{i,j}$ is a CNOT gate between qubits i and j and R_Z^i is a rotation around the Z axis of qubit i . By setting $\gamma_k = 0$, the operator $e^{-i\gamma_k H_C}$ reduces to the identity (up to SWAP gates) and the QAOA circuit produces product states that we efficiently simulate classically. To retain the noise characteristics, we replace the R_Z gates with barriers to prevent the transpiler from removing the CNOT gates; see Fig. 3. Since R_Z is implemented by virtual phase changes [50], the duration and magnitude of all pulses played on the hardware are unchanged. This preserves the effect of T_1 , T_2 , cross-talk, and other forms of errors. In detail, we generate training data with a set of M

random states,

$$|\beta\rangle = \prod_{i=1}^p e^{-i\beta_i H_B} \prod_{(k,l) \in E} CX_{k,l}^2 [p_j X |0\rangle_j + (1-p_j) |0\rangle_j]^{\otimes n}, \quad (3)$$

which are used to measure the observables for the input data \mathcal{X} . Here, each β_i is a uniform random variable in $[0, 2\pi]$ and p_j is a Bernoulli random binary variable that applies an X gate on qubit j if successful. We choose a 1/2 probability of success for p_j . To compute the target data, we use $CX_{k,l}^2 = \mathbb{I}$, the resulting state is a trivial product state for which it is straightforward to efficiently compute the exact expectation values required for the target data \mathcal{Y} .

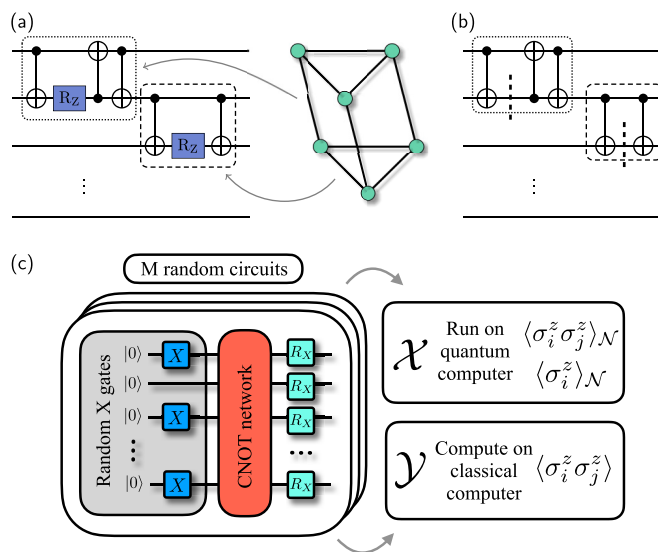


FIG. 3. (a) Part of a QAOA cost operator. The dotted and dashed gates correspond to an R_{ZZ} gate with and without a SWAP gate transpiled to CNOT gates. (b) The R_Z gates are replaced by barriers to prevent CNOT gate cancellation. This preserves the noise structure of the circuit. (c) Training data generation. M random input cuts are created by randomly applying X gates to the qubits. These states are propagated through p alternating networks of CNOT gates, corresponding to simplified $e^{-i\gamma_k H_C}$ operators, and mixer layers $e^{-i\beta_k H_B}$. These circuits are run on hardware to create the input data \mathcal{X} and efficiently simulated classically to generate the ideal output data \mathcal{Y} .

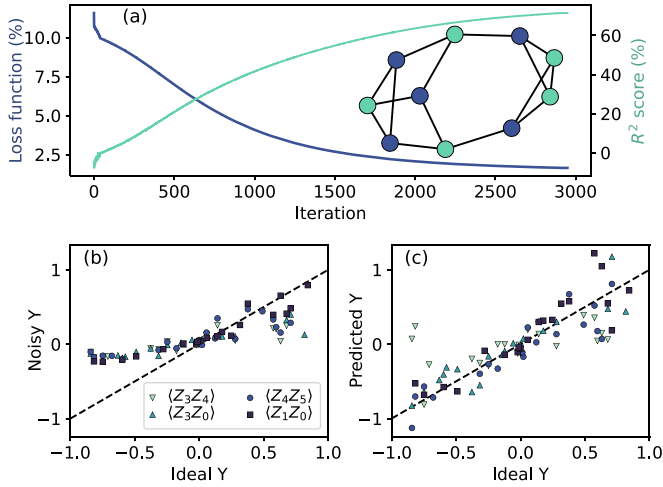


FIG. 4. Training of a FFNN on simulated data. (a) Loss function (dark blue) and R^2 score (light green) as a function of the training iteration. The inset is the ten-node RR3 graph with node color indicating the MAXCUT. (b) Correlators corresponding to edges in the inset graph in (a) before error mitigation versus their ideal value. If the device was noiseless, all the correlators would lie on the dashed line. (c) The correlators of (b) but after error mitigation by the FFNN. If the error mitigation were perfect, all the correlators would lie on the dashed line.

IV. MACHINE-LEARNING ERROR-MITIGATED QAOA

We now apply the FFNN error mitigation discussed in Sec. III and the QAOA execution methods discussed in Sec. II to run depth-two QAOA. We first exemplify the error mitigation in a small ten-qubit simulation and then turn to larger RR3 graphs with 10, 20, 30, and 40 nodes executed on hardware.

A. Simulations

We build a noise model with short-lived qubits. Their T_1 and T_2 times are sampled from a Gaussian distribution with 10 μ s mean and 10 ns standard deviation. Based on these durations, a thermal relaxation noise channel is applied to the CNOT gates lasting $\tau_{\text{CNOT}} = 300$ ns. This is a strong noise model for the 102 CNOT gates in the QAOA circuit as understood, e.g., by $e^{-\tau_{\text{CNOT}}/T_1}$, which gives 97% as proxy for the gate fidelity. The other circuit instructions are noiseless.

We sample 300 random cuts to create the training data following Sec. III C. We train the FFNN with 90% of this data and the other 10% serves as validation data [51]. In this example, the FFNN achieves a mean squared error (MSE) of 3.3% on the training data and an R^2 score of 71.8% on the validation data; see Fig. 4(a). The FFNN thus captures 71.8% of the variation in the validation data. Furthermore, we generate an additional 20 test data points. The corresponding non-error-mitigated correlators are damped towards an expectation value of zero; see Fig. 4(b). The MSE between these $20 \times |E| = 300$ non-error-mitigated and ideal $\langle \sigma_i^z \sigma_j^z \rangle$ correlators is 11%. This number drops to 7% after the correlators are error mitigated with the FFNN. Furthermore, we observe that the error-mitigated correlators better follow the trend set by their ideal values; see Fig. 4(c).

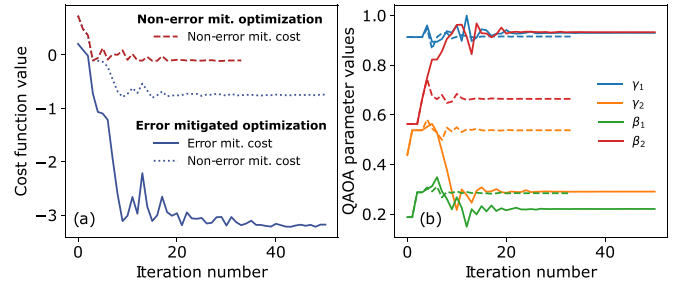


FIG. 5. Optimization of a noisy QAOA. (a) The solid and dotted lines show the error-mitigated cost function and the non-error-mitigated cost function, respectively, during the optimization of the former. The dashed line shows the non-error-mitigated cost function as it is being optimized. (b) QAOA parameters for the error mitigated optimization (solid lines) and the non-error-mitigated optimization (dashed lines).

In a separate simulation, we increase the strength of the noise by lengthening the CNOT gates. At a duration of 400 ns, the FFNN cannot learn an error mitigation since the noise is too strong. We observe that the squared error does not reach low values and the predicted correlators are close to zero (data not shown).

We now optimize a $p = 2$ QAOA with 300 ns CNOT gates twice; once by optimizing the error-mitigated cost function $E_{\mathcal{M}} = \sum_{(i,j) \in E} \langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{M}}$, blue data in Fig. 5(a), and once by optimizing the non-error-mitigated cost function $E_{\mathcal{N}} = \sum_{(i,j) \in E} \langle \sigma_i^z \sigma_j^z \rangle_{\mathcal{N}}$, red data in Fig. 5(a). We use COBYLA with $\theta = (\gamma_1, \gamma_2, \beta_1, \beta_2)$ initialized from a Trotterized quantum annealing schedule [52]. Each circuit is run with 4096 shots.

The error-mitigated energy reaches lower values than the non-error-mitigated energy; see Fig. 5(a). Both optimizations converge within 40 iterations; see Fig. 5(b). Furthermore, when the error-mitigated cost function is optimized, the corresponding non-error-mitigated cost function, dotted blue curve in Fig. 5(a), reaches lower values than a direct optimization of the non-error-mitigated cost function; red dashed curve in Fig. 5(a). This shows that with the error mitigation on, the optimizer finds better values of θ . To further illustrate this, we compute the energy distribution of sampled bitstrings. We compare the distribution of the cost function of each sampled bitstring of the initial and last values of θ labeled θ_0 and θ^* respectively. The sampling is done with the noisy simulator, Figs. 6(a) and 6(c), and a noiseless simulator, Figs. 6(b) and 6(d). Sampling from a noisy $|\theta^*\rangle$ produces a distribution that is nearly identical to the one obtained by sampling from a noisy $|\theta_0\rangle$, see Figs. 6(a) and 6(c). However, sampling from a noiseless $|\theta^*\rangle$ produces a better distribution than sampling from a noiseless $|\theta_0\rangle$, see Figs. 6(a) and 6(b). This suggests that the error mitigation helps find better values of θ despite the fact that we cannot see this by *sampling* bitstrings from noisy QAOA states.

Finally, we repeat these simulations 20 times. In each simulation, we train a FFNN and optimize θ . This produces different optimization results due to the randomness of the noise. The optimization is carried out twice, once on $E_{\mathcal{M}}$ and once on $E_{\mathcal{N}}$. After the optimization, we sample 4096 bitstrings from a noiseless simulation of the QAOA circuit with

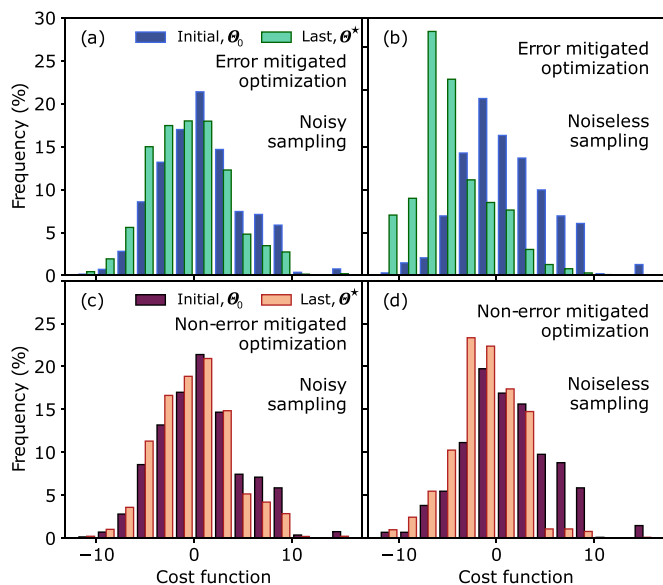


FIG. 6. Distribution of the cost function of each individual sampled bitstring for the optimization in Fig. 5. (a) and (b) correspond to the optimization with the error-mitigated cost function (solid lines in Fig. 5). (c) and (d) correspond to the optimization with the non-error-mitigated cost function (dashed lines in Fig. 5).

the optimized parameters θ^* . We compute the energy distribution of these bitstrings and report the expectation value. This expectation value is -3.78 ± 2.12 and -2.63 ± 2.01 when E_M and E_N is optimized, respectively. These results indicate that error mitigation tends to help the classical optimizer find better QAOA parameter values. However, we observe in seven out of 20 simulations that an optimization of the noisy QAOA cost function E_N produces better parameters, as measured by a noiseless sampling of bitstrings, than an optimization of the error mitigated cost function E_M . In all simulations, except one, the error mitigated cost function E_M has a lower energy than the non-error-mitigated cost function E_N .

B. Hardware

We now optimize the parameters γ_1 , γ_2 , β_1 , and β_2 with COBYLA of a depth-two QAOA for RR3 graphs on superconducting qubit hardware. As a cost function, we minimize the energy $E_M = \sum_{(i,j) \in E} \langle \sigma_i^z \sigma_j^z \rangle_M$ computed with error-mitigated correlators produced by FFNNs. Before each run, the FFNN is trained, as described in Sec. III, with 3000 training points evaluated with 1024 shots each.

For RR3 graphs with 30 nodes, the quantum circuits have a total of 610 ECR gates, 1297 X and \sqrt{X} gates, and 1577 R_Z gates, see Fig. 7. Graphs with 40 nodes have a total of 958 ECR gates. While these circuits are extremely deep and wide, they still contain a significant signal. When running the QAOA optimization on hardware we observe a minimization of E_M for all graphs, see dark purple curves in Figs. 8(a), 8(d), 8(g), and 8(j). The non-error-mitigated cost function E_N (light purple curves) also decreases. In all cases, the optimization of γ_i and β_i converges in about 20 to 40 iterations of COBYLA, see Figs. 8(b), 8(e), 8(h), and 8(k). We compare the distribution of the sampled bitstrings obtained from QAOA circuits with the initial points θ_0 to the distribution obtained with the optimized θ^* . We see an improvement in the distribution, i.e., a bias towards lower values, for all RR3 graphs—compare the dark blue and light teal curves in Figs. 8(c), 8(f), 8(i), and 8(l). This is consistent with the interpretation that there is a meaningful signal in the corresponding circuits. We report the mean μ of each distribution (vertical lines in Fig. 8) as an approximation ratio $\alpha(\mu) = (\mu - \langle H_C \rangle_{\max}) / (\langle H_C \rangle_{\min} - \langle H_C \rangle_{\max})$ contained in $[0, 1]$. The optimized parameters produce an $\alpha(\mu)$ of 71.6%, 64.0%, 59.6%, and 58.3% for the 10, 20, 30, and 40 node graphs.

We optimize the parameters θ with ECR-based circuits since all parameters are in virtual R_Z gates. This preserves the amplitude and duration of all pulses in the schedule thus facilitating noise mitigation. Pulse-efficient transpilation moves the parameters from the R_Z gates into the cross-resonance pulses [17,18]. This shortens the pulse schedule but changes its noise properties. For example, a pulse-efficient transpilation of an $R_{ZZ}(\theta)$ -SWAP pair, as shown in Fig. 3(a), reduces their duration by up to 20%, depending on θ . The shorter schedules produce better bitstrings than the fixed-duration schedules with parameters in R_Z gates. We run the pulse-efficient circuits for the last points θ^* . This results in an improved $\alpha(\mu)$ of 76.0%, 65.5%, 60.8%, and 58.6% over the same circuit without pulse-efficient transpilation for the 10, 20, 30, and 40 node graphs, respectively—compare the dash-dotted red line to the solid teal line in Figs. 8(c), 8(f), 8(i), and 8(l).

The best bitstrings sampled from the distributions shown in Figs. 8(c), 8(f), 8(i), and 8(l) have a cut value of 13 (MAXCUT), 26 (MAXCUT), 35 (0.833), and 44 (0.786) for the 10, 20, 30, and 40 node graphs, respectively. The numbers in parenthesis indicate the approximation ratio. These bitstrings were observed a total of 1619, 124, 474, and 988 times, respectively, out of the 3×4096 shots in the distributions.

To distinguish the impact of hardware noise on the bitstring distribution from limitations of the depth-two QAOA




FIG. 7. Depth-two QAOA circuit of a 30 node RR3 containing 610 ECR gates (dark blue), 1297 single-qubit X and \sqrt{X} gates (blue), and 1577 virtual R_Z gates (teal).

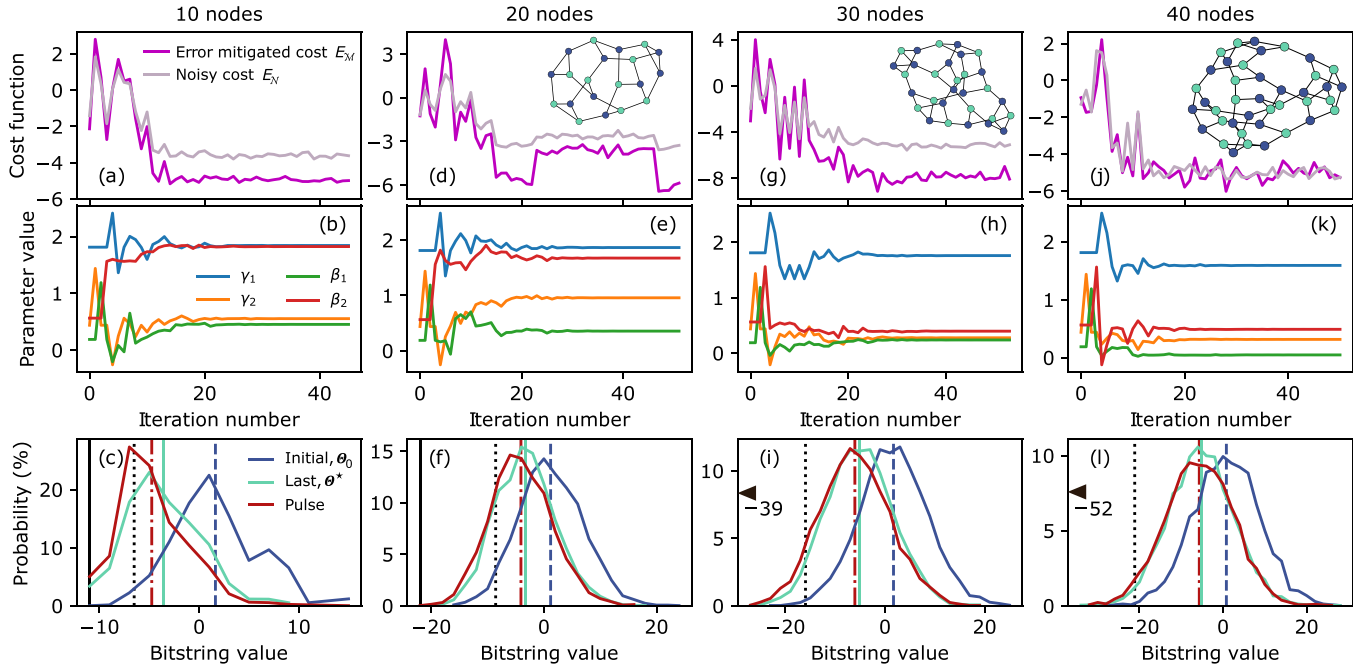


FIG. 8. Depth-two QAOA data acquired on *ibm_brisbane*. The graphs are shown as insets with node colors indicating the MAXCUT found with CPLEX [53]. (a), (d), (g), (j) Error-mitigated cost function $E_{\mathcal{M}}$ obtained from FFNNs at each iteration of COBYLA. The nonmitigated cost function $E_{\mathcal{N}}$, is also shown as $E_{\mathcal{M}}$ is optimized. $E_{\mathcal{N}}$ and $E_{\mathcal{M}}$ thus represent the same quantities as the blue dotted and blue solid lines in Fig. 5(a), respectively. (b), (e), (h), (k) QAOA parameters during the optimization. (c), (f), (i), (l) Distribution of the energy of the sampled bit strings. The dark blue and light teal lines correspond to bitstrings sampled from the QAOA circuits with the initial and last parameters θ_0 and θ^* , respectively. The red lines correspond to sampling from a pulse-efficient circuit. The solid black lines in (c) and (f) indicate the energy of the MAXCUT. In (i) and (l), this energy lies outside of the x -axis range. We indicate it as a triangle with the energy as a number. The black dotted vertical lines indicate the noiseless expectation value obtained with θ^* .

ansatz, we compute the noiseless expectation value of the cost Hamiltonian $H_C = \sum_{(i,j) \in E} \sigma_i^z \sigma_j^z$. We evaluate H_C at the last point θ^* obtained from the noisy hardware optimization. This computation is made fast, even for a 40-node graph, with quantum circuits based on the light cone of each correlator $\sigma_i^z \sigma_j^z$. This method is detailed in Appendix B. The noiseless expectation value is indicated as a dotted line in Figs. 8(c), 8(f), 8(i), and 8(l). The corresponding approximation ratios $\alpha(\mu)$ are 82.8%, 74.2%, 72.6%, and 72.4% for the 10, 20, 30, and 40 node graphs, respectively. These values show the potential improvement in the bitstring value distribution if hardware noise could be reduced.

V. DISCUSSION AND CONCLUSION

Many machine-learning tools can error mitigate an expectation value. The first contribution of this paper is a user-friendly FFNN-based error mitigation strategy with a problem-inspired methodology to generate the training data. Here, the FFNN is trained once before the variational optimization. It tries to match training data acquired at randomly sampled values of β for which all values of γ are set to zero with R_Z gates replaced by barriers. This preserves the circuit structure and makes it easy to simulate classically. We observe that the FFNN performs better on validation data than a linear regression as the circuit size is increased, see Appendix C. The data from the optimization with error mitigation show that the FFNN reduces the effect of noise on the cost function at unseen values of β and γ . Other data-generation approaches

are possible and could be investigated in future work, which may also explore other machine-learning tools such as random forests, as done in Ref. [54].

Our second contribution is to implement nonplanar RR3 graphs on hardware by leveraging the SAT mapping of Matsuo *et al.* [38] and swap networks [20,31]. We observe a meaningful signal for a depth-two QAOA with up to 40 nodes. The swap networks with 2, 4, 6, and 7 layers that we implement on 10, 20, 30, and 40 qubits can generate graph densities of up to 40%, 30%, 27%, and 22%, respectively. The corresponding circuits have impressive gate counts. We attribute hardware improvements of the Eagle quantum processors, shown in Fig. 9, to this success. For example, the T_1 times of *ibm_brisbane* are more than twice as large as those of *ibmq_mumbai*, see Fig. 9(a), which was used in the 27 qubit experiment in Ref. [20]. The cumulative distribution function of the gate error of *ibm_brisbane* and *ibmq_mumbai* is approximately the same, see Fig. 9(c). However, *ibm_brisbane* has 127 qubits while *ibmq_mumbai* only has 27 qubits. This allows us to run the 40-qubit RR3 graph on the *best* line of qubits. For example, the product of the fidelity of the 39 gates on the best line with 40 qubits on *ibm_brisbane* is 76.5%, see Appendix A. By contrast, the product of all 28 two-qubit gate fidelities on *ibmq_mumbai* only reaches 72.8% despite there being fewer gates.

Circuits with 40 nodes can be simulated classically. In particular, computing $\langle H_C \rangle$ of a depth- p QAOA for a RR3 graph produces an effective light cone with at most $\sum_{k=0}^p 2^{k+1}$ qubits in the circuits, i.e., 14 in our $p = 2$ case. This allows




FIG. 9. Properties of the Eagle quantum processors used in this paper, *ibmq_brisbane* and *ibmq_kyiv*, compared to the Falcon processor used in Ref. [20], i.e., *ibmq_mumbai*. The data are presented as cumulative distribution functions (CDFs).

us to confirm that optimizing the error-mitigated correlators produces good variational parameters. As quantum computers increase in size and quality, such a classical verification will no longer be possible. Crucially, our hardware demonstration is much larger than many current experiments which typically employ up to 20 qubits [55]. Furthermore, reproducing even depth-one QAOA samples is classically intractable [56]. Our paper is thus a step towards implementing QAOA on hardware that cannot be classically simulated. Future work must focus on implementing deeper circuits on hardware with more connectivity. Our work also serves as a benchmark to track quantum hardware progress, as done, e.g., with complete graphs [57]. We anticipate that hardware improvements, e.g., increasing T_1 times, and novel architectures, based on, e.g., tunable couplers [58,59], will enable larger simulations.

The depth-one QAOA results, shown in Sec. II, exhibit the parameter concentration already observed in the literature [60–64]. This may be important to quickly generate good yet suboptimal solutions to combinatorial optimization problems without having to optimize the variational parameters for each problem instance [20].

The objective of QAOA is to *sample* good or even optimal solutions from a quantum state $|\theta^*\rangle$ that minimizes the energy H_C , i.e., to find $x_{\text{opt}} = \text{argmin}_x H_C(x)$. The state is obtained by minimizing the *expectation* value of H_C . The error-mitigation method we present helps find good parameters γ and β . However, quantum approximate optimization needs tools that error mitigate *samples*. Noiseless simulations of $\langle H_C \rangle$ computed with the optimal parameters show that a significant gain is obtainable if samples could be error mitigated. Recently, Barron *et al.* showed that a noise dependent sampling overhead produces good solution samples [65]. This makes sense in the context of optimization as long as the total samples drawn is at least less than the 2^n cost of a brute-force search and ideally less than the sampling complexity of the best classical benchmark. Hardware improvements in, e.g., the fidelity of a layer of gates [66] and noise-aware transpilers [67,68], will reduce this sampling overhead. Furthermore, proper QAOA bench-

marks must compare to state-of-the-art solvers [34], such as Gurobi and CPLEX, and randomized rounding algorithms [69]. For example, for RR3 graphs there is an approximation algorithm that achieves an approximation ratio of 0.9326 [70]. Such benchmarking is an important task in itself, which our methods enable on hardware.

For variational algorithms, like the variational eigensolver applied in a chemistry setting [71,72], error-mitigating expectation values are often sufficient, e.g., to compute the energy spectrum of molecules [73,74]. The FFNN-based error mitigation method we present is directly transferable to such settings, which increases its applicability. Furthermore, the transpiler methodology we leverage works with any non-hardware-native block of commuting two-qubit gates. It thus applies to circuits other than QAOA such as graph states [75], and algorithms that implement $e^{-iH_C t}$ including Ising simulations [19].

The code for these results is available at Ref. [76].

ACKNOWLEDGMENTS

S.H.S. acknowledges support from the IBM Ph.D. fellowship 2022 in quantum computing. The authors also thank M. Serbyn, R. Kueng, R. A. Medina, and S. Woerner for fruitful discussions.

IBM, the IBM logo, and *ibm.com* are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available at Ref. [77].

APPENDIX A: RR3 GRAPH TRANSPILATION

Random regular graphs are sparse. When their corresponding QAOA circuit is transpiled to the hardware with predetermined swap layers, certain edges may require a large number of swap layers. By wisely choosing the initial mapping between decision variables and physical qubits, we reduce the number of swap layers needed. A SAT-based approach to this initial mapping problem was proposed by Matsuo *et al.* [38]. Here, the initial mapping problem is formulated as a SAT problem that is satisfiable if $e^{-i\gamma H_C}$ can be routed to hardware with ℓ swap layers. A binary search over ℓ finds the initial mapping that minimizes the number of layers. We label the minimum number of swap layers by ℓ^* .

1. Graph generation

We generate 100 RR3 graphs with n nodes for each $n \in \{10, 20, 30, 40\}$. Each graph is mapped to a line of n qubits with the SAT approach. The distribution of the number of swap layers at the different sizes n is shown in Table I. With the SAT mapping, there are graph instances with 10, 20, 30, and 40 nodes that can be implemented with 2, 4, 6, and 7 swap layers, respectively. This is a large reduction compared to a trivial mapping which typically requires $n - 2$ swap layers [38]. The experiments in the main text are done on graphs that require the smallest number of swap layers.

TABLE I. Number of RR3 graphs that required ℓ^* swap layers after an initial mapping found through solving SAT problems. At each graph size with n nodes, 100 graph instances were generated. For instance, out of 100 random instances of RR3 graphs with 40 nodes, three could be mapped to a line of qubits with seven swap layers.

Number of nodes	Number of swap layers ℓ^*								
	2	3	4	5	6	7	8	9	10
10	26	40	34						
20			13	55	32				
30					18	64	18		
40						3	17	67	13

2. Qubit selection

ibm_brisbane has 127 qubits, i.e., 87 more than the largest graph we study, see Fig. 10. We, therefore, select the best line of qubits to execute the quantum circuits on. For each pair of nodes i and j in the back end’s coupling map, we enumerate all paths of length $|V|$ connecting them. Next, we compute the path fidelity for each path p_k as $\prod_{(i,j) \in p_k} (1 - E_{\text{ECR},i,j})$ and select the best one. Here, $E_{\text{ECR},i,j}$ is the error of the ECR gate between qubits i and j . On *ibm_brisbane*, there are 1336, 15814, 125918, and 754462 lines of 10, 20, 30, and 40 qubits, respectively. The best measured respective path fidelities are 95.9%, 89.5%, 82.8%, and 76.5%.

APPENDIX B: LIGHT-CONE QAOA

For low-depth QAOA and sparse graphs, such as RR3, we can efficiently compute the expectation value $\langle H_C \rangle = \sum_{(i,j) \in E} \langle \sigma_i^z \sigma_j^z \rangle$ by considering the light cone of each correlator $\sigma_i^z \sigma_j^z$. Indeed, for depth-one QAOA, each correlator $\sigma_i^z \sigma_j^z$ is only impacted by the gates applied to nodes in the direct




FIG. 10. Coupling map of *ibm_brisbane* showing the qubit connectivity. The black qubits form the 40-qubit line with the best path fidelity as measured by the quality of the ECR gates.

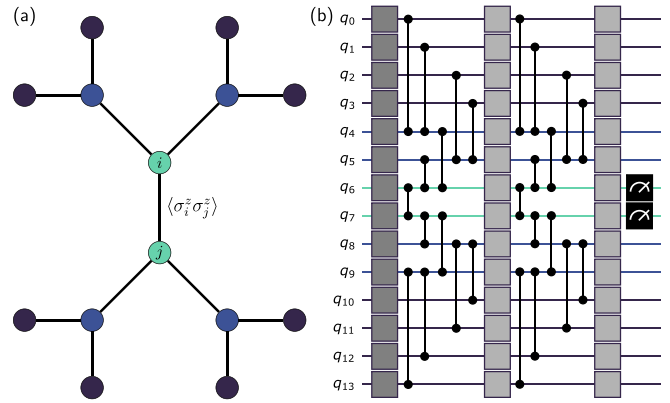


FIG. 11. (a) Subgraph of a RR3 graph with all nodes up to a distance two from the nodes i and j for which we want to compute the correlator $\langle \sigma_i^z \sigma_j^z \rangle$. (b) Quantum circuit for the graph in (a). The dumbbells represent parameterized R_{ZZ} rotations and the dark and light grey boxes are Hadamards and R_x gates, respectively. The wire colors in (b) correspond to the color of the nodes in the graph in (a).

neighborhood of i and j , i.e., distance-one nodes, see Fig. 11. For depth-two QAOA, we must consider all nodes that are at most at a distance of two away from i and j in E . Therefore, to compute $\langle H_C \rangle$ for depth-two QAOA, we create $|E|$ circuits each with at most 14 nodes. In the circuit corresponding to

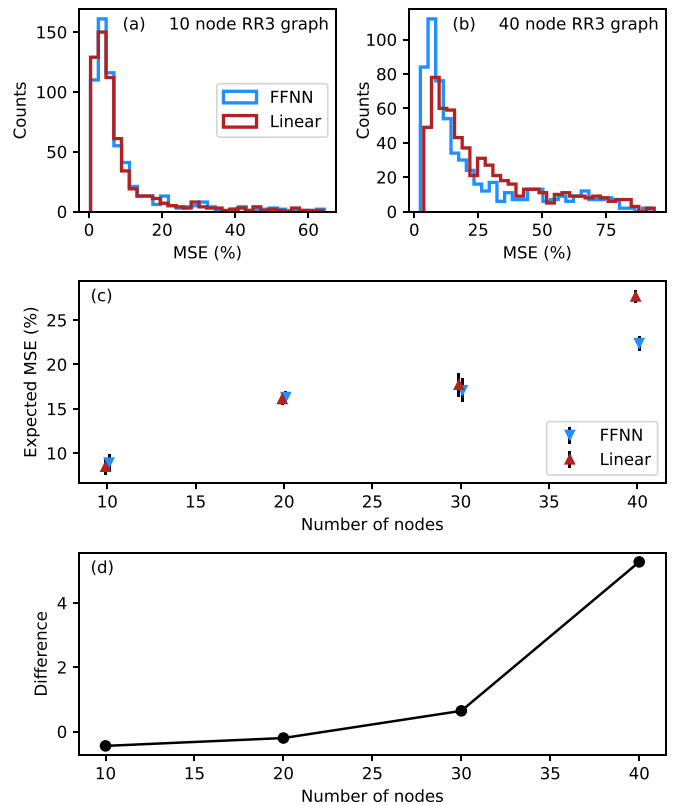


FIG. 12. Comparison of a FFNN and a linear regression. (a), (b) Distribution of the MSE of a validation data set for the ten and 40 node graphs, respectively. (c) MSE averaged over all ten validation data sets as a function of graph size. (d) Difference between the MSE of the FFNN and the linear regression.

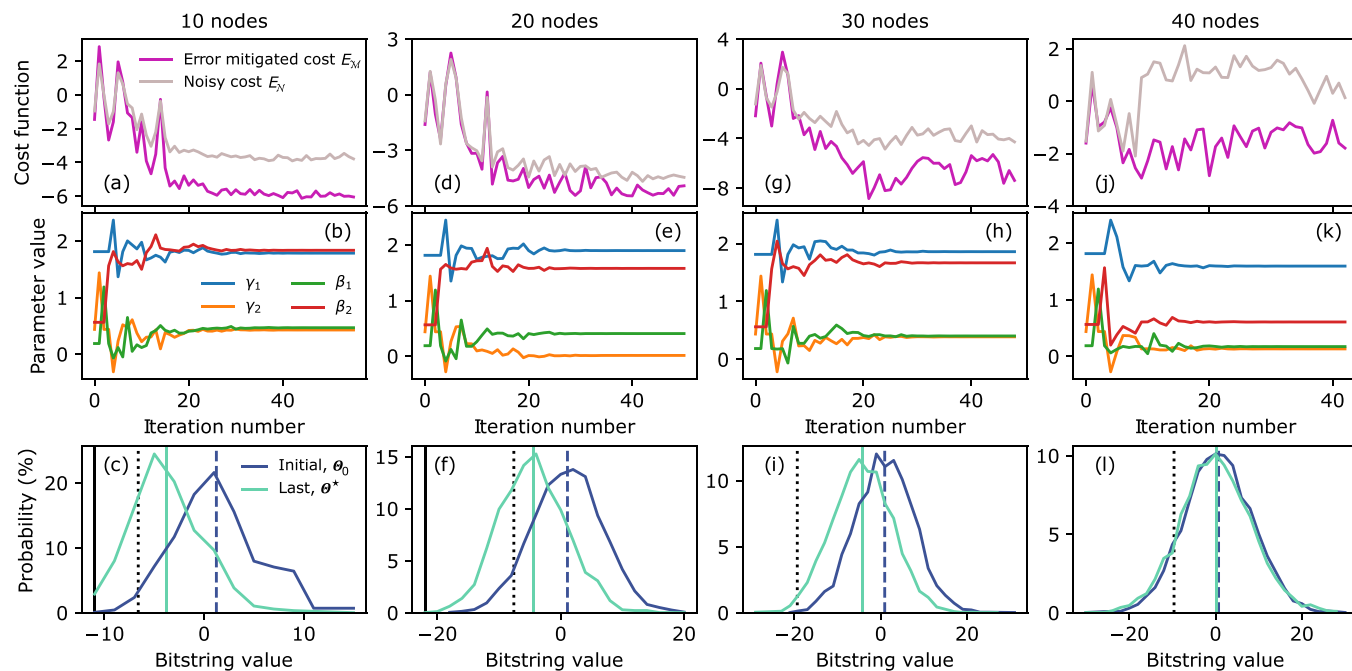


FIG. 13. Run on *ibm_kyiv*. The training data was comprised of 300 circuits and the FFNN had one hidden layer with 100 nodes in all instances. The underlying graphs and other displayed quantities are identical to those in Fig. 8 of the main text.

$\langle \sigma_i^z \sigma_j^z \rangle$, we only measure the qubits that map to nodes i and j , see Fig. 11(b).

APPENDIX C: MODEL COMPARISON

Here, we compare the FFNN to a linear regression. The comparison is done on the hardware-measured data presented

in Fig. 8 of the main text. We resample the 3000 training data points ten times to generate ten training data sets made of 80% of the data, i.e., 2400 sets of $|V|(|V| + 1)/2$ expectation values, and ten validation sets made of the remaining 20% of the data. We train both a FFNN and a linear model on the ten data sets with 80% of the data. Here, we employ the `MLPRegressor` and the `LinearRegression` from `sklearn`.

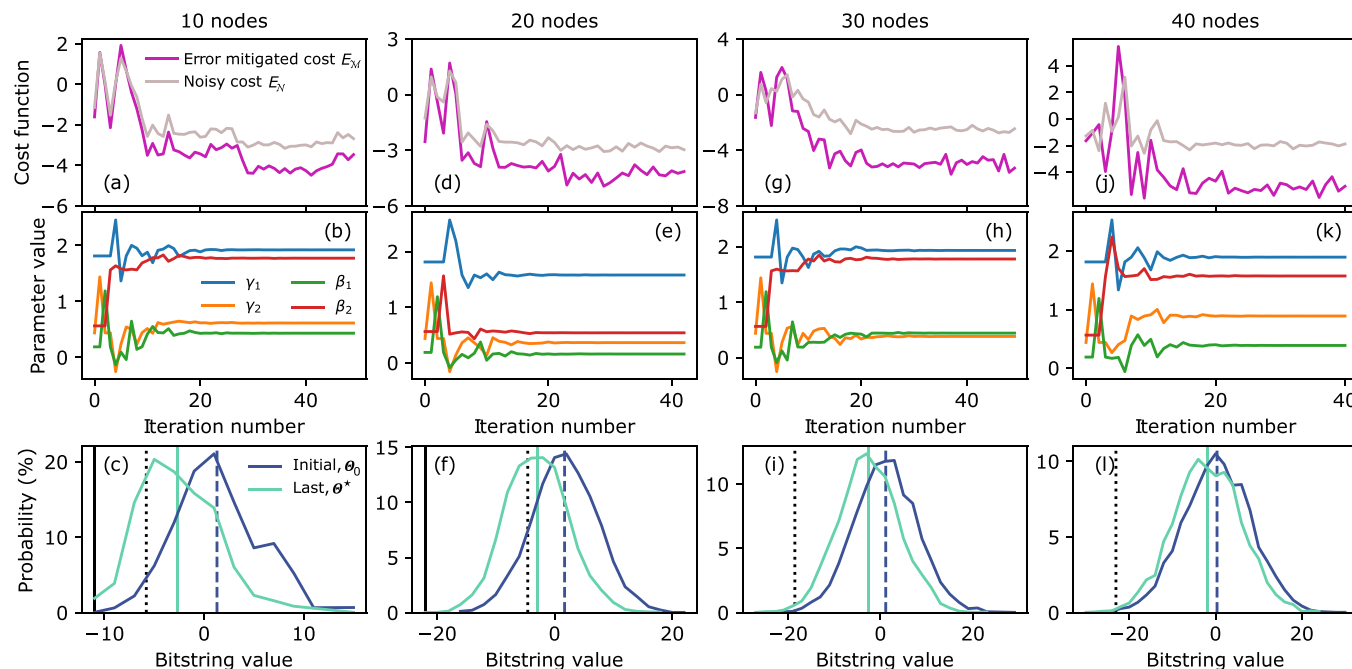


FIG. 14. Run on *ibm_nazca*. The training data were comprised of 3000 circuits and the FFNN had one hidden layer with a number of nodes made of the average of the input and output. The underlying graphs and other displayed quantities are identical to those in Fig. 8 of the main text.

Next, we compute the MSE between the $3|V|/2$ predicted ZZ correlators and their ideal value. This results in a distribution of MSEs across the 600 validation data points for each of the ten validation data sets. We observe an increase in the MSE as the graph size increases. Furthermore, the MSE of the linear model increases faster than the MSE of the FFNN as the graph size increases, as exemplified by comparing Figs. 12(a) and 12(b) for one of the ten validation sets. Next, we compute the average over all MSEs at each graph size and the associated standard deviation of this mean, see Fig. 12(c) and their difference in 12(d). As the number of nodes in the graph increase, the error of the FFNN becomes significantly lower than the error of the linear model. This test suggests that for the particular circuits that we employed, the FFNN performs better than the linear model as the size of the quantum circuit increases. This may be due to effects such as cross-talk and unitary gate errors, e.g., over- and underro-

tations, that are not captured by a linear model like the depolarizing channel [26].

APPENDIX D: ADDITIONAL HARDWARE RUNS

Here, we present depth-two QAOA data acquired on *ibm_kyiv* and *ibm_nazca* in addition to the data acquired on *ibm_brisbane*, see Figs. 13 and 14, respectively. The data acquired on *ibm_nazca* was gathered under the same settings as the data on *ibm_brisbane*. The data acquired on *ibm_kyiv* is produced with a smaller number of training circuits, i.e., 300 instead of 3000, and the hidden layer of the FFNN had 100 nodes for each graph size. By contrast, the FFNN trained for *ibm_brisbane* and *ibm_nazca* had a number of hidden neurons equal to the average of the input and output number of neurons.

-
- [1] A. Abbas, A. Ambainis, B. Augustino, A. Bärtshi, H. Buhrman, C. Coffrin, G. Cortiana, V. Dunjko, D. J. Egger, B. G. Elmegreen, N. Franco, F. Fratini, B. Fuller, J. Gacon, C. Gonciulea, S. Gribling, S. Gupta, S. Hadfield, R. Heese, G. Kircher *et al.*, Quantum optimization: Potential, challenges, and the path forward, [arXiv:2312.02279](https://arxiv.org/abs/2312.02279).
- [2] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [3] D. J. Egger, C. Gambella, J. Marecek, S. McFaddin, M. Mevissen, R. Raymond, A. Simonetto, S. Woerner, and E. Yndurain, Quantum computing for finance: State-of-the-art and future prospects, *IEEE Trans. Quantum Eng.* **1**, 1 (2020).
- [4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [5] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [6] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant, Quantum machine learning in high energy physics, *Mach. Learn.: Sci. Technol.* **2**, 011003 (2021).
- [7] A. D. Meglio, K. Jansen, I. Tavernelli, C. Alexandrou, S. Arunachalam, C. W. Bauer, K. Borrás, S. Carrazza, A. Crippa, V. Croft *et al.*, Quantum computing for high-energy physics: State of the art and challenges. summary of the QC4HEP working group, [arXiv:2307.03236](https://arxiv.org/abs/2307.03236).
- [8] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [9] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, *Phys. Rev. X* **8**, 031027 (2018).
- [10] Y. Li and S. C. Benjamin, Efficient variational quantum simulator incorporating active error minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [11] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature (London)* **567**, 491 (2019).
- [12] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, Digital zero noise extrapolation for quantum error mitigation, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE), Denver, CO, USA* (IEEE, Piscataway, NJ, 2020), pp. 306–316.
- [13] R. LaRose, A. Mari, S. Kaiser, P. J. Karalekas, A. A. Alves, P. Czarnik, M. E. Mandouh, M. H. Gordon, Y. Hindy, A. Robertson *et al.*, Mitiq: A software package for error mitigation on noisy quantum computers, *Quantum* **6**, 774 (2022).
- [14] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [15] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature (London)* **618**, 500 (2023).
- [16] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. C. McKay, Qiskit pulse: Programming quantum computers through the cloud with pulses, *Quantum Sci. Technol.* **5**, 044006 (2020).
- [17] J. P. T. Stenger, N. T. Bronn, D. J. Egger, and D. Pekker, Simulating the dynamics of braiding of Majorana zero modes using an IBM quantum computer, *Phys. Rev. Res.* **3**, 033171 (2021).
- [18] N. Earnest, C. Tornow, and D. J. Egger, Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware, *Phys. Rev. Res.* **3**, 043088 (2021).
- [19] A. Carrera Vazquez, D. J. Egger, D. Ochsner, and S. Woerner, Well-conditioned multi-product formulas for hardware-friendly Hamiltonian simulation, *Quantum* **7**, 1067 (2023).
- [20] 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, *Quantum* **6**, 870 (2022).
- [21] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, From pulses to

- circuits and back again: A quantum optimal control perspective on variational quantum algorithms, *PRX Quantum* **2**, 010101 (2021).
- [22] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, Gate-free state preparation for fast variational quantum eigensolver simulations, *npj Quantum Info.* **7**, 155 (2021).
- [23] D. J. Egger, C. Capecchi, B. Pokharel, P. K. Barkoutsos, L. E. Fischer, L. Guidoni, and I. Tavernelli, Pulse variational quantum eigensolver on cross-resonance-based hardware, *Phys. Rev. Res.* **5**, 033159 (2023).
- [24] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [25] C. Kim, K. D. Park, and J.-K. Rhee, Quantum error mitigation with artificial neural network, *IEEE Access* **8**, 188853 (2020).
- [26] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, Error mitigation with Clifford quantum-circuit data, *Quantum* **5**, 592 (2021).
- [27] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, Learning-based quantum error mitigation, *PRX Quantum* **2**, 040330 (2021).
- [28] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [29] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
- [30] H. Markowitz, Portfolio selection, *J. Finance* **7**, 77 (1952).
- [31] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [32] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and subsystem codes on low-degree graphs with flag qubits, *Phys. Rev. X* **10**, 011022 (2020).
- [33] E. Pelofske, A. Bärttschi, and S. Eidenbenz, Quantum annealing vs. QAOA: 127 qubit higher-order ising problems on NISQ computers, in *High Performance Computing*, edited by A. Bhatele, J. Hammond, M. Baboulin, and C. Kruse (Springer Nature Switzerland, Cham, 2023), pp. 240–258.
- [34] D. Rehfeldt, T. Koch, and Y. Shinano, Faster exact solution of sparse MaxCut and QUBO problems, *Math. Prog. Comp.* **15**, 445 (2023).
- [35] F. Pan, K. Chen, and P. Zhang, Solving the sampling problem of the Sycamore quantum circuits, *Phys. Rev. Lett.* **129**, 090502 (2022).
- [36] A. Kissinger, J. van de Wetering, and R. Vilmart, Classical simulation of quantum circuits with partial and graphical stabiliser decompositions, in *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, Leibniz International Proceedings in Informatics (LIPIcs) (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022), Vol. 232, pp. 51–513.
- [37] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, Tensor network quantum simulator with step-dependent parallelization, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Broomfield, CO, USA (IEEE, Piscataway, NJ, 2022), pp. 582–593.
- [38] A. Matsuo, S. Yamashita, and D. J. Egger, A SAT approach to the initial mapping problem in swap gate insertion for commuting gates, *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* **E106.A**, 1424 (2023).
- [39] M. Streif, S. Yarkoni, A. Skolik, F. Neukart, and M. Leib, Beating classical heuristics for the binary paint shop problem with the quantum approximate optimization algorithm, *Phys. Rev. A* **104**, 012403 (2021).
- [40] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case, [arXiv:2004.09002](https://arxiv.org/abs/2004.09002).
- [41] K. Marwaha and S. Hadfield, Bounds on approximating Max k XOR with quantum and classical local algorithms, *Quantum* **6**, 757 (2022).
- [42] D. S. França and R. García-Patrón, Limitations of optimization algorithms on noisy quantum devices, *Nat. Phys.* **17**, 1221 (2021).
- [43] L. Lao and D. E. Browne, 2QAN: A quantum compiler for 2-local qubit hamiltonian simulation algorithms, in *Proceedings of the 49th Annual International Symposium on Computer Architecture, ISCA '22* (Association for Computing Machinery, New York, 2022), pp. 351–365.
- [44] M. Alam, A. Ash-Saki, and S. Ghosh, Circuit compilation methodologies for quantum approximate optimization algorithm, in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)* (IEEE, Athens, Greece, 2020), pp. 215–228.
- [45] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, *Phys. Rev. B* **81**, 134507 (2010).
- [46] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, *Phys. Rev. A* **93**, 060302(R) (2016).
- [47] L. Viola, E. Knill, and S. Lloyd, Dynamical decoupling of open quantum systems, *Phys. Rev. Lett.* **82**, 2417 (1999).
- [48] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [49] C. M. Bishop and N. M. Nasrabadi, Pattern recognition and machine learning, *J. Electron. Imaging* **16**, 049901 (2007).
- [50] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient Z gates for quantum computing, *Phys. Rev. A* **96**, 022330 (2017).
- [51] The FFNN is implemented with the MLPRegressor from sklearn.
- [52] S. H. Sack and M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *Quantum* **5**, 491 (2021).
- [53] IBM ILOG CPLEX Optimizer.
- [54] H. Liao, D. S. Wang, I. Sitdikov, C. Salcedo, A. Seif, and Z. K. Mineev, Machine learning for practical quantum error mitigation, [arXiv:2309.17368](https://arxiv.org/abs/2309.17368).
- [55] T. Ichikawa, H. Hakoshima, K. Inui, K. Ito, R. Matsuda, K. Mitarai, K. Miyamoto, W. Mizukami, K. Mizuta, T. Mori, Y. Nakano, A. Nakayama, K. N. Okada, T. Sugimoto, S. Takahira, N. Takemori, S. Tsukano, H. Ueda, R. Watanabe, Y. Yoshida *et al.*, A comprehensive survey on quantum computer usage: How many qubits are employed for what purposes? [arXiv:2307.16130](https://arxiv.org/abs/2307.16130).

- [56] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, [arXiv:1602.07674](https://arxiv.org/abs/1602.07674).
- [57] G. C. Santra, F. Jendrzejewski, P. Hauke, and D. J. Egger, Squeezing and quantum approximate optimization, *Phys. Rev. A* **109**, 012413 (2024).
- [58] M. Ganzhorn, G. Salis, D. J. Egger, A. Fuhrer, M. Mergenthaler, C. Müller, P. Müller, S. Paredes, M. Pechal, M. Werninghaus, and S. Filipp, Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform, *Phys. Rev. Res.* **2**, 033447 (2020).
- [59] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim *et al.*, Realization of high-fidelity CZ and ZZ-free iSWAP gates with a tunable coupler, *Phys. Rev. X* **11**, 021058 (2021).
- [60] F. G. S. L. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances, [arXiv:1812.04170](https://arxiv.org/abs/1812.04170).
- [61] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte, Parameter concentrations in quantum approximate optimization, *Phys. Rev. A* **104**, L010401 (2021).
- [62] A. Galda, X. Liu, D. Lykov, Y. Alexeev, and I. Safro, Transferability of optimal QAOA parameters between random graphs, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE), Broomfield, CO, USA* (IEEE, Piscataway, NJ, 2021), pp. 171–180.
- [63] M. Streif and M. Leib, Training the quantum approximate optimization algorithm without access to a quantum processing unit, *Quantum Sci. Technol.* **5**, 034008 (2020).
- [64] R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, Parameter transfer for quantum approximate optimization of weighted MaxCut, *ACM Trans. Quantum Comput.* **4**, 1 (2023).
- [65] S. V. Barron, D. J. Egger, E. Pelofske, A. Bärttschi, S. Eidenbenz, M. Lehmkuehler, and S. Woerner, Provable bounds for noise-free expectation values computed from noisy samples, [arXiv:2312.00733](https://arxiv.org/abs/2312.00733).
- [66] D. C. McKay, I. Hincks, E. J. Pritchett, M. Carroll, L. C. G. Govia, and S. T. Merkel, Benchmarking quantum processor performance at scale, [arXiv:2311.05933](https://arxiv.org/abs/2311.05933).
- [67] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers, in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19* (Association for Computing Machinery, New York, 2019), pp. 1015–1029.
- [68] F. Wagner, D. J. Egger, and F. Liers, Optimized noise suppression for quantum circuits, [arXiv:2401.06423](https://arxiv.org/abs/2401.06423).
- [69] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [70] E. Halperin, D. Livnat, and U. Zwick, MAX CUT in cubic graphs, *J. Algorithms* **53**, 169 (2004).
- [71] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [72] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta *et al.*, Quantum optimization using variational algorithms on near-term quantum devices, *Quantum Sci. Technol.* **3**, 030503 (2018).
- [73] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner *et al.*, Gate-efficient simulation of molecular eigenstates on a quantum computer, *Phys. Rev. Appl.* **11**, 044092 (2019).
- [74] P. J. Ollitrault, A. Kandala, C.-F. Chen, P. K. Barkoutsos, A. Mezzacapo, M. Pistoia, S. Sheldon, S. Woerner, J. M. Gambetta, and I. Tavernelli, Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor, *Phys. Rev. Res.* **2**, 043140 (2020).
- [75] G. J. Mooney, G. A. L. White, C. D. Hill, and L. C. L. Hollenberg, Whole-device entanglement in a 65-qubit superconducting quantum computer, *Adv. Quantum Technol.* **4**, 2100061 (2021).
- [76] https://github.com/eggerdj/large_scale_qaoa.
- [77] <https://www.ibm.com/legal/copytrade>.