Feedback-Based Quantum Optimization

Alicia B. Magann, Kenneth M. Rudinger, Matthew D. Grace, and Mohan Sarovar

Phys. Rev. Lett. **129**, 250502 — Published 13 December 2022

DOI: 10.1103/PhysRevLett.129.250502

Feedback-based quantum optimization

Alicia B. Magann,[1, 2, 3] Kenneth M. Rudinger,[2] Matthew D. Grace,[1] and Mohan Sarovar[1]

[1] *Extreme-Scale Data Science & Analytics, Sandia National Laboratories, Livermore, California 94550, USA*

[2] *Center for Computing Research, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*

[3] *Department of Chemical & Biological Engineering, Princeton University, Princeton, New Jersey 08544, USA*

(Dated: September 21, 2022)

Quantum computers are expected to offer advantages over classical computers for combinatorial optimization. Here, we introduce a feedback-based strategy for quantum optimization, where the results of qubit measurements are used to constructively assign values to quantum circuit parameters. We show that this procedure results in an estimate of the combinatorial optimization problem solution that improves monotonically with the depth of the quantum circuit. Importantly, the measurement-based feedback enables approximate solutions to the combinatorial optimization problem without the need for any classical optimization effort, as would be required for the quantum approximate optimization algorithm (QAOA). We experimentally demonstrate this feedback-based protocol on a superconducting quantum processor for the graph-partitioning problem MaxCut, and present a series of numerical analyses that further investigate the protocol's performance.

Introduction.— Combinatorial optimization has broad and high-value applications in many sectors of industry and science, including for optimization of logistics and supply chain, and drug discovery [1]. Solving general combinatorial optimization problems is NP hard and most practical strategies involve developing good quality approximate solutions. Recently, there has been much interest in approximate solution of combinatorial optimziation problems through mapping to quantum systems, whereby the problem is encoded into an Ising Hamiltonian H_p [2], such that the solution of problem is encoded in the ground state of H_p. Then methods such as quantum annealing [3], or within the quantum circuit model, the quantum approximate optimization algorithm (QAOA) [4], are used to approximately prepare the ground state of H_p. Although there is no rigorous proof of an advantage to using such quantum techniques over classical approximation algorithms, it is widely believed that at some scale of problem such an advantage should exist.

We introduce a new approach to solving combinatorial optimization problems using quantum computers that operates through the use of parameterized quantum circuits and feedback, that is conditioned on qubit measurements at every quantum circuit layer, in order to determine the circuit parameter values at subsequent layers. This Feedback-based ALgorithm for Quantum OptimizatioN (FALQON) makes a direct connection to quantum Lyapunov control (QLC), a control strategy that uses feedback to identify the controls to drive the dynamics of a quantum system in a desired manner [5–13]. Our approach works within the framework of circuit-model quantum computing, but avoids a critical challenge facing the scaling of QAOA, which is the difficulty of optimizing a large number of variational parameters. In fact, it was recently shown that under certain assumptions, this classical optimization problem is itself NP-hard for QAOA [14]. Our feedback-based approach circumvents the need for optimization of variational parameters by using information from iterative measurements.

In the following, we show that FALQON produces a monotonically improving estimate of the combinatorial optimization problem solution, with respect to the depth of the circuit. We then consider the application of FALQON towards solving the MaxCut problem, and present the results of an experimental demonstration on quantum hardware. This is followed by a series of numerical analyses that explore the performance of FALQON for MaxCut on 3-regular graphs. Finally, we examine the required number of repeated circuit evaluations and compare this to the requirements of QAOA in this context. We conclude with a discussion of the tradeoffs between FALQON and QAOA, outline the additional content in our companion paper [15], and look to the future.

Feedback-based algorithm for quantum optimization.— We begin by considering a quantum system whose dynamics are governed by $i\frac{d}{dt}|\psi(t)\rangle = (H_\mathrm{p} + H_\mathrm{d}\beta(t))|\psi(t)\rangle$, where $|\psi(t)\rangle$ is the system state vector, we have set $\hbar = 1$, and H_p and H_d denote the (unitless) "drift" and "control" Hamiltonians, where the latter couples a scalar, time-dependent control function $\beta(t)$ to the system. We seek to minimize $\langle H_\mathrm{p}\rangle = \langle\psi(t)|H_\mathrm{p}|\psi(t)\rangle$ [16], and accomplish this by designing $\beta(t)$ such that

$$\frac{d}{dt}\langle\psi(t)|H_\mathrm{p}|\psi(t)\rangle(t) \leq 0, \quad \forall t \geq 0\,. \tag{1}$$

Evaluating the left-hand-side of Eq. (1), we see that $\frac{d}{dt}\langle\psi(t)|H_\mathrm{p}|\psi(t)\rangle = A(t)\beta(t)$, where $A(t) \equiv \langle\psi(t)|i[H_\mathrm{d}, H_\mathrm{p}]|\psi(t)\rangle$. There is significant flexibility in choosing $\beta(t)$ in order to satisfy Eq. (1), i.e., we may take $\beta(t) = -w\, f(t, A(t))$, for $w > 0$, where $f(t, A(t))$ is any continuous function with $f(t, 0) = 0$ and $A(t)f(t, A(t)) > 0$ for all $A(t) \neq 0$ [17]. Here, we present results for $w = 1$ and $f(t, A(t)) = A(t)$, such that $\beta(t) = -A(t)$. In prac-

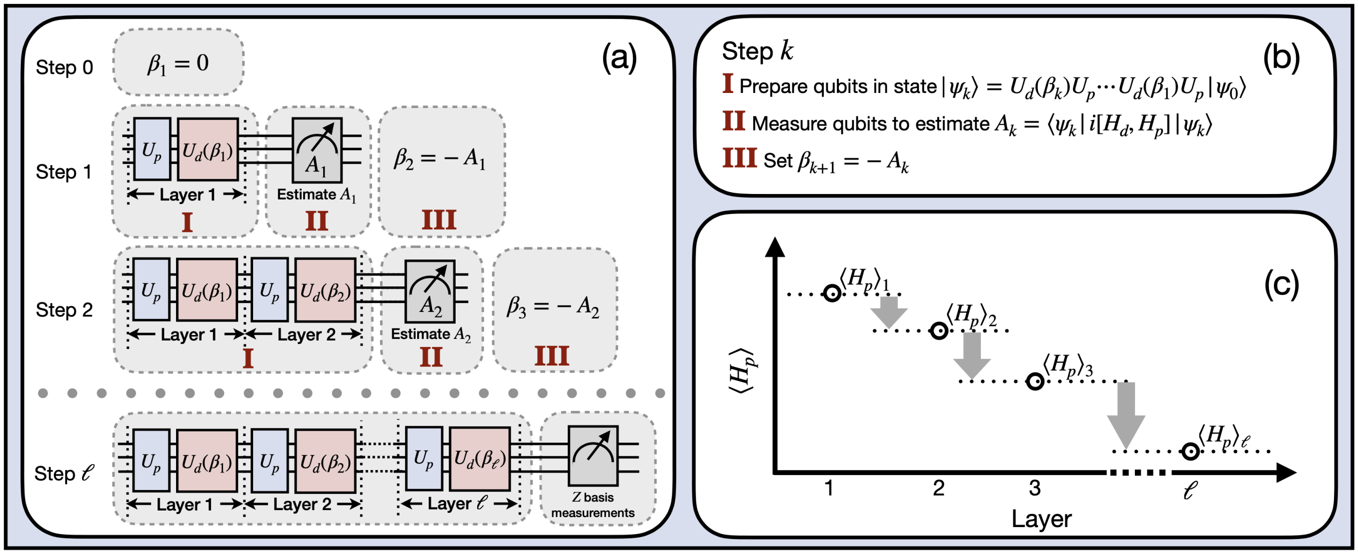

Figure 1. (a) The procedure for implementing FALQON. The initial step is to seed the procedure by setting $\beta_1 = 0$. The qubits are then initialized in the state $|\psi_0\rangle$, and a single FALQON layer is implemented to prepare $|\psi_1\rangle = U_\mathrm{d}(\beta_1)U_\mathrm{p}|\psi_0\rangle$. The qubits are then measured to estimate A_1, whose result is fed back to set $\beta_2 = -A_1$, up to sampling error. For subsequent steps $k = 2, \cdots, \ell$, the same procedure is repeated, as shown in (b): the qubits are initialized as $|\psi_0\rangle$, after which k layers are applied to obtain $|\psi_k\rangle = U_\mathrm{d}(\beta_k)U_\mathrm{p} \cdots U_\mathrm{d}(\beta_1)U_\mathrm{p}|\psi_0\rangle$, and then the qubits are measured to estimate A_k, and the result is fed back to set the value of β_{k+1}. This procedure causes $\langle H_\mathrm{p}\rangle$ to decrease layer-by-layer as per $\langle\psi_1|H_\mathrm{p}|\psi_1\rangle \geq \langle\psi_2|H_\mathrm{p}|\psi_2\rangle \geq \cdots \geq \langle\psi_\ell|H_\mathrm{p}|\psi_\ell\rangle$, as shown in (c), such that the quality of the solution to the combinatorial optimization problem monotonically improves with circuit depth. The protocol can be terminated when the value of $\langle H_\mathrm{p}\rangle$ converges or a threshold number of layers ℓ is reached. Then, after the final step, Z basis measurements on $|\psi_\ell\rangle$ can be used to determine a best candidate solution to the combinatorial optimization problem of interest, by repeatedly sampling from the probability distribution over bit strings induced by $|\psi_\ell\rangle$ and selecting the outcome associated with the best solution.

tice, we assign values to $\beta(t)$ as a feedback loop, where $\beta(t) = -A(t-\tau)$, and τ is a feedback loop time delay.

We now consider alternating, rather than concurrent, applications of H_p and H_d, leading to a time evolution of the form $U = U_\mathrm{d}(\beta_\ell)U_\mathrm{p} \cdots U_\mathrm{d}(\beta_1)U_\mathrm{p}$, where $U_\mathrm{p} = e^{-iH_\mathrm{p}\Delta t}$, $U_\mathrm{d}(\beta_k) = e^{-i\beta_k H_\mathrm{d}\Delta t}$, and $\beta_k = \beta(k\tau - \Delta t)$ for $k = 1, 2, \cdots, \ell$ and $\tau = 2\Delta t$, such that after each period of Δt the applied Hamiltonian alternates between H_p and H_d. We note that for small Δt, this yields a Trotterized approximation to the continuous time evolution of the system. In this Trotterized framework, we again aim to satisfy Eq. (1) by suitably choosing each value of β_k. We note that during the time intervals when H_p is applied, $\frac{d}{dt}\langle H_\mathrm{p}\rangle(t) = 0$; although its value doesn't change, the eigenstates of H_p do accumulate phases during this time, which impact the ensuing dynamics. Meanwhile, during the time intervals when H_d is applied, we recover the same result that $\frac{d}{dt}\langle H_\mathrm{p}\rangle = A(t)\beta(t)$. Consequently, we can ensure that Eq. (1) is satisfied by utilizing the same feedback law, given by $\beta_{k+1} = -A_k$, where $A_k = \langle\psi_k|i[H_\mathrm{d}, H_\mathrm{p}]|\psi_k\rangle$ [18]. In this setting, it is always possible to select Δt small enough such that Eq. (1) is satisfied [15]. However, if Δt is chosen to be too large, Eq. (1) will be violated. Based on this framework, the FALQON algorithm is presented in Fig. 1. The key feature of FALQON is that it is a constructive, optimization-free procedure for assigning values to each β_k according to a feedback law. And by design, the enforcement of Eq. (1) ensures that the quality of the solution to the combinatorial optimization problem under consideration (quantified by $\langle H_p\rangle$) improves monotonically with respect to the depth of the circuit, k.

The circuits used in QAOA have the same alternating structure as those in FALQON, albeit with additional parameters $\gamma_1, \cdots, \gamma_\ell$ that enter into U_p, such that $U_{QAOA} = U_\mathrm{d}(\beta_\ell)U_\mathrm{p}(\gamma_\ell) \cdots U_\mathrm{d}(\beta_1)U_\mathrm{p}(\gamma_1)$. Then, the solution to the original combinatorial optimization problem is sought by minimizing $\langle\psi(\vec{\gamma}, \vec{\beta})|H_\mathrm{p}|\psi(\vec{\gamma}, \vec{\beta})\rangle$ over the set of 2ℓ circuit parameters $\vec{\gamma} = (\gamma_1, \cdots, \gamma_\ell)$ and $\vec{\beta} = (\beta_1, \cdots, \beta_\ell)$ using a classical processor, where $|\psi(\vec{\gamma}, \vec{\beta})\rangle = U_{QAOA}|\psi_0\rangle$. However, we emphasize that FALQON is conceptually distinct from QAOA. Namely, QAOA seeks to minimize $\langle H_\mathrm{p}\rangle$ by classically optimizing over all parameters $\vec{\gamma}, \vec{\beta}$ simultaneously, while FALQON seeks to minimize $\langle H_\mathrm{p}\rangle$ over a sequence of quantum circuit layers, guided by qubit measurement-based feedback, *without classical optimization.*

Applications to MaxCut.— We now consider the application of FALQON towards a quintessential combinatorial optimization problem: MaxCut, which aims to identify a graph partition that maximizes the number of edges in a graph that are cut. For an unweighted graph $\mathcal{G}$, with n nodes and edge set $\mathcal{E}$, the MaxCut problem Hamiltonian is defined on n qubits as $H_\mathrm{p} = -\sum_{i,j\in\mathcal{E}} \frac{1}{2}\big(1 - Z_i Z_j\big)$, while H_d has the standard form $H_\mathrm{d} = \sum_{j=1}^{n} X_j$, such that $i[H_\mathrm{d}, H_\mathrm{p}] = \sum_{i,j\in\mathcal{E}} Y_i Z_j + Z_i Y_j$, where X_j, Y_j, and Z_j denote the Pauli operators acting on qubit j. As such, evaluating the feedback law $\beta_{k+1} = -A_k = -\langle\psi_k|i[H_\mathrm{d}, H_\mathrm{p}]|\psi_k\rangle$ in this setting involves measurements of maximally $n(n-1)$ two-qubit Pauli strings.

As a proof-of-principle, in Fig. 2 we present the results of an experimental demonstration of FALQON on a superconducting quantum processor for a simple instance of the MaxCut problem. In particular, we considered an instance of MaxCut on an unweighted graph composed of $n = 3$ nodes connected by two edges, such that $H_p = -\frac{1}{2}(2 - Z_1Z_2 - Z_2Z_3)$ and $i[H_\mathrm{d}, H_\mathrm{p}] = Y_1Z_2 + Z_2Y_1 + Y_2Z_3 + Z_3Y_2$. The experiment was performed on the publicly accessible `ibmq_manila` processor and utilized three qubits with nearest-neighbor connectivity matching that of the graph under consideration. In this setting, $\ell = 10$ steps of FALQON were performed according to the procedure outlined in Fig. 1, selecting $\Delta t = 0.2$. At each step, one circuit was implemented in order to estimate $\langle H_p\rangle_k$ natively in the computational basis. Two additional circuits were implemented in order to estimate the terms in A_k. For each circuit, the qubits were initialized in the ground state of H_d, and $m = 1024$ shots were taken.

As shown in Fig. 2(a), FALQON was successful in achieving a monotonic decrease of $\langle H_p\rangle$ in this experiment up to layer five (orange point markers). FALQON also achieves a monotonic increase in the success probability of measuring the two degenerate ground states, denoted by ϕ, as shown in Fig. 2(b). The error bars in Fig. 2(a) and (b) present the standard error of the mean, which estimates how much the reported $\langle H_p\rangle_k$ and ϕ_k may deviate from their true values due to finite sampling. Finally, the associated values of β, determined according to the feedback law $\beta_{k+1} = -A_k$, are plotted in Fig. 2(b).

Past layer 5, it is evident that FALQON is no longer able to decrement $\langle H_p\rangle$ using this hardware platform, despite exhibiting a continued monotonic decrease in associated noise-free numerical simulations (blue point markers). This reveals the limitations that hardware noise presents for this problem instance. Looking ahead, we are optimistic that continuous improvements to quantum hardware will pave the way towards applications of FALQON to increasingly complex combinatorial optimization problems.

In the interim, we explore how FALQON performs on larger instances of MaxCut through a series of noise-free numerical illustrations. These illustrations consider unweighted, connected 3-regular graphs with $n \in \{8, 10, \cdots, 20\}$ vertices. For $n \in \{8, 10\}$ we consider all nonisomorphic graphs; for $n \in \{12, 14, \cdots, 20\}$ we consider 50 randomly-generated, nonisomorphic graphs. In

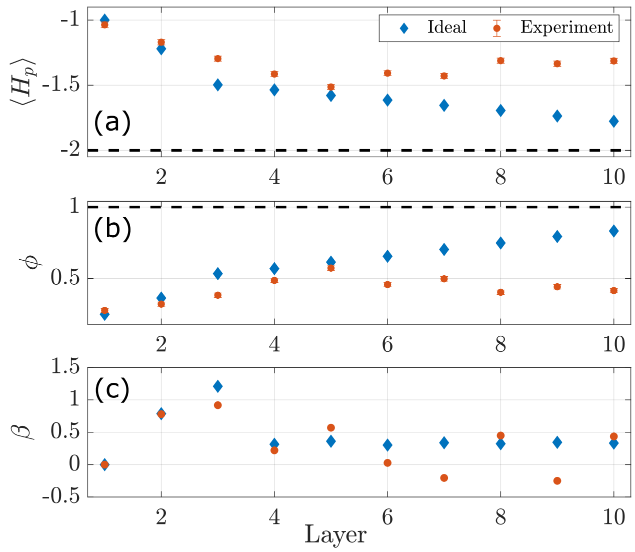

Figure 2. Results from experimental implementation of $\ell = 10$ layers of FALQON on a superconducting quantum processor. For this demonstration, FALQON is applied to an $n = 3$ qubit instance of MaxCut on an unweighted graph. Panel (a) shows that FALQON is successful in achieving a monotonic decrease of $\langle H_p\rangle$ over layers $k = 1, \cdots, 5$ in this experiment (orange point markers), noting that the global minimum value for this problem instance is $\langle H_p\rangle_\mathrm{min} = -2$ (dashed black line). In addition, in panel (b) a monotonic increase of the probability, ϕ, of measuring the two degenerate ground states is also observed up to layer $k = 5$ (orange point markers). The error bars in (a) and (b) indicate the standard error. The values of β are plotted in (c). In (a)-(c), the blue point markers correspond to ideal results computed numerically.

our simulations, the qubits are initialized in the ground state of H_d, and the performance of FALQON is quantified using the mean and standard deviations (over the problem instances) of two figures of merit: the approximation ratio, $r_\mathrm{A} = \langle H_\mathrm{p}\rangle/\langle H_\mathrm{p}\rangle_\mathrm{min}$ and the success probability of measuring the (potentially degenerate) ground state(s) $\{|q_{0,i}\rangle\}$, $\phi = \sum_i |\langle\psi|q_{0,i}\rangle|^2$. We relate the performance to two reference values: $r_\mathrm{A} = 0.932$, corresponding to the highest approximation ratio that can currently be guaranteed using a classical approximation algorithm (i.e., the algorithm of Goemans and Williamson [19]), and $\phi = 0.25$, which implies that on average, four repetitions will be needed in order to obtain a sample corresponding to the ground state. Our only free parameter is the time step Δt, which is tuned to be as large as possible, a value we call the *critical* Δt and denote by Δt_c, as long as the condition in Eq. (1) is met for all problem instances considered. Our results are collected in Fig. 3. In Fig. 3(b), the mean values of $\beta_1, \beta_2, \cdots$ are plotted as a function of layer for different values of n, according to the legend in Fig. 3(c), with the shading representing the standard deviation. We find that with increasing n, the shape of the resultant β curves follows a clear trend,

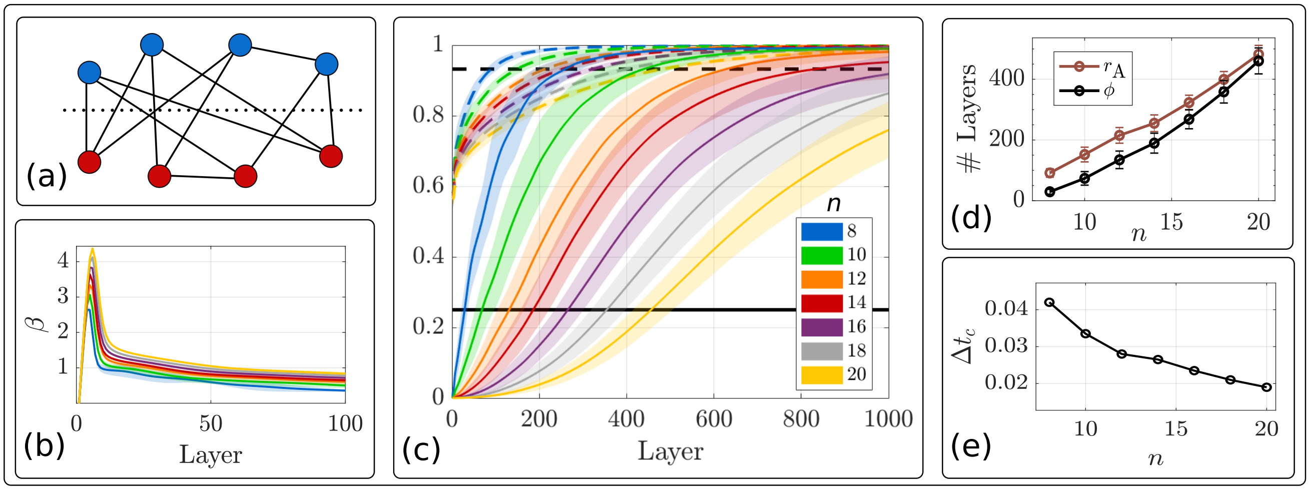

Figure 3. (a) Pictorial representation of MaxCut on a 3-regular graph with 8 vertices. (b) Mean β values are plotted as a function of layer for different n values, with shading showing the standard deviations. (c) The performance of FALQON, as quantified by the approximation ratio (dashed curves) and the success probability of measuring the degenerate ground state (solid curves) is shown for different values of n. (d) The mean number of layers needed to achieve the reference values of $r_{\mathrm{A}} = 0.932$ (dashed curve) and $\phi = 0.25$ (solid curve) is shown; error bars report the associated standard deviation. (e) The critical Δt values for different problem sizes are plotted.

and the standard deviation decreases. In Fig. 3(c), the associated r_{A} and ϕ results are shown (dashed and solid curves, respectively), and the associated reference values are plotted in black. For the cases considered here, we find that FALQON consistently leads to monotonic convergence towards very high r_{A} and ϕ values as a function of layer. To determine how the requisite circuit depths scale with the problem size, in Fig. 3(d) we plot the average number of layers required to achieve the reference values of r_{A} and ϕ as a function of n. Finally, in Fig. 3(e) we plot Δt_c for each value of n under consideration. The scaling of the required number of layers and Δt_c seems nearly linear, even up to $n = 20$, indicating a favorable runtime scaling of the FALQON algorithm, at least for this class of MaxCut problems. We remark that in addition to the analyses presented here, we also tested the performance of FALQON on weighted 3-regular graphs, and identified instances where the r_{A} and ϕ convergence is enhanced by introducing one of three possible heuristic modifications to the FALQON algorithm. Details can be found in Appendix A.

In our companion paper [15], we present a sampling complexity comparison between FALQON and QAOA in the context of MaxCut, as quantified by the total number of samples (i.e., circuit repetitions) that are required, denoted N_s. When a gradient algorithm is used for QAOA, $N_s^{QAOA} = \mathcal{O}(mq(\ell)\ell)$, where m denotes the number of samples needed to estimate the expectation value of a two-qubit Pauli string P_j, and for simplicity, m is assumed to be independent of P_j and q denotes the number of classical optimization iterations. For gradient-free methods, $N_s^{QAOA} = \mathcal{O}(mq(\ell))$. Meanwhile, in FALQON we find $N_s^{FALQON} = \mathcal{O}(md\ell)$, where d denotes the degree of the graph. This suggests that FALQON has a more favorable sampling complexity than QAOA for cases where the number of QAOA optimization iterations $q(\ell)$ exceeds $d\ell$ in general, or d when a gradient algorithm is utilized. Further details can be found in [15].

Discussion and outlook.– We have introduced FALQON as a constructive, feedback-based algorithm for solving combinatorial optimization problems using quantum computers. Importantly, FALQON performs optimization without the need for an expensive classical optimization loop. We have demonstrated its performance on current quantum hardware and provided numerical analyses of its performance towards finding the maximum cut of regular graphs. By studying the performance with respect to layer and the problem size n, our numerical analyses show that FALQON converges to very high approximation ratios and success probabilities with a favorable scaling of resources with respect to n, suggesting that FALQON may be a useful heuristic algorithm for this class of problems.

Our findings also suggest that FALQON can require relatively deep circuits in order to achieve this convergence, relative to the shallow circuits typically considered in QAOA. In our companion article [15], we provide an in-depth analysis of the tradeoffs in the performance and resource requirements of FALQON and QAOA, and discuss the resource regimes where each of these methods can be expected to offer advantages. In short, we expect QAOA to be favorable in settings where suitable classical optimization resources are available and quantum resources are restricted to the regime of shallow circuits. Meanwhile, FALQON performs well for deep circuits and does not require any classical optimization resources, meaning that there is no rising classical cost as the quantum circuit depth is increased. This indicates that in settings where deep circuits are feasible, FALQON is a new heuristic that could offer a considerable advantage.

In addition, our companion paper [15] also includes the following other important elements. (1) We present an analysis of convergence criteria for the algorithm. (2) For the analysis presented here we have assumed ideal, noiseless access to the expectation values A_k that dictate the feedback signal, β_k, however, in [15] we show that FALQON is robust to noise in this quantity stemming from finite sample estimates of these expectations. This robustness ultimately stems from the flexibility in choosing β to satisfy Eq. (1). (3) We compare the performance of FALQON and QAOA for a fixed number of circuit repetitions, and we also explore how FALQON can also be used to seed QAOA by identifying a set of initial QAOA parameters that can serve as the starting point for subsequent iterative optimization. We show that this seeding procedure is useful in settings with limited circuit depth, in cases where FALQON fails to converge on its own, and in cases where QAOA fails to converge on its own due to difficulty with effective initialization of the optimization procedure. (4) We numerically demonstrate FALQON on weighted MaxCut, detail some possible extensions to the protocol, and analyze the relationship between FALQON and quantum annealing protocols.

Finally, we note that FALQON can be applied to combinatorial optimization problems beyond MaxCut, *e.g.,* [20], and could have broader implications for quantum variational algorithms. That is, it is possible to develop feedback-based alternatives of variational ansatzë for other applications such as electronic structure or machine learning [21], and these would have the benefit of needing no classical optimization resources, at the cost of requiring measurements whose results condition the feedback.

Acknowledgments.– We gratefully acknowledge discussions with C. Arenz, L. Brady, L. Cincio, T.S. Ho, L. Kocia, O. Parekh, H. Rabitz, and K. Young. MS also thanks M. Fenech. This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under the Quantum Computing Application Teams program. A.B.M. also acknowledges support from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-FG02-97ER25308, as well as support from Sandia National Laboratories Laboratory Directed Research and Development Program. M.D.G. also acknowledges support from the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. SAND2021-9086 O.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

[1] C. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Dover Publications, 1998).

[2] A. Lucas, Ising formulations of many np problems, Frontiers in Physics **2**, 10.3389/fphy.2014.00005 (2014).

[3] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: methods and implementations, Reports on Progress in Physics **83**, 054401 (2020).

[4] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm (2014), arXiv:1411.4028 [quant-ph].

[5] R. Kosloff, A. D. Hammerich, and D. Tannor, Excitation without demolition: Radiative excitation of ground-surface vibration by impulsive stimulated raman scattering with damage control, Phys. Rev. Lett. **69**, 2172 (1992).

[6] M. Sugawara and Y. Fujimura, Control of quantum dynamics by a locally optimized laser field. application to ring puckering isomerization, J. Chem. Phys. **100**, 5646 (1994).

[7] M. Sugawara and Y. Fujimura, Control of quantum dynamics by a locally optimized laser field. multi-photon dissociation of hydrogen fluoride, Chem. Phys. **196**, 113 (1995).

[8] Y. Ohtsuki, Y. Yahata, H. Kono, and Y. Fujimura, Application of a locally optimized control theory to pump-dump laser-driven chemical reactions, Chem. Phys. Lett. **287**, 627 (1998).

[9] D. J. Tannor, R. Kosloff, and A. Bartana, Laser cooling of internal degrees of freedom of molecules by dynamically trapped states, Faraday Discuss. **113**, 365 (1999).

[10] M. Sugawara, General formulation of locally designed coherent control theory for quantum system, J. Chem. Phys. **118**, 6784 (2003).

[11] S. Grivopoulos and B. Bamieh, Lyapunov-based control of quantum systems, in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No.03CH37475)*, Vol. 1 (2003) pp. 434–438 Vol.1.

[12] M. Mirrahimi, G. Turinici, and P. Rouchon, Reference trajectory tracking for locally designed coherent quantum controls, J. Phys. Chem. A **109 11**, 2631 (2005).

[13] V. Engel, C. Meier, and D. J. Tannor, Local control theory: Recent applications to energy and particle transfer processes in molecules, in *Advances in Chemical Physics* (John Wiley & Sons, Ltd, 2009) pp. 29–101.

[14] L. Bittel and M. Kliesch, Training variational quantum algorithms is np-hard, Phys. Rev. Lett. **127**, 120502 (2021).

[15] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Lyapunov control-inspired strategies for quantum combinatorial optimization (2021), arXiv:2108.05945 [quant-ph].

[16] In reference to QLC, $\langle H_\mathrm{p}\rangle$ essentially serves as a control Lyapunov function, though in practice it may not be strictly positive or meet all of the requirements to be true Lyapunov function.

[17] S. Cong and F. Meng, A survey of quantum lyapunov control methods, Sci. World J. **2013** (2013).

[18] Estimating each A_k can be accomplished by first expanding it in the Pauli operator basis as $A_k = \langle\psi_k|i[H_\mathrm{d}, H_\mathrm{p}]|\psi_k\rangle = \Sigma_{j=1}^{N}\alpha_j\langle\psi_k|P_j|\psi_k\rangle$, where α_j are scalar coefficients and P_j are Pauli strings, and then measuring the expectations of each P_j in order to evaluate the weighted sum. The value of N depends on H_p and H_d. For the MaxCut examples we consider, $N \le n(n-1)$.

[19] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, J. ACM **42**, 1115 (1995).

[20] D. Wakeham and J. Ceroni, Feedback-Based Quantum Optimization (FALQON), `https://pennylane.ai/qml/demos/tutorial_falqon.html` (2021), accessed: 2021-05-21.

[21] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, Variational quantum algorithms, Nat. Rev. Phys. **3**, 625 (2021).

[22] R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, Parameter transfer for quantum approximate optimization of weighted maxcut (2022), arXiv:2201.11785 [quant-ph].

[23] K. Beauchard, J. M. Coron, M. Mirrahimi, and P. Rouchon, Implicit lyapunov control of finite dimensional schrödinger equations, Syst. Control. Lett. **56**, 388 (2007).

[24] S. Zhao, H. Lin, J. Sun, and Z. Xue, An implicit lyapunov control for finite-dimensional closed quantum systems, Int. J. Robust Nonlinear Control **22**, 1212 (2012).

APPENDIX A: HEURISTIC IMPROVEMENTS

Our numerical illustrations involving MaxCut on unweighted 3-regular graphs show that FALQON converges to very high approximation ratios and success probabilities. However, we also tested the performance of FALQON on weighted graphs, and were able to identify problem instances where r_A appears to converge to very high values, while the convergence of ϕ is less favorable, i.e., $\beta \to 0$ prior to $\phi \to 1$, indicating that β tends to zero prematurely. Like behavior has been found in numerical studies of QAOA, where the inclusion of edge weights leads to the appearance of many poor-quality local minima in the optimization landscape [22]. To cope with these situations, we introduce three heuristic modifications that can be used to enhance the performance of FALQON.

The first modification is to incorporate random "kicks" into β. For some β_c of our choosing, for all $\beta_k < \beta_c$, with probability P_k we set $\beta_k = \beta_c$. We choose $\beta_c = 1$ and $P_k = (1-\beta_k)\alpha_k$, where $\alpha_k = 0.1\sin^2(\frac{\pi k}{2\ell} - \frac{\pi}{2})$ is designed to decrease to zero as a function of circuit depth.

We also consider a second heuristic inspired by QLC, where the use of a reference perturbation $\lambda(t)$ in the control $\beta(t)$, such that $H(t) = H_\mathrm{p} + (\lambda(t)+\beta(t))H_\mathrm{d}$, has been considered in order to improve convergence [17, 23, 24]. In this setting, we may define System (a) as a system with drift Hamiltonian H_p and control Hamiltonian H_d, and System (b) as the perturbed system with drift Hamiltonian $H_{\mathrm{p},(b)}(t) \equiv H_\mathrm{p} + \lambda(t)H_\mathrm{d}$ and control Hamiltonian H_d. Then, the time-derivative of $\langle\psi(t)|H_{\mathrm{p},2}(t)|\psi(t)\rangle$, allows us to define $\beta(t) = -A(t)$ as usual to ensure $\frac{d}{dt}\langle\psi(t)|H_{\mathrm{p},2}(t)|\psi(t)\rangle \le 0$. Within this framework, if System (b) converges asymptotically to the ground state of $H_{\mathrm{p},2}(t)$, and if $\lambda(t) = 0$ when this occurs, then System (b) becomes System (a), such that the ground state of $H_{\mathrm{p},(b)}(t)$ is also the ground state of H_p, and the method has converged successfully to the desired state. In practice, $\lambda(t)$ can be chosen to be a slowly-varying reference function that tends to 0 as $t \to \infty$. This framework can be translated into a modified version of FALQON by discretizing as before; for our numerical illustrations, we chose $\lambda_k = \alpha_k$.

Then, using this second heuristic as a baseline, we can define a third heuristic that uses an iterative QLC procedure to successively refine β in a manner that is free of any classical optimization [12]. The procedure begins by implementing the standard FALQON framework and obtaining a set of β values for ℓ layers. Then, these initial $\beta = \beta^{(0)}$ values are set as a reference perturbation $\lambda^{(1)}$, and a new set of $\beta^{(1)}$ values are obtained using the second heuristic approach described above. Then, a new reference perturbation is defined as $\lambda^{(2)} = \lambda^{(1)} + \beta^{(1)}$, and the process is repeated. If ℓ is selected to be large enough such that $\beta_\ell = 0$, this iterative procedure guarantees a monotonic improvement of $\langle H_p \rangle$ with respect to iteration. For further details, we refer the reader to our companion paper [15].

To illustrate these heuristic modifications, in Fig. 4 we present the performance of FALQON with and without these modifications when solving a MaxCut problem on a weighted, 4-regular graph with $n = 8$ nodes using $\Delta t = 0.08$, where the edge weights are drawn from a uniform distribution between 0 and 1.

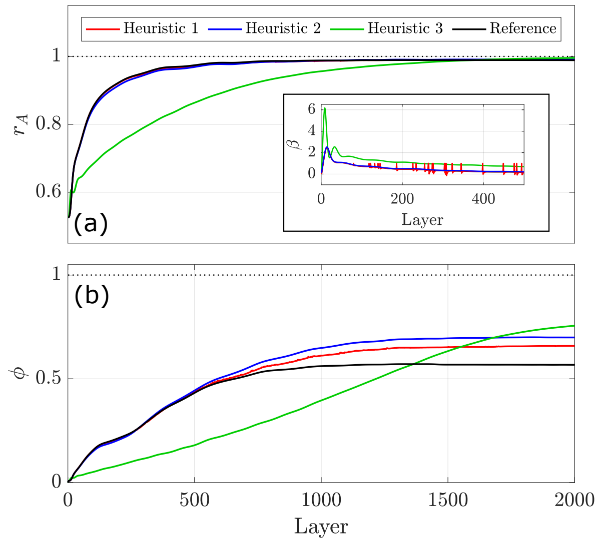

Figure 4. Standard FALQON (black) is compared against the three heuristic modifications, which incorporate random kicks (red), a reference perturbation (blue), and three iterations of the iterative QLC procedure (green), for MaxCut on a weighted, 4-regular graph with $n = 8$ vertices. The r_{A} and ϕ results are shown in (a) and (b), respectively; associated values of β are plotted for the first 400 layers in the inset.